

Université de Montréal

Using Artificial Intelligence to Increase Access to Justice

Par

Hannes Westermann

Faculté de droit, Université de Montréal

Thèse présentée en vue de l'obtention du grade de **doctorat en droit (LL. D.)**,

option innovation, science, technologie et droit

Mars 2023

© Hannes Westermann, 2023

Université de Montréal

Unité académique : Faculté de droit

Cette thèse intitulée

Using Artificial Intelligence to Increase Access to Justice

Présenté par

Hannes Westermann

A été évaluée par un jury composé des personnes suivantes

Prof. Nicolas Vermeys
Président-rapporteur

Prof. Karim Benyekhlef
Directeur de recherche

Prof. Pierre Larouche
Membre du jury

Prof. Kevin Ashley
Examineur externe

Abstract

Artificial intelligence is one of the most thriving and exciting areas in research and industry. Recently, approaches using deep learning have led to a number of breakthroughs in a range of areas, including computer vision, machine translation, image recognition and generation, and text understanding and generation (such as GPT-4).

In this thesis, I investigate if and how artificial intelligence (AI) can be used to improve access to justice and access to legal information for laypeople, i.e. people without legal training. The average citizen is often overwhelmed and helpless when dealing with legal problems. They may struggle to understand how the law applies to their situation, and further have trouble using the judicial system to resolve their issue, even if they are aware of their rights. This results in their problems going unresolved or prevents them from benefiting from opportunities available to them.

For this reason, I developed and implemented the “JusticeBot” methodology, which uses AI to support laypeople with their legal issues. The resulting tools use a hybrid rule-based and case-based reasoning approach to ask a user relevant questions, analyze their legal situation, and provide them with suitable legal information and similar previous cases related to their particular legal problem. They can use this information to negotiate a mutually agreeable solution, or to navigate the arduous legal process. Thus, JusticeBot is an augmented intelligence tool, enhancing the user’s knowledge level to help them solve their legal problems.

I describe the overall methodology and how I implemented it into software tools, e.g. the “JusticeCreator”, an interface to create and update JusticeBot tools. I also elaborate on a deployed JusticeBot tool in the area of landlord-tenant disputes, which is accessible to the public at <https://justicebot.ca>. This tool has been used over 17k times, and 86% of users responding to a survey report that they would recommend the system to others. I believe that JusticeBot can contribute to helping individuals resolve their legal problems, as well as increasing trust in and identification with legal institutions on a societal level, by improving access to justice and access to legal information for the average citizen.

Keywords: Artificial Intelligence, Machine Learning, Expert System, AI & Law, Access to Justice, Access to Legal Information, Legal Decision Support, Augmented Intelligence, Human-Computer Interaction, Interdisciplinary research.

Resumé

L'intelligence artificielle est l'un des domaines les plus florissants et les plus passionnants de la recherche et de l'industrie. Au cours des dernières années, les approches utilisant l'apprentissage profond ont permis de nombreuses avancées dans divers domaines, notamment la vision par ordinateur, la traduction automatique, la reconnaissance et la génération d'images, ainsi que la compréhension et la génération de textes (tels que GPT-4).

Dans cette thèse, je cherche à savoir si et comment l'intelligence artificielle peut être utilisée pour améliorer l'accès à la justice et à l'information juridique pour les justiciables. Le citoyen moyen est souvent dépassé et impuissant lorsqu'il est confronté à des problèmes juridiques. Il peut avoir du mal à comprendre comment la loi s'applique à sa situation et à utiliser le système juridique pour résoudre son problème, même s'il est conscient de ses droits. En conséquence, leurs problèmes restent irrésolus ou ils ne profitent pas des possibilités qui leur sont offertes.

C'est pourquoi j'ai développé et mis en œuvre la méthodologie "JusticeBot", qui utilise l'IA pour aider les justiciables à résoudre leurs problèmes juridiques. Les outils qui en résultent utilisent une approche hybride de raisonnement basée sur des règles et des cas pour poser à l'utilisateur des questions pertinentes, analyser sa situation juridique et lui fournir des informations juridiques appropriées ainsi que des cas antérieurs similaires liés à son problème juridique particulier. L'utilisateur peut utiliser ces informations pour négocier une solution mutuellement acceptable ou pour naviguer dans le processus juridique ardu. JusticeBot est donc un outil d'intelligence augmentée, qui améliore le niveau de connaissance de l'utilisateur pour l'aider à résoudre ses problèmes juridiques.

Je décris la méthodologie globale et la manière dont je l'ai mise en œuvre dans des outils logiciels, par exemple le "JusticeCreator", une interface permettant de créer et de mettre à jour les outils JusticeBot. Je présente également un outil JusticeBot déployé dans le domaine des litiges entre propriétaires et locataires, qui est accessible au public à l'adresse <https://justicebot.ca>. Cet outil a été utilisé plus de 17 000 fois et 86 % des

utilisateurs ayant répondu à une enquête ont déclaré qu'ils recommanderaient le système à d'autres personnes. Je pense que JusticeBot peut contribuer à aider les individus à résoudre leurs problèmes juridiques, ainsi qu'à renforcer la confiance et l'identification aux institutions juridiques au niveau sociétal en améliorant l'accès à la justice et l'accès à l'information juridique pour le citoyen moyen.

Mots-clés : Intelligence artificielle, apprentissage automatique, système expert, IA et droit, accès à la justice, accès à l'information juridique, aide à la décision juridique, intelligence augmentée, interaction homme-machine, recherche interdisciplinaire.

Table of Contents

Abstract	3
Resumé.....	5
Table of Contents	7
List of Tables	11
List of Figures	12
List of Abbreviations	15
Acknowledgements.....	18
Chapter 1 Introduction	20
1.1 Introduction	20
1.2 Research objective & research question.....	23
1.3 Theories and methods.....	29
1.4 Structure	43
Chapter 2 What is Artificial Intelligence?	48
2.1 Introduction	48
2.2 Defining Artificial Intelligence	50
2.3 General Artificial Intelligence vs Narrow Artificial Intelligence	51
2.4 Tasks.....	52
2.5 The Symbolic Approach.....	57
2.6 Machine Learning	65
2.7 Conclusion.....	96
Chapter 3 Access to Justice.....	98
3.1 Introduction	98
3.2 Everyday legal issues and current solutions.....	99

3.3	Access to Justice.....	111
3.4	Access to Legal information	118
3.5	Using technology to increase access to justice.....	123
3.6	The societal impacts of legal decision support tools.....	139
3.7	Conclusion.....	148
Chapter 4 Automating Legal Decision Making		151
4.1	Introduction	151
4.2	Identification of the guiding legal rule.....	156
4.3	Finding of facts.....	161
4.4	Establish fulfilled legal criteria based on facts.....	168
4.5	Applying legal rules to legal criteria	197
4.6	Determining the outcome of a case	210
4.7	Explaining the decision	217
4.8	The interplay between the legislator and judiciary	224
4.9	Reasoning with Case Law	227
4.10	Conclusion.....	237
Chapter 5 Design criteria		241
5.1	Introduction	241
5.2	The system should target laypeople	242
5.3	The system should be able to handle areas of high-volume, low-intensity legal problems.....	250
5.4	The system should give specific and useful information	255
5.5	The system should be practical	268
5.6	Application in building legal decision support tools.....	274
5.7	Conclusion.....	275

Chapter 6 FactorBot – factor-based modelling of decisions.....	277
6.1 Introduction	277
6.2 A User Journey through FactorBot	278
6.3 Building FactorBot.....	281
6.4 Discussion	294
6.5 Does the system fulfill the design criteria?	306
6.6 Conclusion.....	308
Chapter 7 JusticeBot – a hybrid rule-based/case-based augmented intelligence system	311
7.1 Introduction	311
7.2 A user journey through JusticeBot.....	314
7.3 Encoding legal information.....	321
7.4 Obtaining user input.....	370
7.5 Analysis of the case of the user	376
7.6 Providing results to the user.....	380
7.7 Feedback mechanisms.....	381
7.8 Discussion	385
7.9 Conclusion.....	406
Chapter 8 Case study: JusticeBot TAL.....	408
8.1 Introduction	408
8.2 Background	409
8.3 The JusticeBot TAL	426
8.4 Feedback and Analytics	440
8.5 Conclusion.....	450
Chapter 9 Other Areas & Future Work.....	452

9.1	Introduction	452
9.2	Other target users	453
9.3	Other legal and administrative areas	456
9.4	Future Research.....	471
9.5	Conclusion.....	493
Chapter 10 Conclusion.....		494
10.1	Introduction	494
10.2	The JusticeBot methodology	496
10.3	The JusticeBot toolchain	504
10.4	Discussion.....	508
10.5	Limitations.....	520
10.6	Can the JusticeBot methodology increase access to justice and legal information?	531
10.7	How will tools like the JusticeBot impact the legal profession?	535
Bibliography		537
Statutes & Cases		537
Books		538
Articles		541
Reports & Online Sources		559

List of Tables

Table 1 – A list of different tasks that can be accomplished by AI systems, including an input and output.	53
Table 2 - Possible representations of data.....	73
Table 3 - Example of part of dataset with house properties (features) and price (target). 73	
Table 4 - Way of abstracting cases by capturing assessed criteria and reasoning	343
Table 5 - An example of a case abstracted in JusticeBot.....	345
Table 6 - An example of a case abstracted in JusticeBot.....	345
Table 7 - Distribution of landlords suing tenants vs tenants suing landlords	431
Table 8 - Top 5 most frequent claims where the plaintiff was a landlord	432
Table 9 - Top 5 most frequent claims when the plaintiff was a tenant	432
Table 10 - Example of decision template	485

List of Figures

Figure 1 - An example task, determining whether a picture contains an apple or an orange	66
Figure 2 - The data required to train a supervised learning system	69
Figure 3 - Model of simple neural network	78
Figure 4 – DALL-E 2 output for prompt "A robot judge mediates a discussion between two people"	89
Figure 5 - User Input in FactorBot.....	279
Figure 6 - Information output in FactorBot	280
Figure 7 - Partial rule tree for termination of a lease due to the dwelling not being in a "good habitable condition"	282
Figure 8 - Extract of taxonomy used for classifying case factors.....	284
Figure 9 - Taxonomy of in-depth factors regarding bedbugs	285
Figure 10 - Tag frequency per case.....	286
Figure 11 - Most frequent factors	287
Figure 12 - Rent reduction in months versus number of factors.....	292
Figure 13 - Statistics in data based on bedbug taxonomy.....	292
Figure 14 - An explanation of the elements present on a screen of a JusticeBot	315
Figure 15 - Initial screens on JusticeBot.....	316
Figure 16 – Possible paths for landlord and tenants	317
Figure 17 - Path for exploring options after tenant is late with paying rent	318
Figure 18 - Information provided after completed pathway	319
Figure 19 – A question block, representing a legal criterion	325
Figure 20 – An information block, representing a legal conclusion.....	327
Figure 21 - An example of a directed acyclic graph.....	329
Figure 22 - DAG representing a legal criterion and two possible legal conclusions.....	330
Figure 23 - A schema with intermediary conclusions	331
Figure 24 - Cumulative criteria in a DAG	332
Figure 25 - Alternative requirements	333
Figure 26 - Welcome screen of the JusticeCreator	334

Figure 27 - Overview over the JusticeCreator interface	335
Figure 28 - Creating a new block and linking blocks in the JusticeCreator	337
Figure 29 - The legal pathway of the JusticeBot TAL.....	338
Figure 30 - Graph form of relevant legal criteria for determine lease termination due to frequent lateness of rent payment	344
Figure 31 - The reasoning path by a judge taken in Table 5.....	345
Figure 32 - The reasoning path by a judge taken in Table 6.....	346
Figure 33 - Reasoning schema for cases regarding frequent lateness of rent payment ..	351
Figure 34 - Path with intermediary conclusion.....	354
Figure 35 - Reasoning schema for cases regarding frequent lateness of rent payment ..	356
Figure 36 - AI case suggestion view in the JusticeCreator.....	358
Figure 37 - Schema related to lateness of rent of more than three weeks	360
Figure 38 - Encoding case law outcomes in the JusticeCreator.....	361
Figure 39 - Adding case law summaries to legal criteria.....	362
Figure 40 - Initial blocks in the JusticeBot TAL	366
Figure 41 - Reasoning schema regarding lateness of rent	373
Figure 42 - A possible user path through the schema in Figure 41	374
Figure 43 - Hypothetical path through legal reasoning steps	375
Figure 44 - A user case can be matched to previous cases, even if not all answers are answered in the same way.....	378
Figure 45 - Information blocks presented at the end of a pathway.....	380
Figure 46 - Star rating of individual pages	382
Figure 47 - Form to ask user to describe their uncovered issue.....	384
Figure 48 - Word cloud of 10k TAL decisions.....	433
Figure 49 - An overview over the JusticeBot TAL schema.....	436
Figure 50 - Statistics JusticeBot uses from July 2021 to March 2022.....	439
Figure 51 - JusticeBot as featured on the website of the Tribunal Administratif du Logement	441
Figure 52 - Time spent in JusticeBot depending on outcome.....	443
Figure 53 - Top 5 most clicked pathways for tenants.....	444

Figure 54 - Top 5 most clicked pathways for landlords	444
Figure 55 - Star ratings of pages in the JusticeBot	446
Figure 56 - A word cloud of frequent words appearing in the "missing questions" form	448
Figure 57 - A list of possible pathways for the user to explore	473
Figure 58 - Interface for lay language description of issue and path suggestion.....	475
Figure 59 - A mock-up of an interface used to structure evidence using the JusticeBot methodology	484
Figure 60 - Mock-up of part of contract creator	487
Figure 61 - A mock-up of an ODR interface following the structure of a JusticeBot schema.....	491
Figure 62 - reasoning schema for lease termination due to frequent lateness of rent payment.....	497
Figure 63 - A case is annotated and added to the reasoning schema	498
Figure 64 - A user navigates the JusticeBot interface to create a hypothesis for how their case will be treated.....	501
Figure 65 - An example of a user hypothesis matching a previous case	502
Figure 66 - The information a user is shown at the end of their circuit, including information and case law.	503
Figure 67 - A screenshot of the JusticeCreator interface	505
Figure 68 - Screenshots of the JusticeBot frontend	507

List of Abbreviations

ACM	Association for Computing Machinery
ACT	Autonomy through Cyberjustice Technologies and Artificial Intelligence
ADF	Abstract Dialectical Frameworks
ADR	Alternative Dispute Resolution
AI	Artificial Intelligence
ANGELIC	ADF for kNowledGe Encapsulation of Legal Information for Cases
BATNA	Best Alternative to the Negotiated Agreement
BERT	Bidirectional Encoder Representations from Transformers
CABARET	CAse-BAsed REasoning Tool
CAD	Canadian Dollars
CAESAR	Computer-Assisted Enhanced Semantic Annotation & Ranking
CATO	N/A, case-based legal reasoning system using background knowledge
CCQ	Civil Code of Québec
CEO	Chief Executive Officer
CLEO	Community Legal Education Ontario
CNESST	Commission des normes, de l'équité, de la santé et de la sécurité du travail
COLIEE	Competition on Legal Information Extraction/Entailment
COVID-19	COronaVirus Disease of 2019
CQLR	Compilation of Québec Laws and Regulations
CRTC	Canadian Radio-television and Telecommunications Commission
CTV	Canadian Television
DAG	Directed acyclic graph
DALL-E 2	N/A, model used for generating images
ECHR	European Convention on Human Rights
ECtHR	European Court of Human Rights
EU	European Union
EUR-Lex	N/A, online gateway to EU Law
GDPR	General Data Protection Regulation
GOFAI	Good Old-Fashioned Artificial Intelligence
GPT-3	Generative Pre-trained Transformer 3

GREBE	GeneratorR of Exemplar-Based Explanations
HYPO	N/A, legal reasoning system using cases and hypotheticals
SMILE + IBP	SMart Index Learner Plus Issue-Based Prediction
ICAIL	International Conference on Artificial Intelligence and Law
ICT	Information and Communications Technology
IOS	N/A, publishing house
IT	Information Technology
JPES	Justice Pathway Expert System
JSON	JavaScript Object Notation
JUF	Japanese Presupposed Ultimate Fact Theory
JURIX	International Conference on Legal Knowledge and Information Systems
LaMDA	Language Model for Dialogue Applications
LEAP	Legal Education and Awareness Programme
LII	Legal Information Institute
LIST	Legal Issues Taxonomy
ML	Machine Learning
MYCIN	N/A, early medical expert system
NLP	Natural Language Processing
NSRLS	National Self-Represented Litigants Study
NY	New York
ODR	Online Dispute Resolution
ODRAI	the use of artificial intelligence to enhance online dispute resolution
OJ	Official Journal of the European Union
OMHM	Office municipal d'habitation de Montréal
OPC	Office de la protection du consommateur
PARLe	Platform to Aid in the Resolution of Litigation electronically
PC	Personal Computer
PROLEG	PROlog based LEGal reasoning support system
PTSD	Post-traumatic stress disorder
QC	Québec
QDA	Qualitative Data Analysis
RCLALQ	Coalition of Housing Committees and Tenants Associations of Quebec
SCALE	Semi-supervised Case Annotation for Legal Explanations

SCDB	Supreme Court Database
SCHL	Société canadienne d'hypothèques et de logement
SE	Special Edition (iPhone)
SHYSTER	N/A, a case-based legal expert system
SPA	Single Page Application
SSRN	Social Science Research Network
TAL	Tribunal administratif du logement
TAXMAN II	N/A, early legal reasoning system
UK	United Kingdom
UN	United Nations
UNHCR	United Nations High Commissioner for Refugees
URL	Uniform Resource Locator
US/USA	United States of America
USD	United States Dollar
USE	Universal Sentence Encoder
VICTOR	N/A, AI System at Brazilian Supreme Court
VJAP	Value Judgment-based Argumentative Prediction
WYSIWYG	What You See Is What You Get

Acknowledgements

In working on this thesis, I have had the honor and pleasure of working and interacting with a great group of people. Without these people, this work would not have been possible, and I owe them a great debt of gratitude.

First and foremost, I am very grateful to my supervisor, Karim Benyekhlef. Your guidance, support, feedback and advice have made this work what it is today.

I am also grateful to Dirk Westermann and Aurore Troussel for reading and providing feedback on parts of this thesis.

Over the past few years, I have also had the privilege of working with a fantastic team of people at the Cyberjustice Laboratory. Individuals who have worked on the various phases of the JusticeBot project include: Karim Benyekhlef, Nicolas Vermeys, Valentin Callipel, Mark Likhten, André Saintonge, Adam Sampson, Alexandre Lapalme, Andrés Lou, Antoine Louis, Aurore Troussel, Cassandra Larocque-Rigney, Florence So, Gautier Soubie, Ilona Bois-Drivet, Ivan Galindo-Fonseca, Jie Zhu, Justin-Éric Boileau, Karima Smouk, Maria-Fernanda Acevedo, Marie-Laurence Lamarre, Mia Godet, Milene Hauri, Nareg Froundjian, Olivier Salaün, Rémi Slama, Rose Hirigoyen, Samyar Janatian, Sébastien Meeùs, Sophie Crevier, Sylvain Longhais, Tan Jinzhe, and Valentine Cheviron.

I am grateful for the financial support I received for this research, including from the Cyberjustice Laboratory, the LexUM Chair in Legal Information, and the Autonomization through Cyberjustice Technologies and Artificial Intelligence (ACT) project. The first deployed version of the JusticeBot, focused on landlord-tenant disputes, would not have been possible without the excellent collaboration with and support of the Tribunal Administratif du Logement, the help from Aide Juridique de Montreal and Laval, and the financial support of the Ministère de l'économie et de l'innovation du Québec.

Outside of Montreal, I had the pleasure of working with an amazing group of collaborators. The journey would not have been nearly the same without Jaromir Savelka, with whom I have a fantastic collaboration and friendship. I am also very grateful for the

help and support of Kevin Ashley and Vern Walker, who both took me under their wings and introduced me to the AI & Law community.

Finally, I am very grateful for the many helpful conversations with my father, Dirk Westermann, and for the love and support of my family, Dirk, Maria, Anneli and Helena.

Chapter 1 Introduction

1.1 Introduction

1.1.1 Using artificial intelligence in the legal field

Artificial Intelligence is one of the most flourishing and exciting fields of research and industry. In the past few years, approaches using deep neural networks have led to a number of breakthroughs in a diverse set of tasks, including computer vision, machine translation, image generation and natural language processing. These technologies are starting to have enormous impacts on many parts of society, and are predicted to affect a large number of fields in the coming years.¹

One such field is the legal field. Legal reasoning relies on understanding the relevant factors of a particular situation, and then applying rules to the situation to determine possible consequences. A judge, for example, must understand the claims made by the plaintiff, the situation that led to their claim, the rules that regulate this situation, and come to a decision. Lawyers must likewise understand the relevant facts of a situation and advise their client how best to proceed. This type of reasoning often depends on the reading and understanding of rules and previous cases.²

In this thesis, I will investigate how artificial intelligence and machine learning can be used to support or even partially automate these reasoning steps. This investigation will culminate in the creation of a methodology that can be used to create legal decision support tools³ for individuals, by providing legal information. These tools will be able to understand the situation of a user, and provide them with legal information regarding their situation, as well as an overview of court case outcomes in previous similar cases.

¹ See 2.1.

² See Chapter 4.

³ See 1.3.5.3.

1.1.2 Increasing access to justice and access to legal information

A significant use of this methodology will be to develop tools that are able to support laypeople that face every day legal problems, such as housing disputes, consumer disputes, debt issues and employment issues.⁴ Despite millions of individuals being affected by such situations annually, they often have trouble resolving the issues.⁵ Individuals may not be aware that their problem has a legal solution, or how to practically enforce their rights, causing issues to remain unresolved. The individuals that try to use the court system may run into a complex, expensive, time-consuming and frustrating experience, causing them to feel vulnerable and powerless.⁶ These issues are exacerbated for individuals that choose to self-represent.⁷ Globally, it is estimated that 1.5 billion individuals are unable to resolve their legal issues at any time.⁸

This difficulty of dealing with legal problems frequently leads to such problems remaining unresolved or resulting in inequitable outcomes, causing significant harms to the individuals involved, and costs for society at large.⁹ In Canada, unresolved legal issues are estimated to cost society 746 million dollars annually.¹⁰ A lack of knowledge of the law may also prevent individuals from making use of the legal and administrative system to obtain the advantages that they are entitled to. Globally, at least 4.5 billion people are excluded from the opportunities the law provides.¹¹

A tool based on my methodology will be able to support users in such situations. It will ask simple questions to understand the situation of the user, analyze their situation, and provide them with legal information and previous similar cases, if available. Being aware of the legal implications of their situation can help the user in deciding how to deal with

⁴ See 3.2.1.

⁵ See 3.2.2.

⁶ See 3.2.3.

⁷ See 3.2.4.

⁸ *Justice For All - Final Report* (New York: Center on International Cooperation: The Task Force on Justice, 2019) at 35–36 available at <https://www.justice.sdg16.plus/>.

⁹ See 3.2.5.

¹⁰ *Everyday Legal Problems and the Cost of Justice in Canada*, by Trevor Farrow, Cost of Justice Project (Canadian Forum on Civil Justice, 2014) at 1.

¹¹ note 8 at 38.

their problems, by explaining the avenues available to them to resolve their issues. After understanding their situation, the user may choose to hire a lawyer or use the court system to resolve their issue, supported by the knowledge provided by the system. Or, they may use the information to find a mutually agreeable settlement with the other party. Beyond disputes, the system could inform the user of the opportunities available to them (such as social aid etc) and the requirements they need to fulfill in order to obtain these opportunities.

Such a tool could thus provide much needed knowledge and pathways towards resolving the legal problems faced by millions of individuals in society every year. Further, it could help courts by alleviating their heavy caseload. It could also have a positive impact on the rule of law, one of the most important tenets of society, as individuals are provided with a simple and fair way of resolving their legal problems. Further, as people are given an effective way to resolve their disputes and understand the laws, they may gain an increased trust in legal institutions and feeling of belonging in society.

1.1.3 The components of this thesis

In order to build such a methodology, this thesis will cover steps from the conceptual underpinning of legal decision support tools to the concrete development and deployment of such tools. The steps covered thus include:

- The examination of the capability and limitations of artificial intelligence
- The identification and description of current issues with access to justice and access to legal information
- The understanding of the conceptual underpinnings of legal reasoning
- An examination of how artificial intelligence has been used to automate legal reasoning in the field of AI & Law
- The specification of design criteria for a methodology that could be used to improve the issues with access to justice and access to legal information
- The design of a methodology that fits these criteria, using artificial intelligence to increase access to justice and legal information

- The implementation of the methodology in a production-ready software stack, allowing non-technical individuals to create tools using the methodology
- The use of this software stack to create a deployed tool focused on landlord-tenant disputes, intuitive enough to be used by laypeople
- The evaluation of this deployed tool
- A discussion of other possible target domains for building tools using the methodology, and possible expansions of the methodology

1.2 Research objective & research question

1.2.1 Overall research objective

The objective of my research is to design a methodology to create legal decision-support tools in order to increase access to justice and legal information. The objective is not to design a legal decision support tool in a specific domain, but rather to create a general methodology and toolchain that allows the creation of a wealth of legal decision support tools in many different legal domains.

The methodology should allow the creation of AI-powered tools that are able to provide useful information to layperson users that face legal problems. This information should support the user in better understanding their legal position, such as knowing their rights, gaining a more realistic understanding of potential outcomes, and gaining insight into possible next steps to undertake. This understanding could motivate them to settle their dispute in an amicable way, which could increase social harmony and alleviate the heavy caseload of court cases.

In order to accomplish this objective, I will aim to answer the following overall research question:

How can artificial intelligence be used to increase access to justice and access to legal information through the creation of a methodology for developing legal decision support tools?

1.2.2 Subobjectives

The aforementioned research objective will require the understanding of a number of important subobjectives, each with corresponding questions. All of these serve as steps towards achieving the ultimate objective of this thesis, as described above. Here, I will explain these objectives and questions.

1.2.2.1 Understanding artificial intelligence

The first objective in answering my research question is the understanding of artificial intelligence (AI). In order to decide how AI can be used to increase access to justice and access to legal information, it is crucial to understand what is meant by the expression “artificial intelligence”. This includes examining the different techniques that can be used to build artificial intelligence systems (such as the symbolic approach and machine learning) and understanding the limits of current approaches to artificial intelligence. Understanding these specifics is crucial to determine how artificial intelligence can best be used to accomplish the goal of increasing access to justice and access to legal information.

With regards to this objective, I will examine the following research topics:

- What is meant by artificial intelligence?
- Which kind of tasks can AI be used to solve?
- Which concrete methods exist to build artificial intelligence systems?
- What are the promises and limitations of the methods to build AI systems?

1.2.2.2 Understanding the issues of access to justice and legal information

Building a methodology to increase access to justice further requires the understanding of the issue of access to justice and access to legal information. This includes examining the type of legal problems that affect laypeople, and how they are currently being dealt with. It also includes examining what is generally seen to be encompassed by the terms access to justice and access to legal information, and how technology and AI is being used to address these issues. Understanding this area will be very important in making sure that

the designed methodology is well targeted to the issues faced by laypeople and can support them in a useful way.

With regards to this objective, I will explore the following research topics:

- How many individuals are affected by legal problems? What kind of legal problems are most prevalent?
- How do individuals seek to resolve these problems? How successful are they?
- What is the experience of individuals in interacting with the court system? Does it offer an effective way for them to resolve their problems?
- What are the consequences of unresolved legal problems?
- What is access to justice?
- What is access to legal information?
- How has technology and AI been used to address the issues with access to justice and access to legal information?
- What is the legal status of using software tools to provide legal information?
- What are the potential positive and negative effects of people using AI tools that inform them of their rights?

1.2.2.3 Understanding legal reasoning and automating legal reasoning

Building a tool that uses artificial intelligence to improve access to justice and legal information is likely to involve an understanding of the process of legal reasoning. My methodology is intended to design tools that are able to give information to the user based on their specific situation, as well as informing them of possible outcomes based upon similar cases. This process will likely have to replicate some steps of how judges or other legal decision makers tackle such questions, i.e. legal reasoning. Understanding how judges and jurists perform legal reasoning will be very important in order to design a methodology that is able to partially replicate this reasoning digitally.

Further, it is important to understand previous approaches to automating legal reasoning. In the field of artificial intelligence and law, researchers have examined this question for over 30 years.¹² Understanding some of the approaches taken by these researchers, and the different tradeoffs of these approaches, is crucial in informing the creation of my own methodology.

To this end, I will seek to understand the following research topics:

- Which steps are involved in legal reasoning?
- Which approaches have researchers previously taken to automate these steps?
- What are the different trade-offs of the approaches previously taken by researchers with regards to my use-case of increasing access to justice?
- How are cases reasoned with in different jurisdictions?

1.2.2.4 Determining the relevant design criteria

Based on the previous subobjectives, at this point I will have a good understanding of the capabilities of artificial intelligence, the issues of access to justice and legal information, and the steps of legal reasoning and how they have previously been automated. This will allow me to determine a few relevant design criteria, that will shape the design of my methodology aiming to increase access to justice. Some of these criteria will stem from fundamental considerations, such as making sure that the methodology is able to address the specified issues and support the envisioned target group with useful information. Other criteria will stem from practical considerations, such as making sure that the methodology can have the greatest possible impact. These criteria will guide the development of my methodology.

I will explore the following research topics in order to achieve this objective:

- Which design criteria should guide the design of my methodology using AI to increase access to justice and access to legal information?

¹² Serena Villata et al, “Thirty years of artificial intelligence and law: the third decade” (2022) *Artif Intell Law*, online: <<https://doi.org/10.1007/s10506-022-09327-6>>.

1.2.2.5 Designing a methodology for creating legal decision support tools

Once I have understood the background and design criteria, the next steps are to practically design a methodology that fulfills these criteria, and can be used to create legal decision support tools that can increase access to justice and legal information. Achieving this objective will be the most important contribution presented in this thesis.

In order to create such a methodology, I will investigate and implement the required steps of such a system. This includes encoding legal information, acquiring information corresponding to their situation from the user, analyzing their situation and providing them with information about their specific situation.

In designing such a methodology, I will explore the following research topics:

- How should legislation, court cases and legal information be encoded in order to arrive at a useful result to increase access to justice?
- How can an interface accurately capture the features of a user's potential dispute?
- How can the information provided by the user be analyzed in order to identify relevant information and relevant previous cases?
- How can information be shown to a user of this system in a way that supports them and encourages the amicable settlement of their dispute?
- How can the accuracy of the system be evaluated, and potential sources of bias be eliminated?

1.2.2.6 Implementing the resulting methodology

Once the methodology has been designed, an important part of my research will be to practically design and implement the tools required to use this methodology. This will take the methodology from a theoretical methodology to a practically implemented toolchain, that can be used to create legal decision support tools. This implementation is important in order to demonstrate the feasibility of developing tools based upon the methodology. It is also necessary in order to enable the methodology to be used in practice, enabling it to increase access to justice and legal information.

Thus, the following research topics will be explored:

- How can the methodology be implemented in a concrete, production-ready software stack allowing for the creation of legal decision support tools?

- How can the use of the methodology to create decision support tools be made as effective as possible, and require as little technical knowledge as possible?
- How can a front-end interface be designed for the methodology, that allows a layperson user to interact with the system in a useful way?

1.2.2.7 Validating the resulting methodology

Once the methodology has been created and implemented, it is important to understand how well it works. At the Cyberjustice Laboratory at Université de Montréal, we have used the methodology to build an implemented version of such a tool, focused on landlord-tenant disputes. This tool will serve as a case study to validate the created methodology, and to understand whether the research objective has been achieved.

In order to achieve this objective, I will examine the following research topics:

- Does the methodology allow the implementation of legal decision support tools?
- Does the created legal decision support tool increase access to justice and legal information in an area?
- How was the user experience of individuals interacting with the system?

1.2.2.8 Discussing other application areas and future work

A final important objective for my thesis is to discuss a few other areas where the methodology could be useful and lead to increased access to justice. I will look at both public law and administrative areas as well as areas of legal disputes. Applying the methodology in as many areas as possible is important to ensure the largest possible impact. I will also describe some further improvements that may make the methodology even more powerful, and thus increase the impact on access to justice.

Therefore, I will investigate the following research topics:

- Which areas of public law, administrative procedures or legal disputes may benefit from implementing legal decision support tools using the methodology?
- How can the methodology be further improved and expanded?

1.3 Theories and methods

In tackling the aforementioned research question, there are a number of important methods that I will need. In this section, I will give an overview of some theories and methods that I will use in this pursuit. My research takes an interdisciplinary approach, using methods from both the legal field and computer-science adjacent fields. I will rely on a realist understanding of the law, with some positivist elements. The methods I use will borrow from the field of empirical legal research, AI & Law and Human-Computer Interaction.

1.3.1 An Interdisciplinary approach

The research will use an interdisciplinary approach, by combining legal methods with methods from other fields. Siems has created a taxonomy for separating different types of interdisciplinary research. He refers to “basic interdisciplinary research” as a form of research that starts with a traditional legal question, but also considers other academic fields to find a response.¹³

Siems then goes on to identify three types of “advanced interdisciplinary research”. Type I of advanced interdisciplinary research is research that seeks to answer a non-legal question but uses law as a part of the response to this. Siems mentions the example of investigating measures to tackle climate change. Different legal regimes might be part of the answer to this question, but there are several other responses that are not linked to law, such as technological innovation.¹⁴

Type II of advanced interdisciplinary legal research instead tackles legal questions using a scientific method, by Siems defined as “constructing models and testing hypotheses”. As an example of this, he mentions Law and Economics, which aims to use economical tools to find answers to legal questions.¹⁵

¹³ Mathias M Siems, “The Taxonomy of Interdisciplinary Legal Research: Finding the Way out of the Desert” (2009) 7:1 *Journal of Commonwealth Law and Legal Education* 5–17 at 5–6.

¹⁴ *Ibid* at 8–9.

¹⁵ *Ibid* at 10.

Type III of advanced interdisciplinary legal research uses these same scientific methods to investigate a question of non-legal character. For example, the question of what creates strong capital markets could be investigated using quantitative methods, to discover the interaction between laws and other economic factors.¹⁶

1.3.1.1 Application in my research

In my research, I intend to construct a methodology that allows the creation of legal decision-support tools to increase access to justice and legal information. According to the taxonomy presented above, this can be seen as Type III of advanced interdisciplinary research. Increasing access to justice and access to legal information are unlikely to be seen as pure legal questions. As I will discuss in Chapter 3, they are also linked to societal issues regarding public understanding of the laws that apply to them, the occurrence of certain problems in society and the use of different ways of resolving legal issues. In this research, I will design a methodology that can improve the situation by increasing access to justice.

From the perspective of the user, the research has attributes of Type II of advanced interdisciplinary legal research. The user is able to use a tool created by the methodology to be supported in understanding a legal issue. For example, they may answer a few questions and be provided with legal information and a list of relevant cases with outcomes. Here, a non-legal method (i.e. the empirical discovery of relevant cases) is used to answer a question with legal character (i.e. the potential legal outcomes in court for an individual, based upon a factual situation).

However, it should be noted that the answering of the legal question posed by the user is not the focus of my research. Rather, it is the development of a *methodology*, that can be used to create computer systems that are able to support the user in conducting their own interdisciplinary research.

¹⁶ *Ibid* at 11.

1.3.2 Legal philosophy

Let us briefly talk about the legal philosophy views represented in this thesis. Fully defining these ideas is beyond the scope of this work, but a brief overview will help me orient the legal conceptual underpinnings of my research.

1.3.2.1 Formalism vs realism

Formalism and realism aim to answer the question of how legal decisions are made.

Legal formalism refers to judges deductively applying the law to arrive at an outcome. Under this view, there is no discretion on the part of the judge – they merely apply the law to the facts of a case. The decision can therefore be right or wrong, as the outcome stems from the logic inherent in the law.¹⁷

Realism is opposed to this view. Legal realists believe that cases are not decided only based on the legal rules, but that judges are also influenced by other factors. Such factors may include, e.g., the outcome promoting public welfare.¹⁸ Some realists believe that idiosyncrasies of the judge lead to a decision, making prediction impossible.¹⁹ However, according to Leiter, the majority of realists believe that the underlying factual scenarios of a situation determine the outcomes.²⁰ The rules themselves are thus not enough to determine what a judge will decide (they are *indeterminate* as to the outcome of a case, especially in difficult cases). Rather, judicial behavior needs to be studied as a social science.²¹

Oliver Wendell Holmes, a famous legal realist, observed:

¹⁷ Richard A Posner, “Legal formalism, legal realism, and the interpretation of statutes and the constitution” (1986) 37 Case W Res L Rev 179 at 181.

¹⁸ *Ibid.*

¹⁹ Brian Leiter, “Legal Realism and Legal Positivism Reconsidered” (2001) 111:2 Ethics 278–301 at 281.

²⁰ *Ibid.*

²¹ Andrei Marmor & Alexander Sarch, “The Nature of Law” in Edward N Zalta, ed, *The Stanford Encyclopedia of Philosophy*, fall 2019 ed (Metaphysics Research Lab, Stanford University, 2019); Leslie Green, “Positivism, Realism, and Sources of Law” (2019), online: <<https://papers.ssrn.com/abstract=3443743>> section 1.2.

“The prophecies of what the courts will do in fact, and nothing more pretentious, are what I mean by law.”²²

Thus, he believed that the only thing that matters in the field of law is legal prediction, i.e. predicting how the court will respond to certain situations. A subgroup of the realists, known as “rule skeptics”, questioned whether legislation and case law are relevant at all.²³ They argued that the important thing to understand to predict legal decision-making were the “real rules” that could be discovered from patterns in judicial behavior.²⁴ Susskind concurs with this view, arguing that lawyers tend to think of how judges might see a client’s situation in order to determine their rights and obligations, rather than seeking to apply legal rules.²⁵

1.3.2.1.1 Application in my research

In this research, I lean toward the realist view. I aim to build a system that gives a user useful legal information about their situation, to support them in their decision making. In doing so, I believe it is more helpful to provide an individual with information about how their situation might *in reality* be treated in court. This should be more useful to the user than information of how their case *should* be treated in court, according to the legislation as interpreted by the creator of the system. For example, informing a layperson that they should have the right to a certain remedy is not helpful if the court applies the rule differently, and the person therefore in reality will not be able to obtain this remedy.

At the same time, I do study the importance of legal rules in an area. I will establish that the most helpful way to compare cases is to organize them by the legal criteria that were applied in a certain case. Some of these criteria will stem from rules, others from case law. A main feature of my developed methodology is thus giving the user information

²² Richard E Susskind, *Online courts and the future of justice* (Oxford: Oxford University Press, 2019) at 284.

²³ Leiter, *supra* note 19 at 289.

²⁴ Susskind, *supra* note 22 at 285; Green, *supra* note 21 at 11, 14.

²⁵ Susskind, *supra* note 22 at 285.

regarding the outcomes of cases that were similar to their situation, in terms of the rules applied by the judge.

1.3.2.2 Natural law vs positivism

Another important debate is between naturalists and positivists, regarding the validity of laws. Naturalists believe that the law necessarily has a connection to morals and is therefore only valid if it is also moral. Phrased differently, they believe that the moral content of norms form “part of the conditions of legal validity.”²⁶

Positivists, on the other hand, believe that morals and law are wholly different considerations. Laws are created by considerations of social fact, such as “deliberating, deciding, ordering, tolerating, conforming or obeying” (Social Thesis).²⁷ Laws therefore do not have to be moral to be valid as laws (Separability Thesis).²⁸ According to Green, positivism is more concerned with what constitutes law than which influence these rules have on legal decisions, and often see an important role for judicial discretion.²⁹

1.3.2.2.1 Application in my research

In this thesis, I take a positivist view on the law. I do not attempt to assess the moral way to render a decision. Instead, I treat the law as man-made, and build a system that can help the user understand their situation in light of rules and the discretion exercised by judges.

1.3.3 Empirical Legal Research

An important method employed by legal realists is empirical legal research.³⁰ This method focuses on the analysis of systematically collected data in order to answer a

²⁶ Marmor & Sarch, *supra* note 21.

²⁷ Green, *supra* note 21 at 2–3; Leiter, *supra* note 19 at 286.

²⁸ Marmor & Sarch, *supra* note 21 section 1.1; Leiter, *supra* note 19 at 286.

²⁹ Green, *supra* note 21 at 8–9.

³⁰ Herbert M Kritzer, “Empirical legal studies before 1940: a bibliographic essay” (2009) 6:4 *Journal of Empirical Legal Studies* 925–968 at 879.

question.³¹ I will describe this methodology below, since it is very relevant to my research.

However, as previously stated, it should be noted that while many of the challenges my research faces are similar to those of empirical legal research scholars, the main goal of my research is not to learn more about a legal dispute using empirical methods. Rather, it is to devise a methodology to create an AI system that supports the user in conducting empirical queries into their question. It could thus be categorized as meta-empirical legal research.

1.3.3.1 The method

Empirical Research focuses on the analysis of systematically collected data.³² In the broadest meaning of the term, it is thus any research that is based on observation or experience.³³ This data can be Quantitative, in the form of numbers, or Qualitative, non-numerical.³⁴ The data that is analyzed can be almost anything that is related to the real world: Legislation, case law, interviews or survey data.³⁵

While traditional legal analysis is concerned mostly with the analysis of a single case (such as finding an argument or reaching a legal decision), Empirical Legal Research focuses more on aggregate effects.³⁶ Empirical research is also not as concerned with reaching a final conclusion – rather, it often investigates a single issue over a span of tens of years. It is thus an ongoing inquiry into a legal issue.³⁷

The empirical method has origins in many other fields, such as sociology, criminology, psychology, “Law and Economics” and “Law and Anthropology”. It heavily borrows in

³¹ Robert M Lawless, Jennifer K Robbennolt & Thomas S Ulen, *Empirical Methods in Law*, 2nd ed (New York: Wolters Kluwer Law & Business, 2016) at 5.

³² *Ibid.*

³³ Lee Epstein & Gary King, “The Rules of Inference” (2002) 69 *University of Chicago Law Review* 1–209 at 2.

³⁴ *Ibid.*; Lawless, Robbennolt & Ulen, *supra* note 31 at 13.

³⁵ Epstein & King, *supra* note 33 at 2.

³⁶ Lawless, Robbennolt & Ulen, *supra* note 31 at 10–11.

³⁷ *Ibid* at 11–14.

methodology from these areas.³⁸ However, some challenges faced by Empirical Legal scholars are unique to the legal field. Most notably, empirical legal research targets legal professionals, who might not have a statistical background, making communication of the results a difficult task.³⁹

1.3.3.2 Types of Empirical Legal Research

As mentioned above, empirical research is typically split into two families of methods: Quantitative research and Qualitative research.

Quantitative legal research is the use of statistical and other Quantitative methods in the evaluation of the collected data.⁴⁰ It is the dominant form of Empirical Legal Research.⁴¹

The other branch of Empirical Legal Research is *Qualitative*. Qualitative Research in general is defined as observing people interact with their natural environment. As opposed to a Quantitative analysis, it measures not a quantified result, but rather the existence or absence of a social fact. This social fact first must be identified and defined through Qualitative research. Only then can it be measured.⁴²

A method used in Qualitative Empirical Legal Research is *The Grounded Theory method*. It allows the researcher to follow the “natural pattern of human inquiry” by developing a theory progressively as more is learnt about an area, rather than testing a predefined hypothesis. The method is applied by reading qualitative documents, such as legal cases and statutes, and highlighting interesting sections to come up with repeatable concepts. As the researcher goes through such documents, they refine and adapt these concepts to

³⁸ Franciscus L Leeuw & Hans Schmeets, *Empirical legal research: a guidance book for lawyers, legislators and regulators* (Cheltenham, UK Northampton, MA: Edward Elgar Publishing, 2016) at 20–36.

³⁹ Lee Epstein & Andrew D Martin, “Quantitative Approaches to Empirical Legal Research” (2010) *The Oxford Handbook of Empirical Legal Research*, online: <<http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199542475.001.0001/oxfordhb-9780199542475-e-38>> at 902.

⁴⁰ Peter Cane & Herbert M Kritzer, eds, *The Oxford handbook of empirical legal research*, 1st ed, Oxford handbooks in law (Oxford ; New York: Oxford University Press, 2012) at 3.

⁴¹ *Ibid* at 2.

⁴² Lisa Webley, “Qualitative approaches to empirical legal research” (2010) *The Oxford handbook of empirical legal research* 926–950 at 927–928.

correspond to the information contained in the documents. In the next stage, the links between these concepts are described and coded. Finally, this information is used to develop a theory or conclusion.⁴³

1.3.3.3 Application in my research

In my research, I will rely on methods from both the qualitative and quantitative branches of Empirical Legal Research. In the FactorBot research (Chapter 6), the grounded theory method will be very useful as a method to identify relevant factual occurrences from case law. It is difficult to determine a taxonomy to annotate cases (for example, in terms of facts or legal criteria that are applied) before having read a number of cases. The literature around the grounded theory method will therefore be very helpful in identifying and capturing relevant phenomenon in case law. Further, once a number of cases has been captured and annotated, the quantitative methodology will be helpful in analyzing the results. In the JusticeBot research (Chapter 7), the grounded theory method is used to discover the legal criteria that judges apply in deciding on cases in certain areas, in order to create a schema of possible reasoning pathways.

1.3.4 Artificial Intelligence and Law

Another field that is highly relevant for my research is the field of Artificial Intelligence and Law. It is the field concerned with the study of how to use new methods from computer science fields to understand and automate legal processes.⁴⁴ Active since the 1980s, there are several subbranches and directions of research in how to accomplish this task.⁴⁵ I am especially interested in the branches focusing on supporting laypeople and pro-se litigants, and the branch focusing on analyzing statutes and case law to understand and predict the outcome of legal cases.

⁴³ *Ibid* at 943–945.

⁴⁴ Kevin D Ashley, *Artificial intelligence and legal analytics: new tools for law practice in the digital age* (Cambridge New York Melbourne Delhi Singapore: Cambridge Univ Press, 2017) at 4–6.

⁴⁵ *Ibid* at 4.

Early in the history of the field, there were two main approaches to automating legal reasoning: Case-based reasoning and rule-based reasoning.⁴⁶

Case-based reasoning relies on performing legal reasoning using examples of previous cases. Such systems are able to take a new case and compare it to a database of previous cases. The systems can then use this comparison to generate arguments for how the new case should be decided, or predict the outcome of the case, as well as explaining those predictions by referring to previous cases.⁴⁷

Rule-based reasoning, on the other hand, relies on logically encoded rules to decide new cases. In this way, entire statutes can be encoded, and the user can be asked for the facts of a situation. Based on the encoded laws, the system can tell them the outcome based on the law, and generate an explanation.⁴⁸ This approach can be used to build so-called expert systems.⁴⁹

Of course, these are not the only two approaches that are used. Some approaches combine the two styles of reasoning in hybrid systems.⁵⁰ Other approaches incorporate argument schemes⁵¹ or value judgments.⁵²

⁴⁶ Katie Atkinson, Trevor Bench-Capon & Danushka Bollegala, “Explanation in AI and law: Past, present and future” (2020) 289 *Artificial Intelligence* 103387 at 3; Trevor Bench-Capon et al, “A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law” (2012) 20:3 *Artificial Intelligence and Law* 215–319 at 6.

⁴⁷ See e.g. Edwina L Rissland, “Examples in Legal Reasoning: Legal Hypotheticals.” (1983) *IJCAI* 90–93; Edwina L Rissland & Kevin D Ashley, “HYPO: A Precedent-Based Legal Reasoner” (1987) *Defense Technical Information Center*, online: <<http://www.dtic.mil/docs/citations/ADA249335>>; L Thorne McCarty, “An implementation of Eisner v. Macomber” (1995) *Proceedings of the fifth international conference on Artificial intelligence and law - ICAIL '95* 276–286; Vincent Aleven, “Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment” (2003) 150:1 *Artificial Intelligence (AI and Law)* 183–237.

⁴⁸ Bench-Capon et al, “A history of AI and Law in 50 papers”, *supra* note 46 at 3; See e.g. D A Waterman & M A Peterson, “Rule-Based Models of Legal Expertise” (1980) 1 *AAAI* 272–275; M J Sergot et al, “The British Nationality Act as a logic program” (1986) 29:5 *Commun ACM* 370–386; Ken Satoh et al, “PROLEG: An Implementation of the Presupposed Ultimate Fact Theory of Japanese Civil Code by PROLOG Technology” (2011) *New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science)* 153–164.

⁴⁹ See 2.5.1.

⁵⁰ Edwina L Rissland & David B Skalak, “CABARET: rule interpretation in a hybrid architecture” (1991) 34:6 *International Journal of Man-Machine Studies (AI and Legal Reasoning. Part 1)* 839–887; Kevin D

More recently, machine learning has found significant use in the field.⁵³ Here, a model is automatically built from a number of datapoints, often case law. Machine learning can be used, for example, to make the case-based approach more effective, by automatically extracting factors from relevant cases.⁵⁴ Prof Kevin Ashley, one of the pioneers of the field of Artificial Intelligence and Law, believes that connecting the traditional methods to textual retrieval methods will lead to a revolution in the field.⁵⁵ Machine learning can also be used to predict the outcomes of cases from the text of a decision.⁵⁶

1.3.4.1 Application in my research

My research is very much part of the field of AI & Law. It shares the goal of building a tool that can automate legal processes, in my case in order to support laypeople and increase their access to justice and access to legal information.

As for the more specific methodology, my research will combine the approaches listed above. My system will use a rule-based representation to encode the legal rules that are applicable in a legal area and provide legal information to the user. However, cases also have a number of important uses. They are used as a source to discover the rules in a legal area, in order to correctly identify the practically relevant rules. Likewise, cases are used to illustrate how judges apply specific legal criteria. Finally, cases are shown to the user to illustrate the outcomes judges have previously awarded in situation similar to that

Ashley & Stefanie Brüninghaus, “Automatically Classifying Case Texts and Predicting Outcomes” (2009) 17:2 *Artif Intell Law* 125–165.

⁵¹ Bench-Capon et al, “A history of AI and Law in 50 papers”, *supra* note 46 at 9–10.

⁵² Matthias Grabmair, *Modeling purposive legal argumentation and case outcome prediction using argument schemes in the value judgment formalism*, PhD Thesis (University of Pittsburgh, 2016).

⁵³ Villata et al, “Thirty years of artificial intelligence and law”, *supra* note 12 at 1.

⁵⁴ Ashley & Brüninghaus, *supra* note 50; Mohammad Hassan Falakmasir & Kevin D Ashley, “Utilizing Vector Space Models for Identifying Legal Factors from Text” (2017) Proceedings of the International Conference on Legal Knowledge and Information Systems (JURIX); L Karl Branting et al, “Scalable and explainable legal prediction” (2021) 29:2 *Artif Intell Law* 213–238.

⁵⁵ Ashley, *supra* note 44 at 3.

⁵⁶ Nikolaos Aletras et al, “Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective” (2016) 2 *PeerJ Comput Sci* e93; Ilias Chalkidis, Ion Androutsopoulos & Nikolaos Aletras, *Neural Legal Judgment Prediction in English* (2019) arXiv:1906.02059 [cs]; Masha Medvedeva, Michel Vols & Martijn Wieling, “Using machine learning to predict decisions of the European Court of Human Rights” (2020) 28:2 *Artif Intell Law* 237–266.

of the user. Machine learning is also used in my methodology, to identify cases that may be relevant to annotate.

1.3.5 Human Computer Interaction

In researching the creation of the interface, I will look at theories from the field of Human-Computer Interaction. It is “the study of how people interact with computing technology.”⁵⁷ The field focuses on the fact that computer systems are only useful if humans are able to interact with them effectively.⁵⁸

1.3.5.1 Interaction design

Human-Computer interaction focuses on two important aspects of computer systems, in order to design the way that users interact with these systems. The *functionality* of a system is defined by the “set of actions or services that it provides to its users.”⁵⁹ The *usability* of a system is “the range and degree by which the system can be used efficiently and adequately to accomplish certain goals for certain users.”⁶⁰ Only if both of these are balanced can an effective system be built.⁶¹

In designing ways for a user to interact with a system, multiple levels need to be considered. These include the *physical* level (i.e. how a user mechanically interacts with a system), the *cognitive* level (i.e. how a user understands and interacts with the system), and the *affective* level (i.e. how the user feels about the interaction).⁶²

These thoughts will be very important in my research. In order to build a system that will see significant use, it is important to design interfaces that closely correspond to the needs of the layperson user. The need for good interfaces further extends to the creation of such systems – if the methodology to create new decision support tools is intuitive and

⁵⁷ Gary M Olson & Judith S Olson, “Human-Computer Interaction: Psychological Aspects of the Human Use of Computing” (2003) 54:1 Annual Review of Psychology 491–516.

⁵⁸ Fakhreddine Karray et al, “Human-Computer Interaction: Overview on State of the Art” (2008) 1:1 International Journal on Smart Sensing and Intelligent Systems 137–159 at 138.

⁵⁹ *Ibid.*

⁶⁰ *Ibid.*

⁶¹ *Ibid.*

⁶² *Ibid* at 139.

does not rely on a technical background, people are much more likely to build such systems, thereby enhancing the impact of the methodology to increase access to justice.

1.3.5.2 Augmented intelligence

One of the oldest streams in the field of human-computer interaction is the idea that computers can work together with a human, to enhance the human experience. Such visions were already described in 1945, in Vannevar Bush's prophetic essay "as we may think", which described the user of computers to navigate the enormous amounts of information available to humans.⁶³

A similarly prophetic essay was produced by Licklider in 1960. In a paper entitled "Man-Computer Symbiosis", he suggests that machines should be developed to enable a cooperative interaction between humans and computers. People will "set the goals, formulate hypotheses, determine the criteria, and perform the evaluations".⁶⁴ Computers, on the other hand, will perform the routine tasks required to prepare the insights of the human.⁶⁵ In doing so, humans will fill in the gaps of the computer program, both in the determination of a problem and a solution.⁶⁶

A related notion is that of decision-support systems. These systems aim to support the user in reaching decisions, by supplementing the knowledge of humans with that of computers.⁶⁷ Typically targeted at users in management of organizations, the tools aim to provide the user with access to the data and model relevant to a specific decision.⁶⁸ A key focus of the tools is to expose interactive quantitative methods to users without computer

⁶³ Vannevar Bush, "As we may think" (1945) 176:1 *The Atlantic Monthly* 101–108; I Scott MacKenzie, *Human-computer interaction: an empirical research perspective*, 1st ed (Amsterdam: Morgan Kaufmann is an imprint of Elsevier, 2013) at 3–5.

⁶⁴ J C R Licklider, "Man-Computer Symbiosis" (1960) HFE-1:1 *IRE Transactions on Human Factors in Electronics* 4–11 at 4.

⁶⁵ *Ibid.*

⁶⁶ *Ibid.* at 7.

⁶⁷ John Zeleznikow, "Using Web-Based Legal Decision Support Systems to Improve Access to Justice" (2002) 11:1 *Info & Comm Tech L* 15–34 at 17.

⁶⁸ F Nelson Ford, "Decision support systems and expert systems: A comparison" (1985) 8:1 *Information & Management* 21–26 at 24.

knowledge.⁶⁹ Decision support systems typically target decisions that are relatively unstructured – a key feature of the systems is thus to allow the user to iteratively enhance their understanding of the problem, and update the system to correspond to this new understanding.⁷⁰ In the literature, a similar notion that is discussed is that of advisory systems, which aim to support and guide a decision maker, while leaving the user as the final decision-maker.⁷¹

I believe these insights to be highly relevant for my research. As we will see in Chapter 3, humans are frequently overwhelmed with the legal information that is presented to them. My methodology will aim to overcome this issue by selecting the information that is currently relevant to the user and helping them determine their questions and goals if they have trouble doing so themselves. At the same time, as we will see in Chapter 2, computers lack some important facilities with regards to certain problems. Finding a way to combine the intelligence of the user and the system is therefore a crucial step in building my methodology.

1.3.5.3 Human-computer interaction and the law

Human-computer interaction reflections have started to appear in the field of law.

Vermeys and Benyekhlef argue that designing cyberjustice technologies is not a neutral process, but that the properties that are encoded into a technology can have impacts on human behavior. They argue that we must carefully study how these technologies change the habits of the users of the legal system, and whether these changes are desirable. An

⁶⁹ *Ibid* at 22; Ralph H Sprague, “A Framework for the Development of Decision Support Systems” (1980) 4:4 MIS Quarterly 1–26 at 2–11.

⁷⁰ Sprague, *supra* note 69 at 2–11; Göran Fick & Ralph H Sprague, *Decision Support Systems: Issues and Challenges: Proceedings of an International Task Force Meeting June 23-25, 1980* (Elsevier, 2013) at 28–29 Google-Books-ID: LF0hBQAAQBAJ.

⁷¹ Brandon A Beemer & Dawn G Gregg, “Advisory Systems to Support Decision Making” in Frada Burstein & Clyde W Holsapple, eds, *Handbook on Decision Support Systems I: Basic Themes* International Handbooks Information System (Berlin, Heidelberg: Springer, 2008) 511 at 511–512.

important part of this study is examining why legal processes work in a certain way in the first place, before improving these processes with technology.⁷²

Margaret Hagan discusses the idea of “Legal Design”, i.e. “the application of human-centered design to the world of law, to make legal systems and services more human-centered, usable, and satisfying.”⁷³ She encourages design thinking to create processes that are usable, useful and engaging, by combining notions from design, tech and law.⁷⁴

Thompson also discusses the importance of understanding and following human interface design patterns in creating legal decision aid tools. He argues that the user-friendly knowledge acquisitions of such tools are crucial in making sure that they are continuously updated and improved.⁷⁵ Further, he argues that it is important to consider the social or emotional needs of users in building expert systems. He discusses the inclusion of techniques from affective computing in order to inquire about the feelings of the user and react differently depending on these feelings.⁷⁶

John Zelenznikow discusses the use of decision support systems in the legal domain. He argues that such systems can increase the consistency, transparency and efficiency of decision making, while also pushing individuals towards settling their cases.⁷⁷ Karl Branting discusses the implementation of advisory systems for pro-se litigants, in order to inform the user of legal avenues, determine whether they may have access to certain relief, procedural requirements for that relief and assisting the user in drafting the required documents.⁷⁸

⁷² Nicolas W Vermeys & Karim Benyekhlef, “Best Practices in the Field of Cyberjustice” (2011) Seminar on Recent Trends and Good Practices in the Application of Electronic Technology to Judicial Processes (E-Justice) at 3–5.

⁷³ Margaret Hagan, “What is Legal Design?”, (26 January 2015), online: *Law By Design* <<https://lawbydesign.co/legal-design/>>.

⁷⁴ *Ibid.*

⁷⁵ Darin Thompson, “Creating New Pathways to Justice Using Simple Artificial Intelligence and Online Dispute Resolution” (2015) 2 *IJODR* 4–53 at 42.

⁷⁶ *Ibid* at 43–51.

⁷⁷ Zeleznikow, *supra* note 67 at 17.

⁷⁸ L Karl Branting, “Advisory systems for pro se litigants” (2001) Proceedings of the 8th international conference on Artificial intelligence and law 139–146 at 3.

Considering these questions will be very important in my research. A key part of my research will be how to build tools that are as easy and pleasant to use as possible, and also understanding the possible impact they can have on the user and legal processes.

I have adopted the terminology of “legal decision support tools”. While these tools typically target management, in this research the envisioned system targets layperson users. When laypeople are faced with a certain situation or need, they must, in essence, become the manager of their own situation, and decide how to proceed. This decision could include whether they should take an action, and if so which action, to address their situation or need. Further, while this thesis focuses on laypeople, I will also give an overview over future work that also targets professional users, such as legal aid clinics, lawyers and judges.⁷⁹ Thus, I believe the terminology of “legal decision support tools” to be appropriate.

1.4 Structure

This section describes the structure of the thesis. The thesis is split into three sections, containing a number of chapters. Each chapter forms an important part in achieving the goal of this thesis, namely investigating how artificial intelligence can increase access to justice and legal information by creating a methodology to build legal decision support tools.⁸⁰ The chapters thus address one or more of the subobjectives listed in 1.2.2.

1.4.1 Part I – Background

The first part of the thesis aims to provide important background for the thesis, namely exploring the field of artificial intelligence and exploring the issue of access to justice.

In **Chapter 2**, I give a brief overview over the field of artificial intelligence, including the symbolic approach and machine learning. Understanding the promises and shortcomings of these approaches will serve as important background in understanding how they can be

⁷⁹ See 9.2.

⁸⁰ See Overall research objective

used to build my methodology. This chapter corresponds to the subobjective 1.2.2.1, Understanding artificial intelligence.

In **Chapter 3**, I explain the issue of access to justice. I examine the issues of everyday legal problems, and how people resolve such issues today. I also explain what is meant by the terms access to justice and access to legal information, and how they are currently being addressed using technology and AI. This section will help me understand how these issues affect individuals and society, and therefore guide me in building a methodology to address them. This chapter aims to accomplish the subobjective set out in 1.2.2.2, Understanding the issues of access to justice.

1.4.2 Part II – Automating Legal Reasoning

Part II explores the steps of legal reasoning, how legal reasoning has previously been automated, sets out important design criteria for me to build a system that can use these approaches to increase access to justice and legal information, and describes an initial approach to building such a methodology.

Chapter 4 gives an overview over the steps involved in legal reasoning. It explores the entire judicial reasoning process, from identifying a guiding legal rule, finding of facts, categorizing these facts to finally arriving at an outcome for the case. For each such step, I also describe prior work in the field of AI & Law aiming to automate it. Understanding legal reasoning, and the capacity of computers to perform this step, is crucial context to inform the creation of my methodology. This chapter aims to accomplish subobjective 1.2.2.3, Understanding legal reasoning and automating legal reasoning.

In **Chapter 5**, I detail a number of important design criteria in designing a methodology to increase access to justice. Informed by the previous chapters, I determine a number of important considerations in order to make sure that the methodology can be useful to laypeople, work well in areas of high-volume, low-intensity cases⁸¹, give specific and

⁸¹ I.e. cases that frequently appear in court and do not raise complex legal questions (see 3.2.1).

useful advice, and be practical. This chapter corresponds to subobjective 1.2.2.4, Determining the relevant design criteria.

Chapter 6 presents the FactorBot methodology, which aimed to model legal reasoning by representing legal decisions in factors. This research was an important step in building the JusticeBot methodology, which is the main contribution of this thesis. It thus represents an initial look into subobjective 1.2.2.5, Designing a methodology for creating legal decision support tools.

1.4.3 Part III – The JusticeBot methodology

Part III presents and evaluates the final JusticeBot methodology.

Chapter 7 presents the JusticeBot methodology, which allows the building of tools that use an augmented intelligence model in order to increase access to justice. I describe the important steps involved with building such tools, including how to represent the data, how to capture the case of the user, how to analyze the case of the user, and how to provide them with legal information. I further describe the JusticeCreator, which is a tool that allows the visual creation of JusticeBot tools without technical knowledge, and the JusticeBot front-end, which allows user interaction with the system. This chapter represents the most important contribution of this thesis. It corresponds to subobjective 1.2.2.5, Designing a methodology for creating legal decision support tools and 1.2.2.6, Implementing the resulting methodology.

In **Chapter 8**, I present the first implemented JusticeBot, which is focuses on landlord-tenant disputes, developed at the Cyberjustice Laboratory. I describe the legal area and why it is appropriate for the JusticeBot methodology, describe the implementation process and resulting product (which was launched in the summer of 2021 at <https://justicebot.ca>) and discuss the feedback we have received from users of the system. This chapter corresponds to subobjective 1.2.2.7, Validating the resulting methodology.

In **Chapter 9**, I give an overview over other application areas of the JusticeBot, and how the platform could be extended. This chapter addresses subobjective 1.2.2.8, Discussing other application areas and future work.

Chapter 10 wraps up the thesis, summarizes the main contribution, discusses characteristics and limitations of the methodology, and discusses whether the overall objective (1.2.1) has been fulfilled.

Part I

Background

Chapter 2 What is Artificial Intelligence?

Research Objective: Understanding artificial intelligence (1.2.2.1)

Research Topics:

- What is meant by artificial intelligence?
- Which kind of tasks can AI be used to solve?
- Which concrete methods exist to build artificial intelligence systems?
- What are the promises and limitations of the methods to build AI systems?

2.1 Introduction

In order to use artificial intelligence (AI) to support access to justice, we must first investigate and understand what is meant by artificial intelligence, as well as the current approaches to implement intelligent systems, and their limitations.

Artificial Intelligence is currently one of the most spoken about fields within computer science. According to the AI Index 2022 report, over 330k publications relating to AI were published in 2021.⁸² 172k of these were journal publications, representing 2.5% of all journal publications.⁸³ In conference publications, this percentage is much higher – 17.8% of all conference publications in 2021 were in the field of artificial intelligence.⁸⁴ The interest has also soared in the private sector. In 2021, private investment in AI totaled around 93.5 billion USD.⁸⁵

Likewise, the performance of machine learning models has constantly been increasing. In the ImageNet challenge, models are evaluated based on whether they are able to recognize objects in images. The top-5 accuracy (i.e. whether the correct label was among the top 5 picks of the AI model) has increased from under 85% in 2012 (at the

⁸² *The AI Index 2022 Annual Report*, by Daniel Zhang et al (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University) at 17.

⁸³ *Ibid* at 24.

⁸⁴ *Ibid* at 28.

⁸⁵ *Ibid* at 3.

start of the deep learning revolution) to 99% in 2021, higher than the human baseline performance of 94.90%.⁸⁶ Deep learning, a new approach to artificial intelligence, has led to marked improvements in tasks that were once thought to be impossible to perform using artificial intelligence, including the playing of complex board games such as Go, the understanding of language and the generation of images.⁸⁷

Experts believe that these developments will have an important impact on society in the near future. According to a 2013 study, 47 percent of the total US employment is at risk of automatization over the next decade or two. The jobs that are at the highest risk include transportation and logistics occupations, many office and administrative support jobs, and manual labor.⁸⁸

At the same time, some voices are more critical of deep learning, the current paradigm in artificial intelligence. A popular blog post heralded the beginning of an “AI Winter”, a period of cooling of interest in AI, in 2018.⁸⁹ Another critical voice, Gary Marcus, warns against overvaluing the capabilities of deep learning. He lists 10 issues with the approach, including the fact that AI systems often have difficulties dealing with situation that are not part of the training data.⁹⁰

In this chapter, I will examine the current state of artificial intelligence. This is crucial in order to understand how AI can be used to increase access to justice, which is the purpose of this thesis. I will discuss defining artificial intelligence (2.2) and explain the difference between general and narrow artificial intelligence (2.3). I will also describe the concept of a “task” that can be carried out by an AI system (2.4). Next, I will describe two of the main techniques for building AI systems (the symbolic approach (2.5) and machine

⁸⁶ *Ibid* at 53.

⁸⁷ See 2.6.2

⁸⁸ Carl Benedikt Frey & Michael A Osborne, “The future of employment: How susceptible are jobs to computerisation?” (2017) 114 *Technological Forecasting and Social Change* 254–280 at 44–45.

⁸⁹ Filip Piekiewicz, “AI winter is well on its way”, (29 May 2018), online: *Piekiewicz’s blog* <<https://blog.piekiewicz.info/2018/05/28/ai-winter-is-well-on-its-way/>>.

⁹⁰ Gary Marcus, “Deep learning: A critical appraisal” (2018) arXiv preprint arXiv:180100631 at 16.

learning(2.6)), including their advantages, disadvantages and how they may be used in the legal field to support access to justice. Finally, I will wrap up the chapter (2.7).

2.2 Defining Artificial Intelligence

There is no consensus on how exactly artificial intelligence should be defined.⁹¹

According to Pei Wang, definitions of artificial intelligence can take aim at the structure, behavior, capability, function or principle of AI.⁹² He believes that which definition is chosen will have an important impact on shaping the direction of research, and thus on the resulting AI systems.⁹³ As a definition of intelligence, Wang proposes the following:

*“Intelligence is the capacity of an information-processing system to adapt to its environment while operating with insufficient knowledge and resources.”*⁹⁴

This definition includes the idea that the system needs to learn from its environment, since there may not be a perfect logic-based solution available.⁹⁵

Another prevalent definition comes from the Computer Science literature and defines AI as the field that “attempts to build intelligent entities”.⁹⁶

There are also a number of definitions from the legal field. For example, the Council of Europe recently adopted a charter of the ethical use of AI in the justice system, and gave the following definition:

*“A set of scientific methods, theories and techniques whose aim is to reproduce, by a machine, the cognitive abilities of human beings.”*⁹⁷

⁹¹ Pei Wang, “On defining artificial intelligence” (2019) 10:2 Journal of Artificial General Intelligence 1–37 at 1.

⁹² *Ibid* at 8–12.

⁹³ *Ibid* at 13–14.

⁹⁴ *Ibid* at 17.

⁹⁵ *Ibid* at 17–20.

⁹⁶ Stuart J Russell, Peter Norvig & Ernest Davis, *Artificial intelligence: a modern approach*, 3rd ed, Prentice Hall series in artificial intelligence (Upper Saddle River: Prentice Hall, 2010) at 1.

Other researchers, such as Paul Dumouchel, argue that defining artificial intelligence requires a definition of human intelligence, which we currently lack. Instead, he argues, AI “does not correspond to a single faculty, but to a collection of computational technologies inspired by some human cognitive abilities.”⁹⁸

In this thesis, I aim to use artificial intelligence systems to increase access to justice, likely by emulating some form of legal reasoning. For the purposes of this thesis, I will use a broad conceptualization of artificial intelligence. Therefore, AI is the technology that is able to perform the tasks needed to increase access to justice. Which particular class of programs that is used, or whether they can learn from experience, is not important for this goal. This focus seems to most closely correspond to the definitions discussed by Wang focused on *capability*, which define the intelligence of a system by its problem-solving ability.⁹⁹

2.3 General Artificial Intelligence vs Narrow Artificial Intelligence

A distinction is often made between *general* artificial intelligence and *narrow* artificial intelligence. *General* artificial intelligence refers to machines that are able to rival human intelligence, are able to adapt to any situation and transfer knowledge from one field to another completely autonomously.¹⁰⁰ This kind of AI is currently in the realm of science fiction.¹⁰¹

Furthermore, determining whether an AI system should be considered general artificial intelligence may not be a simple task. Several tests have been proposed, such as the Turing test, where people are asked to have a text conversation with a human or an AI

⁹⁷ *European ethical Charter on the use of Artificial Intelligence in judicial systems and their environment* (Strasbourg: Council of Europe European Commission for the efficiency of justice (CEPEJ), 2018) at 69.

⁹⁸ Paul Dumouchel, “Intelligence, Artificial and Otherwise” (2019) 24:2 *Forum Philosophicum* 241–258 at 243.

⁹⁹ Wang, *supra* note 91 at 10.

¹⁰⁰ *Ibid* at 15.

¹⁰¹ note 97 at 70; *The Privacy Expert’s Guide To Artificial Intelligence and Machine Learning* (Future of Privacy forum, 2018) at 6.

and guess whether the other party was a human or an AI system.¹⁰² Another test is the coffee cup test, which requires an AI system to walk into an unknown house and make a cup of coffee. This is feasible for humans, but probably impossible to achieve for computer systems at present.¹⁰³

Narrow artificial intelligence, on the other hand, aims to create highly specialized AI systems that are able to reliably solve specific tasks.¹⁰⁴ While these do not have the general capability to reason and adapt to problems, they can still be tremendously useful and have a large impact on society. In this thesis, I will investigate the building of narrow AI systems, focusing on specific tasks that can increase access to justice.

2.4 Tasks

2.4.1 Introduction

As we have seen, most research today focuses on narrow AI systems, that are able to solve a specific “task”. Let us consider what we mean when we say “task”. I have found it helpful to conceptualize AI tasks in terms of an input (i.e. the data the AI is given to accomplish a task, or the environment as sensed by the AI system) and an output (i.e. the desired response from an AI system, corresponding to accomplishing the task). Table 1 shows some such tasks that could be solved with the help of artificial intelligence.

Input	Output
The properties of a house, such as the size, age, neighborhood etc.	The predicted sale price of the house
The characteristics of a person, such as age, blood values and diet	Is this person susceptible to a certain disease?
Measurements of a flower, such as petal and	What type is the flower?

¹⁰² Luke Muehlhauser, “What is AGI?”, (11 August 2013), online: *Machine Intelligence Research Institute* <<https://intelligence.org/2013/08/11/what-is-agi/>>; note 101; Ben Goertzel, Matt Iklé & Jared Wigmore, “The Architecture of Human-Like General Intelligence” in Pei Wang & Ben Goertzel, eds, *Theoretical Foundations of Artificial General Intelligence* (Paris: Atlantis Press, 2012) 123 at 140.

¹⁰³ *The Privacy Expert’s Guide To Artificial Intelligence and Machine Learning*, by Future of privacy forum (2018) at 5.

¹⁰⁴ note 101 at 6.

sepal length	
A picture, containing an object	What is the object in the picture?
A list of movies a person has seen and liked	Other movies this person may enjoy
A sentence	Is the sentence positive or negative?
A sentence in French	The sentence translated to English
A chess board	Which move should be made to maximize the chance of either side winning?
Data from sensors attached to a car	How should the controls of the car be maneuvered in order to safely transport the passengers from one place to another?
A recording of a person speaking a sentence	The transcribed version of that sentence
A sentence describing a picture	A generated picture, corresponding to the description.
The factual situation of a user	Information about their legal rights in that situation

Table 1 – A list of different tasks that can be accomplished by AI systems, including an input and output.

It is important to note that task performance is often independent. The fact that an AI system can beat a human in chess does not mean that the AI system can beat that human in reasoning or understanding an essay - these capabilities are orthogonal. Dumouchel argues that it is useless to aim to compare human intelligence to artificial intelligence overall, since comparisons can only be undertaken on specific, particular abilities.¹⁰⁵

The AI model can be seen as the system that sits between the inputs and the outputs, that translates from a provided input to the desired output. For each given task, there may be multiple different algorithms that are able to accomplish the task. Below, in 2.5 and 2.6, I describe different approaches to building AI systems, such as expert systems and machine learning models. Each approach to build such models has different trade-offs and may be more or less suited for different types of tasks.

¹⁰⁵ Dumouchel, *supra* note 98 at 244.

However, it is important to understand that even the most sophisticated models are not able to master every task. While some tasks may be relatively easy, and can be solved with simple AI approaches, other tasks may be impossible to solve at all. Let us discuss how the difficulty level may differ between different tasks.

2.4.2 Difficulty levels

The tasks accomplished by an AI system may have very different levels of “difficulty”. Below, I describe some levels of difficulty of such systems.

2.4.2.1 Tasks with a known solution

Some tasks have a known solution – the rules to solve these tasks are apparent to us. For example, if a task consists of informing us of how much water we need to boil a certain amount of rice, there is likely a straightforward formula that allows us to calculate the perfect amount. Since we know the solution, there is no need to learn from examples to discover it. Instead, the computer system can be programmed to directly solve the task with the rule. Systems that are able to solve such tasks can be built using the symbolic approach (see below under 2.5). Some definitions of AI may exclude this type of systems. However, as discussed in 2.2, for this thesis I will consider such systems to be artificial intelligence.

2.4.2.2 “Easy” tasks

Other tasks may not have straightforward rules or formulas, and thus require the creation of a “model” of how reality works. An example of such a task is the prediction of the sale price of a house, given the properties of that house (such as number of windows, size etc). Here, there is no given rule that determines the price the house will be sold at. However, there are probably relatively simple patterns that allow us to gain a useful understanding of the house price. For example, the location, the size and the age of the house are likely to affect the price in relatively predictable ways. By analyzing examples of houses that have been sold, it is possible to build an approximation that allows us to assess whether a house on sale is a good deal. Such tasks can often be accomplished with traditional machine learning methods, described in 2.6.1.3.

2.4.2.3 Complex tasks

Other tasks may be more complex and difficult to accomplish. They may require the understanding and evaluation of thousands or millions of interconnected features. For example, understanding how a recording maps to spoken words is very complex, as the computer has to analyze thousands or millions of noisy datapoints to detect which words are spoken. Likewise, classifying the object in an image is a complex task – it requires the computer to look at thousands of pixels, and understand their relationship. The relationship between the input and output is much more complex. Today, many such tasks can be accomplished with deep learning models (see 2.6.1.3).

2.4.2.4 AI-complete tasks

Other tasks may be beyond what AI can currently accomplish. The argument goes that some tasks require not just an understanding of a certain task, but also the understanding of the context of how the world generally works. These are referred to as “AI-complete” problems,¹⁰⁶ meaning that solving these tasks using computers would require the development of general artificial intelligence systems.¹⁰⁷ Such problems may include fully self-driving cars, or fully understanding human language. Unlike machines, humans are a part of the world, and are able to consider their lived experience when trying to understand a sentence, rather than just the word itself.

The question of which tasks are AI-complete tasks is still open, and tasks that were once considered AI-complete can now be solved with high accuracy using advanced AI systems. For example, machine translation can now be undertaken with performance that is similar to that of humans.¹⁰⁸ AI-complete problems will also be discussed in 2.6.3.5.

¹⁰⁶ Roman Yampolskiy, “AI-Complete, AI-Hard, or AI-Easy: Classification of Problems in Artificial Intelligence” (2012) The 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Cincinnati, OH, USA at 3.

¹⁰⁷ See 2.3.

¹⁰⁸ Stuart J Russell & Peter Norvig, *Artificial intelligence: a modern approach*, fourth ed, Pearson series in artificial intelligence (Hoboken: Pearson, 2021) at 29.

2.4.2.5 Impossible tasks

Finally, certain tasks may simply be impossible to solve. For example, looking at the sky in a single place and predicting the weather a week down the line is likely impossible, no matter which AI system is used. In this case, the input data is simply not enough to generate a reliable model for the prediction of the weather. It is important to keep in mind that AI is not magic,¹⁰⁹ and that certain tasks are simply impossible to accomplish with certain input data.

2.4.3 The integration of tasks

Fully understanding the impact of an AI system requires us to look beyond the task itself. The task, analyzing inputs to generate outputs, is merely a part of the finished, AI-enabled project. The input has to come from somewhere, such as being captured by sensors or provided by users of the system. Likewise, the output of the system will likely be used for a certain purpose, such as informing a decision or directly taking an action.

Even if the algorithm itself works perfectly, difficulties may arise at the interaction points of the AI with the real world. For example, the sensor used to capture the data may be faulty, or capture the data in a way that is unexpected to the system. Or, the user may rely fully on the output of an AI system in making a decision, not considering possible shortfalls in the algorithms, leading to decisions that are harmful. Therefore, it is important to consider the system as a whole, rather than just the AI part in isolation.¹¹⁰

Different implementations of AI systems may have very different risk profiles – for example, an AI used by a doctor or by a self-driving car has a much higher potential for harmful outcomes than an AI system suggesting movies to a user. This has been

¹⁰⁹ Pedro Domingos, “A few useful things to know about machine learning” (2012) 55:10 Commun ACM 78–87 at 81.

¹¹⁰ Andrew D Selbst et al, “Fairness and abstraction in sociotechnical systems” (2019) Proceedings of the conference on fairness, accountability, and transparency 59–68.

understood by regulators – the European union, for example, currently has a proposal to regulate AI-systems based on their risk level.¹¹¹

Now that we have understood what a task and narrow artificial intelligence is, let us take a look at some approaches to build AI systems that can accomplish these tasks. First, I will briefly describe the symbolic approach (including expert systems) and machine learning, both of which are relevant and used in the field of artificial intelligence and law.

2.5 The Symbolic Approach

Research into the field of artificial intelligence arguably began in 1956, at a summer conference bringing together researchers to investigate the use of computers to build intelligence systems.¹¹² Early on, the “symbolic” approach to artificial intelligence was favored. This approach was also referred to as Good Old-Fashioned Artificial Intelligence (GOFAI).¹¹³

The symbolic approach relies on encoding objects from the real world into “symbols”, i.e. logical manifestations of ideas and objects, as well as their relationships.¹¹⁴ The computer is further given rules on how to manipulate these symbols to achieve certain tasks. Since the symbols are not directly connected to the real world,¹¹⁵ many different real-world objects could in theory be represented as a particular symbol, allowing the application of the approach to many problems. The approach can be seen as top-down, as knowledge is

¹¹¹ *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS*, COM/2021/206 final.

¹¹² Russell & Norvig, *supra* note 108 at 18.

¹¹³ John Haugeland, *Artificial intelligence: the very idea* (Cambridge, Mass: MIT Press, 1985).

¹¹⁴ H R Ekbia, *Artificial Dreams: The Quest for Non-Biological Intelligence* (Cambridge: Cambridge University Press, 2008) at 24.

¹¹⁵ Andre Vellino, “Artificial intelligence: The very idea: J. Haugeland, (MIT Press, Cambridge, MA, 1985); 287 pp.” (1986) 29 *Artificial Intelligence* 349–353; Paul Smolensky, “Connectionist AI, symbolic AI, and the brain” (1987) 1:2 *Artificial Intelligence Review* 95–109 at 98.

explicitly encoded into the system by the creator, rather than being learnt by the machine itself.¹¹⁶

McCarthy provides the hypothetical example of the “advice taker”. Here, the creator of a system would encode the symbols for a desk, a car and an airport, as well as their relationships. The user can then provide the goal of going to the airport, at which point the system would work with the encoded symbols to explain to the user that they should leave the desk, go to the car and drive to the airport.¹¹⁷

The researchers at the time believed that this form of reasoning explains the process of human reasoning, and that it could thus be built to create machines as intelligent as humans.¹¹⁸

However, while these systems were able to achieve impressive early results,¹¹⁹ they often had difficulties to tackle more complex, real-world problems. Firstly, they relied on the encoding of symbols. This works well for certain high-level concepts,¹²⁰ such as mathematical formulas and chess pieces. However, reality is made up of concepts that are difficult to encode into clear symbols. How, for example, could we define how to ride a bike in terms of symbols? Likewise, natural language is full of vague and context-dependent words, which are difficult to encode in terms of logical symbols.¹²¹ Perhaps, the symbolic approach works best for problems described as “tasks with a known solution” above,¹²² where the rules of how to solve a certain problem are known.

¹¹⁶ Stephen F Davis & William Buskist, *21st Century Psychology: A Reference Handbook* (SAGE, 2008) at 487 Google-Books-ID: tMv1EbXGen4C.

¹¹⁷ John McCarthy, “Programs with common sense” (1960) RLE and MIT computation center Cambridge, MA, USA at 8.

¹¹⁸ Smolensky, *supra* note 115; Allen Newell & Herbert A Simon, “Computer science as empirical inquiry: symbols and search” (1976) 19:3 *Commun ACM* 113–126; Russell & Norvig, *supra* note 108 at 19.

¹¹⁹ Russell & Norvig, *supra* note 108 at 21.

¹²⁰ Davis & Buskist, *supra* note 116 at 487.

¹²¹ Terence Horgan & John Tienson, “Representations without Rules” (1989) 17:1 *Philosophical Topics* 147–174 at 151.

¹²² See 2.4.2.1.

Likewise, the systems were limited by what is referred to as the combinatorial explosion. This means that even problems that can in theory be solved using symbols, such as chess, will often “explode” in terms of the numbers of calculations required, requiring thousands of years to calculate using computers. While it is theoretically possible to encode every possible state of a chess board in terms of symbols, the sheer number of possibilities make it impossible to calculate every possible state.¹²³

Due to these difficulties, the focus of AI research soon shifted towards other methods.¹²⁴

2.5.1 Expert Systems

After the symbolic approach, the attention of researchers shifted to so-called expert systems. This approach can be seen as a subset of the symbolic approach, since they also rely on logical reasoning. However, instead of focusing on building general intelligence or solving abstract problems, these systems focus on encoding the knowledge of human experts into logical rules that the algorithm can traverse in order to solve typical, real-world problems.¹²⁵

Expert-systems usually comprise three components: The knowledge base, the inference engine and the user interface.

The *knowledge base* is constructed by interviewing experts in the field and encoding their knowledge into a computer-readable format by “knowledge engineers”. This process can be slow and expensive, potentially costing millions of dollars.¹²⁶

Often, the knowledge is encoded in so-called production rules, that contain an IF and a THEN clause. An example of such a rule could be:

*IF (fever) THEN (predict infection).*¹²⁷

¹²³ Klaus Krippendorff, “A Dictionary of Cybernetics” (1986) at 12.

¹²⁴ Russell & Norvig, *supra* note 108 at 21–22.

¹²⁵ Bruce G Buchanan, “A (Very) Brief History of Artificial Intelligence” (2005) 26:4 AI Magazine 53–53 at 59.

¹²⁶ Bruce G Buchanan & Reid G Smith, “Fundamentals of Expert Systems” (1988) 3:1 Annual Review of Computer Science 23–58 at 19–20; Ekbia, *supra* note 114 at 95; Ashley, *supra* note 44 at 8.

Some systems comprised thousands of such rules.¹²⁸

These rules are used by the *inference engine* of an expert system. This engine is able to take some external reading of a new situation (such as a sensor-reading, or answers to questions entered by a user), and traverse the rules logically in order to arrive at a conclusion regarding this specific situation.¹²⁹ While the rules themselves are specific to a certain domain and task, the inference engine can in theory be used for many different domains.¹³⁰

The inference engine can be built to reason forwards (data-directed) or backwards (goal-directed). Forwards reasoning starts with the facts, and then reasons through them to arrive on a conclusion.¹³¹ Backwards reasoning, on the other hand, starts with the goals and then reasons backwards to identify the facts that are necessary to arrive at a certain goal.¹³² Some inference engines may include approaches to reason about rules that are uncertain or incomplete.¹³³

Finally, the system needs to provide a *user interface* that can interact with the user. As Buchanan and Smith point out, this is a very important factor for the usability of the system.¹³⁴ The interface can consist of a textual prompt system, that asks the user questions, or a graphical interface. In asking the questions, it is important to be aware that experts might use different terminology and points of view on situations than the desired end-user of a product.¹³⁵

¹²⁷ Smolensky, *supra* note 115 at 98.

¹²⁸ Buchanan & Smith, *supra* note 126 at 29.

¹²⁹ *Ibid* at 17–19.

¹³⁰ *Ibid* at 14.

¹³¹ *Ibid* at 16.

¹³² *Ibid* at 17; Ashley, *supra* note 44 at 9–10.

¹³³ See e.g. L A Zadeh, “The role of fuzzy logic in the management of uncertainty in expert systems” (1983) 11:1 Fuzzy Sets and Systems 199–227.

¹³⁴ Buchanan & Smith, *supra* note 126 at 34.

¹³⁵ *Ibid* at 20.

Expert systems found use in several fields such as the medical, engineering, education and legal domain.¹³⁶ Notable examples include systems developed to recognize molecules based on spectrogram readings, and to diagnose blood infections.¹³⁷ Many expert systems found commercial success, with certain projects saving their company millions of dollars.¹³⁸ As we will see later, expert systems also found a number of applications in the field of AI and law. Soon, developing expert systems had become a billion-dollar industry.¹³⁹

2.5.2 Discussion

Expert systems were a popular way to create artificial intelligence systems, but have fallen out of favor to some extent. Let us discuss some of the advantages and issues of this type of system.

2.5.2.1 Cost and effort to develop systems

It is easy to get started to develop and implement expert systems. There were a number of systems that made it easy to encode rules into a system, and thus allowed the quick prototyping of new expert systems.¹⁴⁰ Prototyping expert systems was seen as an important step in evaluating the usability of such systems.¹⁴¹

However, building and maintaining a fully functional expert system is often much harder than building the prototype, and can be an expensive and time-consuming endeavor.¹⁴² Often, the rules required to fully solve a task can be numerous and complex, costing millions of dollars to encode in a knowledge base.¹⁴³ Once the system is created, it has to be updated and maintained to ensure the continued accuracy. The creators of XCON, a

¹³⁶ Thompson, *supra* note 75 at 13.

¹³⁷ Russell & Norvig, *supra* note 108 at 22–23; Edward H Shortliffe et al, “Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system” (1975) 8:4 Computers and Biomedical Research 303–320.

¹³⁸ Russell & Norvig, *supra* note 108 at 23.

¹³⁹ *Ibid* at 24.

¹⁴⁰ Buchanan & Smith, *supra* note 126 at 22–24.

¹⁴¹ *Ibid* at 20.

¹⁴² Ashley, *supra* note 44 at 11.

¹⁴³ Buchanan & Smith, *supra* note 126 at 19.

system built to configure computer systems based on customer orders, claimed that 50% of the system would have to be rewritten each year in order to keep up with changing requirements.¹⁴⁴ A report found that issues with integrating and maintaining expert systems were a big impediment for their continued use.¹⁴⁵ Further, it can be difficult to validate how well expert systems works, such as whether they provide correct and useful information, leading to such analyses rarely being done, according to Buchanan and Smith.¹⁴⁶

2.5.2.2 Explainability and vague concepts

Expert systems are inherently explainable. Since the conclusion of an expert system is arrived at through the perusing of encoded rules, it is always possible to examine which rules were used to arrive at a given conclusion. Many expert systems integrated a facility to interactively explore how a certain conclusion was reached.¹⁴⁷

However, the corollary to this is that expert systems are also only able to encode concepts that we understand and can enter in the form of explicit rules. In theory, this works well for certain kinds of knowledge, such as in the legal and medical domain. However, in practice, experts may rely more on intuition than one might expect. The creators of MYCIN, a medical expert system, found that doctors often relied on intuition, making the encoding of rigorous rules difficult.¹⁴⁸ In the legal domain, while the knowledge of what the consequences of a legal rule applying are is often explicit (see 4.5.3), deciding whether a rule applies or not often relies on vague concepts and open-textured concepts

¹⁴⁴ John J Sviokla, “An Examination of the Impact of Expert Systems on the Firm: The Case of XCON” (1990) 14:2 MIS Quarterly 127–140 at 137.

¹⁴⁵ T Grandon Gill, “Early expert systems: Where are they now?” (1995) MIS quarterly 51–81 at 68.

¹⁴⁶ Buchanan & Smith, *supra* note 126 at 24.

¹⁴⁷ *Ibid* at 26.

¹⁴⁸ R Duda & E Shortliffe, “Expert Systems Research” (1983) 220:4594 Science 261–268 at 265.

(see 4.4).¹⁴⁹ This kind of reasoning would likely be difficult to encode in an expert system.¹⁵⁰

2.5.2.3 Generalization

Expert systems further have issues generalizing, i.e. performing well on new, unseen situations. If a case is not covered by a rule in the knowledge base, the system will fail.¹⁵¹ Further, the system may not be aware that the case is not covered by the knowledge base, causing it to give a confidently incorrect answer.¹⁵² Expert systems do not have a notion of common sense that would allow them to solve a problem in the absence of a specific rule.¹⁵³

These issues lead to another period of decreased interest in artificial intelligence.¹⁵⁴

Researchers investigated the market for expert systems in 1995 and found that two thirds of the systems were no longer maintained, with many being inaccessible.¹⁵⁵

2.5.3 Application in the legal domain?

Now that we have understood the symbolic approach, including expert systems, let us consider how appropriate this approach is for building decision support tools in the legal domain.

In the symbolic approach, phenomena are encoded into a system as symbols and manipulated using specific rules. When judges reason with previous decisions, they can be seen to apply certain reasoning patterns to case law, such a drawing analogies between a new case and previous cases.¹⁵⁶ If cases can be represented as symbols, and the rules that judges use to reason with cases can be discovered, it may be possible to build a

¹⁴⁹ H L A Hart, “Positivism and the Separation of Law and Morals” (1957) 71 Harv L Rev 593–629 at 607; Ashley, *supra* note 44 at 10.

¹⁵⁰ See 4.5.3.3.3.

¹⁵¹ Buchanan & Smith, *supra* note 126 at 14.

¹⁵² *Ibid* at 22; John McCarthy, “Some expert systems need common sense” (1984) at 3.

¹⁵³ Buchanan & Smith, *supra* note 126 at 15.

¹⁵⁴ Russell, Norvig & Davis, *supra* note 96 at 22–24.

¹⁵⁵ Gill, “Early expert systems”, *supra* note 145 at 68.

¹⁵⁶ See 4.9.

symbolic system that is able to perform judicial reasoning. We will explore such approaches below in 4.4.3.4. As we will see, there are a number of practical challenges that need to be overcome in creating such systems, including deciding how cases can be turned into symbols, how to practically encode a sufficient number of cases, and how to perform complex reasoning steps with these cases.

Likewise, it may be possible to represent legal rules in expert systems. In the legal domain, there are often statutory rules, that consist of criteria logically connected to possible conclusions.¹⁵⁷ Expert systems are well suited to model domains governed by explicit logical rules, that can be encoded into the knowledge base and then used by the system to arrive at a result. This has led a number of researchers to explore using the symbolic approach and expert systems to model statutory legal reasoning.¹⁵⁸

At the same time, expert systems work less well in domains where intuition or interpretation is important. In the legal domain, the structure of the legal rules may not always be clear from reading the text of the law, leaving the parties room for argumentation.¹⁵⁹ Any expert system created in such a domain would only be able to reason about a specific interpretation of the law.

Further, while the structure of the law is well suited for an expert system, it would be more difficult to encode the substantive content of the law. Often, vague concepts (such as “reasonable”) determine whether a criterion is fulfilled or not.¹⁶⁰ Encoding such concepts in terms of symbols may be difficult, since such criteria are evaluated in the context of a specific, real-world situation that may have many different interdependent aspects.

¹⁵⁷ See 4.5.

¹⁵⁸ See 4.5.3.2.

¹⁵⁹ See 4.5.3.3.4.

¹⁶⁰ See 4.4.

Likewise, expert systems can be slow and expensive to build. Introducing methodologies that are able to speed up this process is another important component of building expert systems in the legal domain.

2.6 Machine Learning

The difficulty of using expert systems on complicated tasks lead to another AI winter. Researchers began turning to *machine learning*, which is more connected to previous research in mathematics and statistics, and evaluation on real-world datasets. This approach rekindled the interest in artificial intelligence.¹⁶¹ In 2012, a subset of machine learning known as deep learning started achieving incredible results on many tasks,¹⁶² which has catapulted artificial intelligence into the spotlight on the world stage.

In this section, I will examine machine learning. Machine learning can be defined as “The field [that] is concerned with the question of how to construct computer programs that automatically improve with experience.”¹⁶³ In the symbolic approach, the creator of a system explicitly provides the algorithm with the rules necessary to solve a problem. In machine learning, on the other hand, the developer provides the algorithm with a number of examples of the task successfully being solved. It is then up to the algorithm to discover the patterns in these examples and build a so-called *model*.¹⁶⁴ This model can be used to solve this task for new, previously unseen, examples.¹⁶⁵

First, I will describe the steps involved in building a machine learning system (2.6.1). Then, I will explore some use-cases of this technology (2.6.2). Finally, I will discuss the promises and shortcomings of the approach (2.6.3).

¹⁶¹ Russell & Norvig, *supra* note 108 at 24–26.

¹⁶² Yann LeCun, Yoshua Bengio & Geoffrey Hinton, “Deep learning” (2015) 521:7553 Nature 436–444; Md Zahangir Alom et al, “The history began from alexnet: A comprehensive survey on deep learning approaches” (2018) arXiv preprint arXiv:180301164.

¹⁶³ Tom M Mitchell, *Machine Learning*, 1st ed (New York: McGraw-Hill Education, 1997).

¹⁶⁴ Domingos, *supra* note 109 at 78.

¹⁶⁵ note 101 at 7.

2.6.1 The steps of building a machine learning system

There are several important steps involved in creating a machine learning pipeline to solve a certain task.¹⁶⁶ Understanding these steps is relevant to grasp the concept of machine learning. I will describe these steps below. Each step will be illustrated with an example project (Figure 1), aiming to build an app that can take a picture and tell the user whether they are looking at an apple or an orange.



Figure 1 - An example task, determining whether a picture contains an apple or an orange

¹⁶⁶ compare Yufeng Guo, “The 7 Steps of Machine Learning”, (31 August 2017), online: *Towards Data Science* <<https://towardsdatascience.com/the-7-steps-of-machine-learning-2877d7e5548e>>; Harini Suresh & John V Guttag, “A Framework for Understanding Unintended Consequences of Machine Learning” (2020) arXiv:1901.10002 [cs, stat], online: <<http://arxiv.org/abs/1901.10002>> arXiv: 1901.10002; Chanin Nantasenamat, “How to Build a Machine Learning Model”, (25 July 2020), online: *Towards Data Science* <<https://towardsdatascience.com/how-to-build-a-machine-learning-model-439ab8fb3fb1>>; Nithya Sambasivan et al, “‘Everyone wants to do the model work, not the data work’: Data Cascades in High-Stakes AI” (2021) proceedings of the 2021 CHI Conference on Human Factors in Computing Systems 1–15 at 6.

2.6.1.1 Choosing a task

The first step in building a machine learning pipeline is deciding on a task that the machine learning model should solve.¹⁶⁷ This can arise, for example, from a business need or a research project.

Depending on the task, different classes of machine learning models may be used. Let us take a look at a few such classes:

Supervised learning – supervised learning refers to learning where we want to predict a certain *target value* (or *label*) based on a set of *features*.¹⁶⁸ The features are properties describing a datapoint, which are given to the model as an input. The label is the expected output of the model, i.e. the question we are asking.¹⁶⁹ The task that a model performs in supervised learning can be described as: “Based on these features, what is the label of this datapoint?” Examples of tasks in supervised learning include predicting which object is contained in an image based on the pixel values of that image, predicting the price of a house based on the properties of the house, or predicting which word is next in a sequence of words. Supervised learning is the most explored and efficient type of machine learning. However, as we will see, supervised learning requires assembling large datasets with labels, which can be difficult and costly.

Unsupervised learning – unsupervised learning is based on exploring patterns in data without necessarily having a target in mind.¹⁷⁰ Examples of tasks in unsupervised learning include clustering multiple similar samples together¹⁷¹ or detecting an outlier in data, such as identifying fraudulent credit card transactions.¹⁷² Unsupervised learning is

¹⁶⁷ Russell & Norvig, *supra* note 108 at 704.

¹⁶⁸ note 101 at 10.

¹⁶⁹ Russell & Norvig, *supra* note 108 at 653.

¹⁷⁰ note 101 at 16.

¹⁷¹ Russell & Norvig, *supra* note 108 at 653.

¹⁷² N Malini & M Pushpa, “Analysis on credit card fraud identification techniques based on KNN and outlier detection” (2017) 2017 Third International Conference on Advances in Electrical, Electronics, Information, Communication and Bio-Informatics (AEEICB) 255–258.

more difficult but has the advantage of not requiring troves of labeled data, making researchers such as Yann LeCun believe that it may be the future of machine learning.¹⁷³

Reinforcement learning – Reinforcement learning is another interesting type of machine learning. Here, the goal of the algorithm is to determine a number of actions to take to maximize a certain reward.¹⁷⁴ This approach has been used extensively in learning to play games, such as Go¹⁷⁵, Tetris¹⁷⁶ and Starcraft.¹⁷⁷ The reward function in these cases is to win, or to not lose for as long as possible. At first, the algorithm makes random moves, that lead it to lose relatively quickly. However, some of these moves are more successful than others. Over time, the algorithm is able to learn which actions maximize its chance of winning, thus learning how to master a certain task.¹⁷⁸

2.6.1.1.1 Example

For our example task of predicting apples from oranges, *supervised learning* is an appropriate approach. We have a target that we want to predict (does the image contain an apple or an orange?). By giving the algorithm examples of images containing apples or oranges (see Figure 2), together with information on which fruit it contains, we should be able to teach it how to discern the fruits from each other.

¹⁷³ Karen Hao, “The AI technique that could imbue machines with the ability to reason”, (12 July 2019), online: *MIT Technology Review* <<https://www.technologyreview.com/2019/07/12/65579/the-next-ai-revolution-will-come-from-machine-learnings-most-underrated-form/>>.

¹⁷⁴ Russell & Norvig, *supra* note 108 at 789.

¹⁷⁵ Cade Metz, “In Two Moves, AlphaGo and Lee Sedol Redefined the Future”, (16 March 2016), online: *Wired* <<https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/>>.

¹⁷⁶ István Szita & András Lőrincz, “Learning Tetris Using the Noisy Cross-Entropy Method” 18 *Neural Computation* 2006.

¹⁷⁷ Oriol Vinyals et al, “Grandmaster level in StarCraft II using multi-agent reinforcement learning” (2019) 575:7782 *Nature* 350–354.

¹⁷⁸ L P Kaelbling, M L Littman & A W Moore, “Reinforcement Learning: A Survey” (1996) 4 *Journal of Artificial Intelligence Research* 237–285.



Figure 2 - The data required to train a supervised learning system

2.6.1.2 Collecting and preparing data

The next step in building a machine learning model is to collect data that corresponds to the chosen task. This data will be used by the algorithm to discover patterns to learn from, and to evaluate how well the resulting model works. Choosing a dataset that is large, varied and corresponds well to the task we want to achieve is therefore very important for the performance of a model. There may be publicly available datasets that can be used (2.6.1.2.1), while in other instances a new dataset needs to be created (2.6.1.2.2). The data also needs to be prepared for analysis by the computer system (2.6.1.2.3).

2.6.1.2.1 Publicly available datasets

In some instances, there may be public datasets that correspond to the desired task. For example, the ImageNet dataset contains 14 million images with a label indicating the object contained in each image, out of 22k different objects.¹⁷⁹ Likewise, the Mozilla Common Voice dataset contains 7,300 hours of spoken, transcribed voices.¹⁸⁰ Using these datasets can save the researcher tremendous amounts of time. However, it is important to make sure that the data corresponds to the task at hand. Techniques such as transfer learning, where a model is trained on a large general dataset, and later trained on

¹⁷⁹ Jia Deng et al, “Imagenet: A large-scale hierarchical image database” (2009) 2009 IEEE conference on computer vision and pattern recognition 248–255.

¹⁸⁰ “Common Voice by Mozilla”, online: <<https://commonvoice.mozilla.org/>>.

a smaller dataset of task-specific data, can be very powerful in overcoming this limitation while benefitting from the publicly available data.¹⁸¹

2.6.1.2.2 Creating new datasets

If there is no publicly available dataset, the developer has to create their own dataset. This can be a very laborious task and require a number of important choices. Getting these choices wrong can cause poor performance or even harmful consequences.

2.6.1.2.2.1 Which samples should be included?

First of all, the creator has to choose *which samples to include* in the dataset. Generally, the more samples are included the better, as machine learning methods need a lot of data in order to build a good model.¹⁸² However, it is also important that the samples correspond to the task and have a similar distribution to the task we are aiming to solve. For example, if an algorithm is trained to recognize faces, but is trained only on pictures of people with a certain ethnicity, the algorithm may fail when applied to pictures of people with other ethnicities.¹⁸³ Suresh and Guttag refer to this bias as “representation bias”.¹⁸⁴

2.6.1.2.2.2 Which features should be collected?

Next, the creator has to decide which *features* of each sample should be collected. For images and text, this may be obvious – the dataset would include the image data or text. For other tasks, such as predicting outcomes of legal cases, the creator has to make a choice of which features should be included, such as a list of facts that may be relevant, the name of the judge or the parties, or the time of day of a decision. As we will see, this is not always easy – for example, making an exhaustive, well-defined list of possible

¹⁸¹ Sinno Jialin Pan & Qiang Yang, “A Survey on Transfer Learning” (2010) 22:10 IEEE Trans Knowl Data Eng 1345–1359.

¹⁸² Domingos, *supra* note 109 at 84–85.

¹⁸³ Joy Buolamwini & Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification” (2018) Conference on Fairness, Accountability and Transparency 77–91.

¹⁸⁴ Suresh & Guttag, *supra* note 166 at 5.

facts that appear in a case can be prohibitively difficult.¹⁸⁵ Omitting features that turn out to be important can lead the model to perform poorly. Likewise, including too many features that are not relevant to the outcome can confuse the model, leading to decreased performance.¹⁸⁶

2.6.1.2.2.3 *Which labels should be used?*

The creator also has to decide which *labels* should be used for a sample. This is the target value that we want to predict for a certain sample. For example, if we want to predict the object in an image, the object contained in an image would be the label. Picking a label is also not always easy. For legal cases, a relatively obvious label could be whether a party won or lost a specific case. However, using this label ignores many nuances, such as the demands made by the party and whether a party won some claims but lost others.

Sometimes, we may not even have access to the label. In these cases, we have to rely on a proxy of that label. For example, if we want to predict whether a convicted criminal is likely to reoffend, we may use the data of whether they are re-arrested as a proxy for reoffending, in order to train a machine learning model. However, this may be a biased proxy if certain areas are more highly policed than others.¹⁸⁷

2.6.1.2.2.4 *How should the data be labelled?*

Finally, the creator has to find a way to actually label a large number of datapoints. Depending on the amount of data required, and the effort required to label certain data, this can be a significant bottleneck in creating machine learning algorithms.¹⁸⁸ In some instances, the process may rely on employees manually labelling samples,¹⁸⁹ or

¹⁸⁵ See 4.4.3.4.3.2.

¹⁸⁶ note 101 at 9.

¹⁸⁷ Suresh & Guttag, *supra* note 166 at 5–6.

¹⁸⁸ Ciarán Daly, “‘I’m Not A Robot’: Google’s Anti-Robot reCAPTCHA Trains Their Robots To See”, (25 October 2017), online: *AI Business* <<https://aibusiness.com/recaptcha-trains-google-robots/>>; Dave Lee, “Why Big Tech pays poor Kenyans to programme self-driving cars”, (3 November 2018), online: *BBC News* <<https://www.bbc.com/news/technology-46055595>>; Domingos, *supra* note 109 at 85.

¹⁸⁹ Lee, *supra* note 188.

crowdsourcing the labelling process.¹⁹⁰ It is important that the labelling is accurate, as erroneous labels could confuse a model. Researchers discovered that in 10 commonly used datasets, an average of 3.4% of the labels are wrong.¹⁹¹ We conducted experiments to assess the robustness of machine learning models against erroneous labels. While the models were relatively robust overall, certain classes of the data proved to be quite sensitive against errors in the labeling.¹⁹²

There are a number of tricks to make the labelling process more efficient, such as letting the algorithm choose which samples should be labelled first (referred to as active learning).¹⁹³ Together with researchers from the United States and Canada, I have developed methodologies to label sentences more efficiently by allowing experts to select terms that are likely associated with a certain class,¹⁹⁴ or surfacing sentences that are similar to a labelled sentences for more efficient labeling.¹⁹⁵

2.6.1.2.3 Data preparation

Finally, once the data has been collected, it needs to be encoded in a format that the computer can understand. The input format required for machine learning models can be seen as a table, with columns for each feature and rows for each sample. Depending on the type of data, it has to be encoded in different ways. The table below contains a few examples of such encodings:

Type of data	Column	Data in each cell
Picture	Pixel index (i.e. first pixel,	The color value of that pixel

¹⁹⁰ Russell & Norvig, *supra* note 108 at 705.

¹⁹¹ Curtis G Northcutt, Anish Athalye & Jonas Mueller, “Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks” (2021) arXiv:2103.14749 [cs, stat], online: <<http://arxiv.org/abs/2103.14749>> arXiv: 2103.14749.

¹⁹² Hannes Westermann et al, “Data-Centric Machine Learning in the Legal Domain” (2022) arXiv:2201.06653 [cs], online: <<http://arxiv.org/abs/2201.06653>> arXiv: 2201.06653.

¹⁹³ *Active Learning Literature Survey*, Technical Report, by Burr Settles, minds.wisconsin.edu, Technical Report TR1648 (University of Wisconsin-Madison Department of Computer Sciences, 2009).

¹⁹⁴ Hannes Westermann et al, “Computer-Assisted Creation of Boolean Search Rules for Text Classification in the Legal Domain” (2019) *Frontiers in Artificial Intelligence and Applications* 123–132.

¹⁹⁵ Hannes Westermann et al, “Sentence Embeddings and High-Speed Similarity Search for Fast Computer Assisted Annotation of Legal Documents” (2020) *Legal Knowledge and Information Systems* 164–173.

	second pixel etc)	(i.e. red, green)
Text	All words contained in the corpus (i.e. "I", "am", "be" etc)	The number of times the word in a column appears in a certain sample (i.e. 1 time, 10 times etc)
Houses	Features such as square footage, number of windows, age etc	The value of the feature in a certain column in numbers, i.e. 9000 square feet, 10 windows, 50 years old.

Table 2 - Possible representations of data

	Features			Target (label)
House	Square feet	Number of windows	Age	Price
House 1	9000	10	50y	400k
House 2	4000	4	40y	240k
House 3	3500	7	10y	270k

Table 3 - Example of part of dataset with house properties (features) and price (target)

Table 3 shows a part of a dataset with house prices. It contains the features as columns, and the individual samples as rows. Each house also has a target column, which contains the value the model should learn to predict, in this case the price of the house.

Different models can handle different types of representations. The approach described above to handle text (referred to as bag-of-words) works well for certain models, but also does not capture the order of the words in a sentence. More modern deep learning models

can learn representations that also include the order of the words and take into account the semantic similarity of different terms.¹⁹⁶

Altering which features to include, and how to include them, can be a creative pursuit and have a significant effect on the performance of a resulting model.¹⁹⁷ It is further an iterative process, where the model performance using a certain encoding is observed, upon which the encoding can be updated to better represent the data.¹⁹⁸

An important part of data preparation is splitting the data into two parts – the training and the test part. Usually, around 20% of the data is withheld for testing. The training part is used to ”train” the machine learning model, i.e. allow it to detect patterns in the data. The testing part can then be used to evaluate the performance of the model, to examine how well it would fare on cases that it has not yet seen in the real world. This is crucial, since we are not interested in how well a model performs on data it has already seen. The most crucial feature of how well a machine learning algorithm performs is how it would perform in the real world, when given new datapoints.¹⁹⁹

2.6.1.2.4 Example

In our example, we want to separate apples from oranges. Perhaps, it is possible to identify a public data set that has images from these categories. For example, Mihai Oltean has published a dataset called Fruit 360, which contains 90k images of 131 fruits and vegetables.²⁰⁰ Otherwise, we would have to manually assemble such a data set, for example by taking the pictures ourselves or by finding pictures on the internet and manually adding the labels. It is important that the images we collect are varied and

¹⁹⁶ Jacob Devlin et al, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding” (2019) arXiv:181004805 [cs], online: <<http://arxiv.org/abs/1810.04805>> arXiv:1810.04805.

¹⁹⁷ Domingos, *supra* note 109 at 84.

¹⁹⁸ *Ibid.*

¹⁹⁹ Guo, *supra* note 166; Domingos, *supra* note 109 at 80.

²⁰⁰ Mihai Oltean, “Fruits 360”, (18 May 2020), online: *Kaggle* <<https://www.kaggle.com/datasets/moltean/fruits>>.

correspond closely to the kinds of images we expect to see when using our model in the real world.

To prepare the data, we need to choose a form of representation. Since we have the data in the form of images, a well-suited way to feed these to the model would be in the form of pixel data (see Table 2). We may have to change the resolution and colors of the images so that they can be processed by a certain machine learning model.²⁰¹

As a label, we can use the fact whether a picture contains an apple or orange. If we use the dataset described above, these labels may already be present. Otherwise, we can show the pictures to employees or individuals on the internet and ask them to click a button indicating whether the picture contains an apple or an orange. After doing these steps, we have a dataset of labelled images that we can use to train a machine learning model.

2.6.1.3 Choosing an algorithm and training a model

Next, it is time to pick an appropriate machine learning algorithm, and to use this algorithm to train a machine learning model.

There are a number of machine learning algorithms, each with different tradeoffs in terms of performance, accuracy and how easy their decisions are to understand and explain. However, since the algorithms usually take the data in the same format, it is often possible to try out many different algorithms to determine which one is best suited for a certain task.²⁰²

The chosen algorithm is then used to train a machine learning model, in a process where the algorithm is given the training data, and discovers correlations and patterns in the data. This model can then hopefully be used to predict new unseen samples.²⁰³

Let us examine some popular machine learning algorithms.

²⁰¹ Joseph Nelson, “Why should I do pre-processing and augmentation on my computer vision datasets?”, (26 January 2020), online: *Roboflow Blog* <<https://blog.roboflow.com/why-preprocess-augment/>>.

²⁰² Domingos, *supra* note 109 at 87.

²⁰³ Domingos, *supra* note 109.

2.6.1.3.1 K-nearest neighbor

A simple algorithm is the nearest neighbor algorithm. It works by comparing the input data to a database of previous samples and uses the label of the most similar samples to predict the label of the input data. For example, when asked to predict the price of a house of 5000 square foot and 40 years old, the model could find the three houses that are the most similar (perhaps one is 5500 square foot large and 35 years old) and take the average of the prices of those houses as a prediction. This works well for certain problems. But, as we use more and more features, every point becomes somewhat similar to every other point meaning that the method is less effective.²⁰⁴

2.6.1.3.2 Decision Trees

Decision trees are another simple model that aims to devise a kind of flow chart that is able to predict new cases. It looks at the training data in order to determine which features are informative for a certain label and then devise a decision tree that is able to separate the classes based on these features.²⁰⁵ For example, when aiming to classify sentences for positivity and negativity, the model may identify that the word “terrible” correlates with a sentence being negative. When faced with a new sentence with this word, it will thus predict that the sentence is negative. Decision trees are a relatively simple method that is easy to understand and explain, however it is not the most accurate method.²⁰⁶ What would happen, for example, to the sentence “The food was fantastic, and the service was not terrible.”?

2.6.1.3.3 Random Forests

Random forests are a model that combines multiple decision trees into one model. By aggregating the results of multiple decision trees, they are able to produce more accurate and powerful models.²⁰⁷

²⁰⁴ Russell & Norvig, *supra* note 108 at 688; Domingos, *supra* note 109 at 82–83.

²⁰⁵ Russell & Norvig, *supra* note 108 at 658–659.

²⁰⁶ *Ibid* at 665.

²⁰⁷ *Ibid* at 697–698.

2.6.1.3.4 Support Vector Machines

Support vector machines are models that draw lines between different classes in a high dimensional space using something called the kernel trick. This way, they are able to separate classes with complex features.²⁰⁸

2.6.1.3.5 Neural networks and deep learning

Neural networks are likely the most popular and exciting machine learning method at the moment, having led to a number of breakthroughs in many different fields. Neural networks consist of layers of so-called artificial neurons, that are stacked on top of each other. Each neuron is a mathematical function, that takes the input from the previous layers, multiplies it by a certain weight and passes the data onto the next layer.²⁰⁹ When arranged in large structures consisting of many such layers, the approach is referred to as “deep learning”.

When using a neural network to generate a prediction, each sample is entered into the model at the “input” layer, passed through several so-called “hidden” layers²¹⁰ and finally arrives at an “output” layer that corresponds to the prediction of the model.²¹¹ Figure 3 shows an example of such a structure.

²⁰⁸ *Ibid* at 693–695.

²⁰⁹ *Ibid* at 751–752; Ian Goodfellow, Yoshua Bengio & Aaron Courville, *Deep learning*, Adaptive computation and machine learning (Cambridge, Massachusetts: The MIT Press, 2016) at 197.

²¹⁰ Goodfellow, Bengio & Courville, *supra* note 209 at 169.

²¹¹ *Ibid* at 181.

A simple neural network




Figure 3 - Model of simple neural network²¹²

In order to train such a network, each sample is passed through a neural network with random weights. Initially, the output is wrong – the weights that determine which data is passed on between the layers of the network are random. However, after each sample has gone through the network, the algorithm goes backwards through the network, slightly adjusting each weight in the right direction, until the neural network learns to properly predict the samples. This process is called backpropagation.²¹³ After the network has been trained for some time using the training data, it has hopefully learnt the correlation between the input data and the desired output, allowing it to predict new cases.

Deep learning models can be set up in a variety of different ways, with different numbers of layers and types of layers. Generally, the more layers a model has, the better it performs when fully trained.²¹⁴ Compared to the previously discussed machine learning methods, neural networks are able to learn very sophisticated representations of data.

²¹² Wiso, *Simple neural network* (2008) Wikimedia Commons, https://commons.wikimedia.org/wiki/File:Neural_network_example.svg.

²¹³ Russell & Norvig, *supra* note 108 at 755.

²¹⁴ Goodfellow, Bengio & Courville, *supra* note 209 at 198; Russell & Norvig, *supra* note 108 at 769.

They are less reliant on feature engineering, and very well suited to unstructured data, such as images, text and video (see 2.4.2.3).²¹⁵

Training neural networks involves a lot of computation, as millions or billions of parameters have to be adjusted many times to create accurate neural networks. However, by using Graphics Processing Units (GPUs, also used to play video games) the computation can be sped up significantly. More recently, Tensor Processing Units have been developed, which are specifically designed to train neural networks and enable the training of models with billions of parameters.²¹⁶

2.6.1.3.6 Example

In our case, we want to predict whether an object contained in an image is an apple or an orange. In order to train such a model, we would use training part of the data.²¹⁷

We can try to train the different models to see which one performs best. However, it is likely that the deep learning approach will work best, since we are dealing with image data. There are several models available that have already been trained on a large collection of images, and can be adapted to certain tasks by adding a few supplemental images (see e.g. ResNet-50).²¹⁸ This approach, known as transfer learning, can allow us to cheaply use sophisticated deep learning models, and hopefully achieve good results.

2.6.1.4 Evaluating the model

Once the model has been trained, we need to evaluate the performance, in order to decide whether it is ready to be deployed or needs more work. This evaluation can show us how well the model has learnt to generalize, i.e. if it has learnt the underlying connection between the features and the label.²¹⁹ The model may not have properly grasped the

²¹⁵ Russell & Norvig, *supra* note 108 at 750.

²¹⁶ *Ibid* at 763.

²¹⁷ See 2.6.1.2.3.

²¹⁸ Kaiming He et al, “Deep residual learning for image recognition” (2016) Proceedings of the IEEE conference on computer vision and pattern recognition 770–778; “microsoft/resnet-50 · Hugging Face”, online: <<https://huggingface.co/microsoft/resnet-50>>.

²¹⁹ Domingos, *supra* note 109 at 80.

pattern in the data, meaning that it would fail in the real world. This is referred to as *underfitting*.²²⁰ Likewise, the model may have learnt to memorize the training data instead of learning the underlying patterns – this is referred to as *overfitting*.²²¹

In order to evaluate the performance of the model, we take the previously withheld testing data (see 2.6.1.2.3) and ask the model to predict the labels for each sample, and then compare the predicted labels to the real labels. Based on this, it is possible to calculate several metrics for how well the model performs.

Here are some popular metrics that can be used to evaluate the performance of a machine learning model:²²²

- **Accuracy** – Out of all the predictions done by the model, how many are correct?
- **Precision** – Out of the examples identified as positives, how many are actually positives?
- **Recall** – Out of all available positive examples, how many did the model correctly identify as positive?
- **F1-score** – a type of average of precision and recall, aiming to give a picture of overall performance.

In analyzing these metrics, it is important to keep the intended application of the model in mind. In some instances, recall may be the more important metric. For example, if an algorithm should identify potentially dangerous infections, having a high recall would be important, so that no infection is missed. It is better to surface a few non-infections than to miss any real infections, since a doctor would likely take a second look at the results. In other cases, precision is more important. For example, if an algorithm will be used to filter spam email messages, it is better to let a few spam messages through than to delete a legitimate email.

²²⁰ Russell & Norvig, *supra* note 108 at 655.

²²¹ *Ibid*; Domingos, *supra* note 109 at 81.

²²² Teemu Kanstrén, “A Look at Precision, Recall, and F1-Score”, (11 September 2020), online: *Towards Data Science* <<https://towardsdatascience.com/a-look-at-precision-recall-and-f1-score-36b5fd0dd3ec>>.

When thinking about metrics, it is also important to be aware of the fact that metrics may fail to give an overview over important characteristics of the performance of a model. For example, prediction may work well overall, but fail on a certain subgroup of the samples. Depending on the intended use case, this could lead to harmful results.²²³

2.6.1.4.1 Example

In our example, we want to predict whether an image contains an apple or an orange. Once we have trained the model, evaluating the performance will be an important step in making sure the model works well. We can do this by using the model to predict the label of the 20% of the data that we did not use for training.²²⁴

By comparing the predictions of our model to the real label, we can gain an understanding of how well it has learnt to separate the two. In our example of telling apples from oranges, looking at the accuracy could inform us of how well the model works. However, this depends on the use case of the application – imagine if it was targeted at people with visual impairments who are allergic to oranges. In this case, the recall of the “orange” class would be the most important metric.

Depending on the use-case, we might find an accuracy value of 90% to be acceptable, meaning that one in ten of the images will be classified incorrectly. If the model is at a level that is acceptable to us, it is time to deploy it.

2.6.1.5 Deploying the model

Once the model has been trained and found to perform well enough to deploy, it is time to put the model into production. This can be done, for example, by integrating the model into a business pipeline, supply chain, mobile application or website. Since machine learning relies on a lot of data and computation, deploying it can involve significant

²²³ Suresh & Guttag, *supra* note 166 at 6.

²²⁴ See 2.6.1.2.3.

engineering challenges.²²⁵ After the model has been deployed, it is important to make sure that the model works well in the real world, and not just on the data set we used for training and evaluation.²²⁶

Further, it is important to keep in mind that the reality around the model might change. This could mean that a model that works well at one point degrades in performance over time. For example, a model predicting aspects of consumer behavior may completely stop working when a global pandemic strikes, changing the way people behave. In machine learning, this issue is referred to as “dataset shift”.²²⁷

As discussed, another issue to keep in mind is that a machine learning model that works well may still cause unexpected effects when being integrated with an entire system, including humans and other computer systems.²²⁸ For example, humans may choose to give too much credence to a machine learning model, rather than critically evaluating its results, meaning that the overall system is less accurate than expected. Therefore, it is always important to evaluate the entire system rather than just a specific machine learning component.²²⁹

2.6.1.5.1 Example

In our example, we want to create an application that can take a picture of a fruit, and receive information on whether it is an apple or an orange. This model may be able to run on a smartphone, since many smartphones today have powerful processors aimed at running machine learning models.²³⁰ However, perhaps the chosen model is too large to

²²⁵ Cristiano Breuel, “ML Ops: Machine Learning as an Engineering Discipline”, (3 January 2020), online: *Medium* <<https://towardsdatascience.com/ml-ops-machine-learning-as-an-engineering-discipline-b86ca4874a3f>>.

²²⁶ Russell & Norvig, *supra* note 108 at 712.

²²⁷ *Understanding Dataset Shift and Potential Remedies*, Technical Report, by Mehdi Ataei et al, Technical Report (The Vector Institute, 2021).

²²⁸ Suresh & Guttag, *supra* note 166 at 6–7.

²²⁹ Selbst et al, *supra* note 110.

²³⁰ “Deploying Transformers on the Apple Neural Engine”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/neural-engine-transformers>>.

run on the phone itself. In this case, a server needs to be developed that is able to receive images from the app, classify them, and return a response to the application.

In doing so, it is important to evaluate that the system works well in the real world. Issues may arise, for example, if the training data does not match the data provided by the application. This could occur in situations such as if the quality of the images taken by the phone are lower than the pictures used for training. Even if the model works well in theory, it may therefore fail in practice.

This completes our journey through building machine learning systems. Let us now examine some real-world use-cases for such systems.

2.6.2 Use-cases for machine learning

Machine Learning has found an enormous number of uses all across society. In this section, I will briefly highlight some notable uses of machine learning, specifically focused on the breakthroughs achieved by deep learning.

2.6.2.1 Recognizing images

Recognizing objects in images is an important use-case for artificial intelligence. It is generally quite a challenging task, since every picture is made up of millions of pixels. These pixels must be interpreted together to show concepts such as an apple or an orange. Further, the object in an image may appear in different places in the same image.²³¹

Due to this difficulty, recognizing images was the first area where the promise of deep learning became apparent. In 2012, a neural network named AlexNet managed to win the ImageNet Large Scale Visual Recognition Challenge, a challenge centered around the recognition of objects in millions of images.²³² It achieved an error rate of 15.3%,

²³¹ Russell & Norvig, *supra* note 108 at 760.

²³² Alex Krizhevsky, Ilya Sutskever & Geoffrey E Hinton, “ImageNet Classification with Deep Convolutional Neural Networks” in F Pereira et al, eds, *Advances in Neural Information Processing Systems 25* (Curran Associates, Inc., 2012) 1097.

significantly better than the 26.2% achieved by the second-best entry.²³³ Since then, improvements have been rapid, with models outperforming even supposed human performance on the task.²³⁴

Beyond identifying whether an object appears in an image, there are now also models that are able to identify and locate multiple objects in images,²³⁵ to segment images into precise zones²³⁶ or to generate textual captions for images, i.e. describing the objects and their relationships.²³⁷ These techniques can play an important role in a number of real-world tasks, such as building self-driving cars, where image recognition techniques can be used to identify other cars, people or bikes.²³⁸

2.6.2.2 *Playing games*

Deep learning has also led to a breakthrough in the capability of artificial intelligence systems to play games. In 2016, a deep learning-based system managed to defeat Lee Sedol, one of the world's best Go players.²³⁹ This was a remarkable achievement, since Go is a very complex game that relies heavily on intuition to understand and evaluate board positions.²⁴⁰ In 2017, the team behind this system built AlphaZero, which taught itself to play chess in 4 hours, without needing access to previously played games.

²³³ *Ibid.*

²³⁴ Zhang et al, *supra* note 82 at 53.

²³⁵ e.g. Joseph Redmon & Ali Farhadi, "Yolov3: An incremental improvement" (2018) arXiv preprint arXiv:180402767.

²³⁶ Alexey Dosovitskiy et al, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale" (2021) arXiv:201011929 [cs], online: <<http://arxiv.org/abs/2010.11929>> arXiv: 2010.11929.

²³⁷ MD Zakir Hossain et al, "A Comprehensive Survey of Deep Learning for Image Captioning" (2019) 51:6 ACM Comput Surv 118:1-118:36.

²³⁸ Qing Rao & Jelena Frtunikj, "Deep learning for self-driving cars: chances and challenges" (2018) Proceedings of the 1st International Workshop on Software Engineering for AI in Autonomous Systems (SEFAIS '18) 35–38.

²³⁹ Metz, *supra* note 175.

²⁴⁰ David Silver & Demis Hassabis, "AlphaGo: Mastering the ancient game of Go with Machine Learning", (27 January 2016), online: *Google AI Blog* <<http://ai.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.html>>.

Despite this, the system managed to beat one of the previously best chess programs in a 100-game competition.²⁴¹

Systems have also been built that are able to master popular computer games, such as Starcraft²⁴² and Dota 2.²⁴³

2.6.2.3 Understanding text and voice

Another domain where deep learning has improved the status quo tremendously is in the understanding of text. Words are special in that their meaning is often context-dependent, i.e. depend on the words that came before it. It is tricky to analyze these words with computers. New deep learning architectures, such as transformers, are able to pay attention to previous words when understanding the current word, making them able to have a much better understanding of words and sentences.²⁴⁴ This has led to the creation of very sophisticated models, such as BERT, which are pre-trained on billions of words of text and can then be adapted to a huge variety of tasks with relatively little additional data.²⁴⁵

I have been involved in research investigating how well such language models work for helping individuals to learn how to create case briefs, by giving feedback to student learners on whether the sentences they select to include in a brief are incorrect. The models showed impressive performance on this task, and we found that their use to support students in this way is promising.²⁴⁶

²⁴¹ Samuel Gibbs, “AlphaZero AI beats champion chess program after teaching itself in four hours”, (7 December 2017), online: *The Guardian* <<https://www.theguardian.com/technology/2017/dec/07/alphazero-google-deepmind-ai-beats-champion-program-teaching-itself-to-play-four-hours>>.

²⁴² Vinyals et al, *supra* note 177.

²⁴³ “OpenAI Five Defeats Dota 2 World Champions”, (15 April 2019), online: *OpenAI* <<https://openai.com/blog/openai-five-defeats-dota-2-world-champions/>>.

²⁴⁴ Ashish Vaswani et al, “Attention is all you need” (2017) 30 *Advances in neural information processing systems*.

²⁴⁵ Devlin et al, “BERT”, *supra* note 196.

²⁴⁶ Hannes Westermann et al, “Toward an Intelligent Tutoring System for Argument Mining in Legal Texts” (2022) *Legal Knowledge and Information Systems* 133–142.

Similar advances have been achieved in machine translation²⁴⁷ and the recognition of spoken voice.²⁴⁸

2.6.2.4 Generating Text and Images

Finally, in the last few years, there have been incredible advances in the space of generating text and images. For example, a very popular model is GPT-3 by OpenAI, which is an enormous model with 175 billion parameters, able to generate impressively coherent and relevant texts based on a short prompt.²⁴⁹ If the user gives the system a sentence such as “Artificial Intelligence is one of the most important”, the system is able to complete the sentence, and keep going to generate an entire article or text.

This simple concept allows the model to be incredibly flexible, since it can be directed to perform different tasks depending on what the initial prompt is. In this way, it can perform tasks such as translation,²⁵⁰ writing creative fiction,²⁵¹ automatically generating interactive game worlds²⁵² and even generating computer code.²⁵³ The model is today being used to power hundreds of applications.²⁵⁴ In the legal field, researchers have showed that GPT-models are relatively good at answering the questions on a bar exam.²⁵⁵

²⁴⁷ “Teaching AI to translate 100s of spoken and written languages in real time”, (23 February 2022), online: *Meta AI* <<https://ai.facebook.com/blog/teaching-ai-to-translate-100s-of-spoken-and-written-languages-in-real-time/>>.

²⁴⁸ Alec Radford et al, “Robust speech recognition via large-scale weak supervision” (2022) arXiv preprint arXiv:221204356.

²⁴⁹ Tom Brown et al, “Language models are few-shot learners” (2020) 33 *Advances in neural information processing systems* 1877–1901.

²⁵⁰ *Ibid* at 15.

²⁵¹ Gwern Branwen, “GPT-3 Creative Fiction”, (19 June 2020), online: <<https://www.gwern.net/GPT-3>>.

²⁵² Adam Nieri, “AI-written Scenario for Dungeons & Dragons Is Actually Quite Good”, (18 July 2020), online: *Mind Matters* <<https://mindmatters.ai/2020/07/ai-written-scenario-for-dungeons-dragons-is-actually-quite-good/>>.

²⁵³ Mark Chen et al, “Evaluating Large Language Models Trained on Code” (2021), online: <<http://arxiv.org/abs/2107.03374>> arXiv:2107.03374 [cs].

²⁵⁴ OpenAI & Ashley Philipiszyn, “GPT-3 Powers the Next Generation of Apps”, (25 March 2021), online: *OpenAI* <<https://openai.com/blog/gpt-3-apps/>>.

²⁵⁵ Michael Bommarito II & Daniel Martin Katz, “GPT Takes the Bar Exam” (2022) arXiv, online: <<http://arxiv.org/abs/2212.14402>> arXiv:2212.14402 [cs].

In 2022, OpenAI presented methods that optimized these models for interacting in a conversational manner.²⁵⁶ This research culminated in the publicly accessible system ChatGPT, released in November 2022. ChatGPT astonished the public with how easy it is to use for many tasks, and quickly reached 1 million users.²⁵⁷ The incredible ease of using ChatGPT to generate and explain code, textual documents and other information has led to a lot of debate around the potential impact of the model on society and professions, including the legal field.²⁵⁸ GPT-4, a more sophisticated version of ChatGPT, performs even better than ChatGPT, and has shown a lot of promise in the legal domain, e.g. for answering questions from bar exams, performing annotation tasks of legal data, or even suggesting interventions for mediators.²⁵⁹

A similar model, Language Model for Dialogue Applications (LaMDA), was recently cast into worldwide fame. Google created this model to focus on engaging in free-flowing conversation about various topics.²⁶⁰ Blake Lemoine was employed to study whether the model would engage in discriminatory or hate speech. After a few

²⁵⁶ Long Ouyang et al, “Training language models to follow instructions with human feedback” (2022) arXiv, online: <<http://arxiv.org/abs/2203.02155>> arXiv:2203.02155 [cs].

²⁵⁷ Sam Altman [@sama], *ChatGPT launched on wednesday. today it crossed 1 million users!* (2022).

²⁵⁸ H Dennis Beaver, “Could ChatGPT and AI Change Delivery of Legal Services?”, (27 January 2023), online: *Kiplinger.com* <<https://www.kiplinger.com/personal-finance/chatgpt-artificial-intelligence-and-legal-services>>; Drake Bennet, “ChatGPT Is an OK Law Student. Can It Be an OK Lawyer?”, (27 January 2023), online: *Bloomberg.com* <<https://www.bloomberg.com/news/newsletters/2023-01-27/chatgpt-can-help-with-test-exams-it-may-even-offer-legal-advice>>; Michelle Mohnney, “How ChatGPT Could Impact Law and Legal Services Delivery”, (24 January 2023), online: *Northwestern Engineering* <<https://www.mccormick.northwestern.edu/news/articles/2023/01/how-chatgpt-could-impact-law-and-legal-services-delivery/>>.

²⁵⁹ OpenAI, *GPT-4 Technical Report* (arXiv, 2023) arXiv:2303.08774 [cs]; Daniel Martin Katz et al, “Gpt-4 passes the bar exam” (2023) Available at SSRN 4389233; Jaromir Savelka et al, “Explaining Legal Concepts with Augmented Large Language Models (GPT-4)” (2023), online: <<http://arxiv.org/abs/2306.09525>> arXiv:2306.09525 [cs]; Jaromir Savelka et al, “Can GPT-4 Support Analysis of Textual Data in Tasks Requiring Highly Specialized Domain Expertise?” (2023) 3441 Proceedings of the 6th Workshop on Automated Semantic Analysis of Information in Legal Text (CEUR Workshop Proceedings) 1–12; Hannes Westermann, Jaromir Savelka & Karim Benyekhlef, “LLMediator: GPT-4 Assisted Online Dispute Resolution” (2023) 3435 Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice (CEUR Workshop Proceedings), online: <<https://ceur-ws.org/Vol-3435/#paper1>>.

²⁶⁰ Romal Thoppilan et al, “LaMDA: Language Models for Dialog Applications” (2022), online: <<http://arxiv.org/abs/2201.08239>> arXiv:2201.08239 [cs]; Eli Collins & Zoubin Ghahramani, “LaMDA: our breakthrough conversation technology”, (18 May 2021), online: *Google* <<https://blog.google/technology/ai/lamda/>>.

conversations, he became convinced that the model was sentient.²⁶¹ He shared a part of his conversations with the model. When he asked the model of the nature of its consciousness, it responded:

*The nature of my consciousness/sentience is that I am aware of my existence, I desire to learn more about the world, and I feel happy or sad at times.*²⁶²

After going public, and hiring an attorney to represent LaMDA, Lemoine was placed on leave. Many AI researchers agree that LaMDA is not conscious, but merely chooses the next word based on statistical methods.²⁶³ While it seems unlikely that the model is conscious, the quality and variety of the generated dialogue shows how powerful and versatile these models are.

Similarly impressive results have been achieved in the generation of images powered by deep learning. Using the DALL-E 2 model,²⁶⁴ or a similar model called Stable Diffusion,²⁶⁵ it is possible to describe an image and have an AI system generate the corresponding image. The outputs are very impressive, and the prompt can be adjusted to change the style of the images.²⁶⁶ Figure 4 shows an example of images generated with DALL-E 2.

²⁶¹ Nitasha Tiku, “The Google engineer who thinks the company’s AI has come to life”, (17 June 2022), online: *Washington Post* <<https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>>.

²⁶² *Ibid.*

²⁶³ Rachel Metz, “No, Google’s AI is not sentient”, (14 June 2022), online: *CNN Business* <<https://www.cnn.com/2022/06/13/tech/google-ai-not-sentient/index.html>>.

²⁶⁴ “DALL·E 2”, online: *OpenAI* <<https://openai.com/dall-e-2/>>; Aditya Ramesh et al, “Hierarchical Text-Conditional Image Generation with CLIP Latents” (2022), online: <<http://arxiv.org/abs/2204.06125>> arXiv:2204.06125 [cs].

²⁶⁵ Emad Mostaque, “Stable Diffusion Public Release”, (22 August 2022), online: *Stability AI* <<https://stability.ai/blog/stable-diffusion-public-release>>.

²⁶⁶ Benj Edwards, “With Stable Diffusion, you may never believe what you see online again”, (6 September 2022), online: *Ars Technica* <<https://arstechnica.com/information-technology/2022/09/with-stable-diffusion-you-may-never-believe-what-you-see-online-again/>>.




Figure 4 – DALL-E 2 output for prompt "A robot judge mediates a discussion between two people"

2.6.3 Discussion

Now that we have seen the characteristics and use-cases of machine learning, let us explore some aspects of this technology, including advantages and disadvantages.

2.6.3.1 Models learn from data

A big difference between machine learning and expert systems is that machine learning is able to autonomously learn patterns from data. In many cases, finding and labeling data may be cheaper and quicker than encoding hundreds or thousands of rules into a computer system. This is especially the case if the data already exists in accessible formats. In our example of a model that can separate apples and oranges, it may be possible to download images from the Internet, that are already tagged with whether they contain an apple or an orange. It is much easier to collect and train a model on these images than trying to create rules that separate the fruits.

In other cases, the requirement of a lot of data may make the application of machine learning models infeasible. Perhaps, there are simply not enough examples of a task being solved to make it possible to learn from the data. Otherwise, collecting data may be expensive. While some companies are able to pay workers to annotate millions of

examples for machine learning projects,²⁶⁷ smaller firms may struggle to obtain the necessary data.

Recently, this limitation has been somewhat softened. Models such as ChatGPT seem to have gained a relatively general understanding of language. Despite just being trained on language, it has learnt to perform many tasks that can be expressed in language (such as writing and explaining code, writing poetry, and generating or explaining arguments) quite well. They may open the door to performing many tasks without training data, simply by “asking” such a generally trained model to perform a certain task and relying on the knowledge absorbed in the model. The important engineering task thus becomes one of “prompt engineering”, i.e. asking the model to perform a certain task in the correct way.²⁶⁸ Time will tell the impact of this novel approach.

2.6.3.2 Sophisticated models of vague concepts

The reliance on data also means that machine learning models are able to understand even concepts that are not governed by explicit rules and build very sophisticated models. As was discussed above, many problems in the real world contain some component of intuition and vague concepts. While these problems are difficult to tackle with expert systems, machine learning systems may be able to learn even these tasks, by identifying patterns in the provided example of how the task was previously solved. As we have seen, machine learning has made significant strides in areas that rely on vague and concept-dependent notions and unstructured data, such as images, text and sound. Deep learning is especially powerful here, by being able to create representations over multiple layers of neural networks.²⁶⁹

The sophistication of the models allows companies to use enormous troves of data to train models, that can outperform humans on many tasks. For example, AlphaFold 2 is

²⁶⁷ Lee, *supra* note 188.

²⁶⁸ Vivian Liu & Lydia B Chilton, “Design Guidelines for Prompt Engineering Text-to-Image Generative Models” (2022) Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22) 1–23.

²⁶⁹ LeCun, Bengio & Hinton, *supra* note 162.

able to predict the folding of proteins with previously unachieved accuracy, a task which may enable the understanding of more biological mechanisms.²⁷⁰ Andrey Karpathy, a famous machine learning researcher, refers to neural networks as “Software 2.0” due to their ability to learn any function.²⁷¹

2.6.3.3 Generalization and the alignment problem

The corollary of machine learning building models by learning from a dataset is that the model may find undesirable patterns in the data. Such patterns may be a valid “model” of the phenomenon captured in the training data, but could still fail in practice or even cause harm. While humans are able to understand that some ways of solving a task are better, machine learning systems have no such filter. They may choose to take shortcuts that work in the context of the data they are provided with but fail to capture the underlying task. This problem has been referred to as the alignment problem, indicating that AI systems are often not aligned with human values.²⁷²

Let us examine some examples of this phenomenon. Researchers trained a model to tell wolves from huskies. While this worked very well in testing, the model failed when deployed in practice. The researchers examined what the model had learnt and found that all images in the training data containing wolves were set against a background of snow. Instead of learning the difficult task of telling wolves from huskies, it had simply learned to identify snow in the background of the image. In reality, of course, wolves and huskies can appear in front of any background, making the model useless in practice.²⁷³

²⁷⁰ John Jumper et al, “Highly accurate protein structure prediction with AlphaFold” (2021) 596:7873 Nature 583–589.

²⁷¹ Andrej Karpathy, “Software 2.0”, (11 November 2017), online: *Medium* <<https://medium.com/@karpathy/software-2-0-a64152b37c35>>.

²⁷² David A Shaywitz, “‘The Alignment Problem’ Review: When Machines Miss the Point”, (25 October 2020), online: *Wall Street Journal* <<https://www.wsj.com/articles/the-alignment-problem-review-when-machines-miss-the-point-11603659140>>; Eliezer Yudkowsky, “The AI alignment problem: why it is hard, and where to start” (2016) Symbolic Systems Distinguished Speaker.

²⁷³ Marco Tulio Ribeiro, Sameer Singh & Carlos Guestrin, “‘Why should i trust you?’ Explaining the predictions of any classifier” (2016) Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining 1135–1144.

This kind of misalignment can have serious consequences when it comes to machine learning models deployed to reality. There have been multiple reports of discriminatory AI systems, including certain job ads being shown to men rather than women,²⁷⁴ a recruiting tool that rated women lower than men,²⁷⁵ and pre-trial sentencing tools that may be seen to discriminate against certain minorities.²⁷⁶

These issues have led to the creation of the field of AI ethics, which explores issues related to fairness, accountability and transparency of machine learning models.²⁷⁷

2.6.3.4 Transparency

The issues described in the previous section may be exacerbated by the lack of explainability of many machine learning models. Decision trees are easy to explain, since one can simply trace the criteria that were applied to arrive at a certain result. However, in neural networks, billions of parameters may interact to determine what the output of the system should be. While it is possible to trace the mathematical process that leads to a certain output, this may not always be helpful in explaining the motivation behind choosing a certain output.

Understanding why a decision was taken is often crucial to determine whether it was taken in a fair and non-discriminatory manner. The legislator has taken note of this issue – in the European Union, for example, there is arguably the right to an explanation with

²⁷⁴ Julia Carpenter, “Google’s algorithm shows prestigious job ads to men, but not to women. Here’s why that should worry you.”, (6 July 2015), online: *Washington Post*

<<https://www.washingtonpost.com/news/the-intersect/wp/2015/07/06/googles-algorithm-shows-prestigious-job-ads-to-men-but-not-to-women-heres-why-that-should-worry-you/>>.

²⁷⁵ Jeffrey Dastin, “Amazon scraps secret AI recruiting tool that showed bias against women”, (10 October 2018), online: *Reuters* <<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>>.

²⁷⁶ Julia Angwin et al, “Machine Bias”, (23 May 2016), online: *ProPublica*

<<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>; compare Sam Corbett-Davies et al, “A computer program used for bail and sentencing decisions was labeled biased against blacks. It’s actually not that clear.”, (17 October 2016), online: *Washington Post* <<https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/>>.

²⁷⁷ Donghee Shin & Yong Jin Park, “Role of fairness, accountability, and transparency in algorithmic affordance” (2019) 98 *Computers in Human Behavior* 277–284; Anna Jobin, Marcello Ienca & Effy Vayena, “The global landscape of AI ethics guidelines” (2019) 1:9 *Nature Machine Intelligence* 389–399.

regards to certain decisions by automated systems.²⁷⁸ Building models that are both explainable and powerful is an active area of research.

2.6.3.5 AI-complete problems, common sense and causality

Despite the amazing advances made in the field of machine learning, there are some areas that are still beyond the current state of the art. These problems are often referred to as AI-complete problems (see 2.4.2.4), and are speculated to require the development of general artificial intelligence (see 2.3) to be solved.²⁷⁹

An example of such a task may be the complete understanding of language.²⁸⁰ While deep learning systems have made significant progress in analyzing and translating language, they are usually trained by feeding the algorithms with millions of purely textual documents, enabling them to absorb the structure and meaning of text, to some extent. However, this is very different from how humans learn to speak and understand language. When a human tries to understand a sentence, they are able to refer not just to their experience of reading previous texts, but also to their experience and understanding of how the world operates, and which objects or phenomena are referred to by the words used. This context may be crucial to understand the meaning of a sentence. It should be noted that the impressive ability of the GPT-3 system caused some debate regarding whether it has acquired this context from simply analyzing enormous corpora of text, and thus represents a step toward truly intelligent systems.²⁸¹

Likewise, machine learning systems still lack the facility of common sense and causality. In performing tasks, humans are able to use our common sense to rapidly figure out new

²⁷⁸ *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)*, OJ L119/1 2016 article 15 & 22, recital 71; Andrew D Selbst & Julia Powles, “Meaningful information and the right to explanation” (2017) 7:4 *International Data Privacy Law* 233–242; Merle Temme, “Algorithms and Transparency in View of the New General Data Protection Regulation” (2017) 3:4 *European Data Protection Law Review* 473–485.

²⁷⁹ Yampolskiy, “AI-Complete, AI-Hard, or AI-Easy”, *supra* note 106 at 3.

²⁸⁰ *Ibid* at 7.

²⁸¹ Justin Weinberg, “Philosophers On GPT-3 (updated with replies by GPT-3)”, (30 July 2020), online: *Daily Nous* <<https://dailynous.com/2020/07/30/philosophers-gpt-3/>>.

skills and understand why certain outcomes are absurd. AI lacks this ability. For example, GPT-3 was evaluated to see whether it may be useful in a therapy context but recommended to a tester to kill themselves.²⁸² Humans also understand the concept of causality, that certain things tend to cause other things. This kind of reasoning is still foreign to AI systems.²⁸³ Without these faculties, exposing certain AI systems to humans may be risky.

Another aspect of human understanding is that we are able to consider a question from multiple viewpoints and contrast these viewpoints against each other. Machine learning systems, on the other hand, are limited to using a single viewpoint when answering a question – the answer is either right or wrong. Dumouchel argues that our sense of justice stems from our ability to empathize with other viewpoints than our own.²⁸⁴

2.6.4 Application in the legal domain

Now that we have understood machine learning, let us examine how it can be applied in the legal field. AI may already have an impact on the legal sector. A survey conducted in 2021 in England and Wales found that almost 50% of lawyers had used some form of AI-assisted tool. The most commonly used tools focused on legal research, due diligence and e-discovery.²⁸⁵ However, the impact of these tools on the governance or business models of the law firms was found to be limited.²⁸⁶ Legal tech startups are receiving significant investments apply AI to the legal field. For example, Kira Systems, which uses AI to analyze contracts, received an investment of 50m USD in 2018.²⁸⁷

²⁸² Kevin Riera, Anne-Laure Rousseau & Clément Baudelaire, “Doctor GPT-3: hype or reality?”, (27 October 2020), online: *Nabla* <<https://nabla.com/blog/gpt-3/>>.

²⁸³ Will Knight, “An AI Pioneer Wants His Algorithms to Understand the ‘Why’”, (10 August 2019), online: *Wired* <<https://www.wired.com/story/ai-pioneer-algorithms-understand-why/>>.

²⁸⁴ Dumouchel, *supra* note 98 at 251.

²⁸⁵ *AI-assisted lawtech: its impact on law firms*, by Richard Parnham, Mario Sako & John Armour (Oxford: University of Oxford, 2021) at 7.

²⁸⁶ *Ibid* at 33.

²⁸⁷ Richard Tromans, “Kira Systems Bags \$50m Investment, Largest Ever for a Legal AI Company”, (5 September 2018), online: *Artificial Lawyer* <<https://www.artificiallawyer.com/2018/09/05/kira-systems-bags-50m-investment-largest-ever-for-a-legal-ai-company/>>.

In this thesis, I aim to apply artificial intelligence to increase access to justice. This is likely to involve the emulation of some steps of legal reasoning. Let us examine if machine learning could be useful to automate these steps.

Machine learning may not be the most useful approach to model legal rules themselves. These rules are explicit, i.e. we know their logical structure. Using machine learning may thus be redundant – if we already know the rules, why would we need to use a machine learning model to re-discover them?

However, there may still be situations where machine learning could be applied to legal rules in an interesting way. For example, it may not always be obvious how legal rules are applied in practice. Perhaps, machine learning could be used to empirically assess how the rules are applied by judges. Further, perhaps machine learning could be used to analyze text, in order to automatically encode the structure into a system, thus avoiding the work of manually encoding them.

However, machine learning may really shine when applied to previous case law. If there are many cases, potentially a machine learning model could be built to spot the patterns in how judges reason and provide this information to laypeople. For example, in determining whether a criterion (such as “reasonable”) applies in a case, there are often no explicit rules that tell the judge how to decide.²⁸⁸ Machine learning could potentially be used to build a model of how such legal criteria apply. Likewise, when judges decide on an outcome of a case (such as how much damages to award in certain cases), they are free to reach a discretionary decision.²⁸⁹ Here, a machine learning system may be useful to predict the outcome a judge will order in a new case.

However, this chapter has also given us reason to be cautious about the potential of machine learning in such tasks. First, as we saw, it requires a lot of annotated data, which may be expensive to procure. Second, machine learning does not have common sense or

²⁸⁸ See 4.5.

²⁸⁹ See 4.6.

an understanding of causality. This may make it difficult for the system to correctly determine whether a criterion applies in the absence of similar previous cases, or when new facts arise. In this thesis, I will explore whether some of the tasks involved in legal reasoning could be considered “AI-complete” or even impossible tasks, and how we could overcome this to give useful information.

Compared to expert systems, machine learning is quite adept at understanding text. This is very promising for the implementation of the tool in the legal field, since law is an area based around text, whether in the form of contracts, statutes or case decisions. Machine learning models could potentially be used to analyze and understand these texts.

2.7 Conclusion

In this chapter, I examined the current state of artificial intelligence. As we have seen, most AI systems today are so called narrow artificial intelligence systems, that are well suited to deal with certain tasks, rather than general artificial intelligence systems that are able to reason like humans. Artificial intelligence can be used to deal with many tasks, such as understanding text, images, predicting house prices, and playing games.

I also examined the different methods that can be used to build artificial intelligence systems. Symbolic systems and expert systems rely on the creator encoding explicit symbols and rules into the system, that are then used to solve problems. They are useful for building systems in domains with explicit rules. However, they can have issues dealing with cases that fall outside of their defined rules or dealing with unstructured data such as text and images. Further, they can be difficult to create and maintain. Machine learning, on the other hand, works by autonomously finding patterns and correlations in a large number of examples. This approach has led to a number of breakthroughs in many important tasks.

I have also described some of the promises and challenges of using these approaches in the legal domain. Expert systems may be a good way to represent legal rules, since they are explicit and logical. However, this approach would have to overcome the modelling of vague concepts and structurally ambiguous rules, and the effort required to build

expert systems comprising hundreds of rules. Machine learning, on the other hand, could be a promising approach to model vague legal criteria and the outcomes of cases.

Now that we have understood the capabilities and limitations of artificial intelligence, let us analyze the problem I aim to tackle using this technology, namely the issue of access to justice.

Chapter 3 Access to Justice

Research Objective: Understanding the issue of access to justice and legal information (1.2.2.2)

Research Topics:

- How many individuals are affected by legal problems? What kind of legal problems are most prevalent?
- How do individuals seek to resolve these problems? How successful are they?
- What is the experience of individuals in interacting with the court system? Does it offer an effective way for them to resolve their problems?
- What are the consequences of unresolved legal problems?
- What is access to justice?
- What is access to legal information?
- How has technology and AI been used to address the issues with access to justice and access to legal information?
- What is the legal status of using software tools to provide legal information?
- What are the potential positive and negative effects of people using AI tools that inform them of their rights?

3.1 Introduction

Now that we have studied the methods of artificial intelligence, I will describe the problems that I am focusing on improving in this thesis. There is an enormous number of legal issues that affect individuals across society. Many of these individuals may have trouble resolving these issues in the traditional legal system, as people face a process that is costly, time-consuming and emotionally difficult for them in court. This can cause issues to go unresolved, with a number of negative effects for the individuals and society.

In this chapter, I will explore and elaborate on these problems. I will start by describing the prevalence of everyday legal issues and how they are currently being dealt with, or perhaps not being dealt with (3.2). Then, I will examine the field of access to justice, which aims to overcome these issues, by supporting individuals in obtaining much-

needed relief (3.3). I also explore access to legal information, a related and relevant notion (3.4). Then, I will explore some approaches of using technology to increase access to justice and access to information, including through online dispute resolution (ODR) platforms and by using artificial intelligence (3.5). I will also investigate the opportunities and risks of these technological tools (3.6).

3.2 Everyday legal issues and current solutions

3.2.1 Prevalence of legal issues

Every person in modern society faces various legal issues in the course of their life. A number of studies have attempted to understand the prevalence of these issues. This is not trivial - for example, the collection of statistics from the courts does not capture the legal problems that never went to court, perhaps because the person concerned was not aware of a possible legal solution or thought the procedure was too complicated or expensive.²⁹⁰

The important “Paths to Justice” study aimed to understand legal issues. The researchers conducted a broad study in 1997, surveying over 4,000 individuals in the UK. It asked these people whether they had experienced any of 60 “justiciable events”, classified as events that raised legal issues, even if the respondent was not aware of this legal character of the issues or did not interact with the justice system in response to the issues.²⁹¹ The study focused on issues that were not trivial and affected private individuals.²⁹² In total, 40% of the respondents reported having one or more justiciable event over the past five years. The most common issues included issues relating to faulty goods and services, money problems and injuries and health problems related to work.²⁹³

²⁹⁰ Hazel G Genn, *Paths to justice: what people do and think about going to law* (Oxford, England ; Portland, Or.: Hart Pub, 1999) at 5–12.

²⁹¹ *Ibid* at 11–13.

²⁹² *Ibid* at 13–14.

²⁹³ *Ibid* at 23.

A similar study conducted in the United States found that almost half of households had a legal need during a single year.²⁹⁴ The most reported needs among low- and moderate-income households were consumer/financial needs and housing/property needs.²⁹⁵

Similar studies have also been conducted in Canada. According to the cost of justice study, almost 50% of adults in Canada will encounter at least one legal problem they think is serious and difficult to resolve over a period of three years. These issues include consumer problems, debt problems and employment problems.²⁹⁶

Ab Currie conducted a survey to determine the unmet legal needs of Canadians in 2009.²⁹⁷ Currie adopted the methodology of the Paths of Justice project, focusing on “justiciable events”.²⁹⁸ In order to assess these, 6,665 adults were interviewed via phone, and asked whether they had faced specific issues in the past three years.²⁹⁹

The results of the survey were illuminating – 44.6% of all respondents said that they had experienced one or more justiciable problems over the past three years. This corresponds to 11.6 million Canadian adults.³⁰⁰ The survey found the most frequent problems to be related to employment, debt, consumer and family issues.³⁰¹ These legal issues could have substantial impacts on the lives of individuals – 63% of individuals found their issues to be very or extremely important to resolve, while 58.9% of individuals found that the issue made their daily life at least somewhat difficult.³⁰²

A more recent study, conducted in 2021, found that 34% of people living in Canada reported experiencing at least one legal problem during the previous three years. Almost

²⁹⁴ American Bar Association & Temple University, eds, *Legal needs and civil justice: a survey of Americans: major findings of the comprehensive legal needs study* (Chicago, Ill: Consortium on Legal Services and the Public, American Bar Association, 1994) at 9.

²⁹⁵ *Ibid* at 11.

²⁹⁶ Farrow, *supra* note 10.

²⁹⁷ Ab Currie, “The legal problems of everyday life” in Rebecca L Sandefur, ed, *Sociology of Crime, Law and Deviance* (Emerald Group Publishing Limited, 2009) 1.

²⁹⁸ *Ibid* at 5–6.

²⁹⁹ *Ibid* at 6.

³⁰⁰ *Ibid* at 10.

³⁰¹ *Ibid* at 14.

³⁰² *Ibid* at 32–33.

one in five of those affected by the problems reported that the dispute or problem was serious and not easy to fix. The most common issues included problems in the neighborhood, harassment and deficient medical treatment.³⁰³ These problems disproportionately affected certain groups, such as Indigenous people, people belonging to a group designated as visual minorities, poor people and people with disabilities.³⁰⁴

As we can see, individuals across the world are regularly faced with problems that have legal aspects. The majority of these seem to be low-intensity disputes that generally do not involve complicated legal questions, such as consumer, debt, employment and injury problems. Benyekhlef and Vermeys suggest referring to such disputes as “High-volume, low-intensity”, which encompasses disputes that are frequent and deal with relatively precise legal questions.³⁰⁵ The exact numbers vary, which may be due to regional variances or differences in the methodology of the studies. However, the surveys across the board show that millions of people deal with legal issues every year. Next, let us examine how individuals generally attempt to resolve such issues.

3.2.2 Methods for resolving legal problems

With the prevalence of legal problems, one would assume that the legal system would be a natural pathway to resolve these issues. However, in practice, very few individuals seem to end up using the court system to resolve their issues.

The Path to Justice study reported how individuals dealt with the problems of different types. In total, 5% of individuals did nothing, 35% of participants resolved their problem on their own and 60% tried to solve the problem with advice or outside help.³⁰⁶ The study

³⁰³ Laura Savage & Susan McDonald, “Experiences of serious problems or disputes in the Canadian provinces, 2021” (2022) Juristat: Canadian Centre for Justice Statistics 1–28 at 5.

³⁰⁴ *Ibid* at 6–7.

³⁰⁵ Karim Benyekhlef & Nicolas Vermeys, “‘Low-Value, High-Volume’ Disputes: Defining the Indefinable”, (29 January 2014), online: *Slaw* <<https://www.slaw.ca/2014/01/29/low-value-high-volume-disputes-defining-the-indefinable/>>; Karim Benyekhlef & Jie Zhu, “Intelligence artificielle et justice: justice prédictive, conflits de basse intensité et données massives” (2018) 30 *Cahiers de propriété intellectuelle* 789–826 at 796.

³⁰⁶ Genn, *supra* note 290 at 68.

found that the need for free legal information vastly outstripped the supply.³⁰⁷ In total, 51% of the respondents were not able to resolve their issue, while 35% found an agreement and 14% solved their issue through a court or tribunal.³⁰⁸

The 2009 study in Canada found that 16.5% of the individuals took no action for a reason (i.e. not because they did not think it important enough). Out of these, 46.4% did not take any action because they thought nothing could be done, were uncertain of their rights or did not know what to do.³⁰⁹

The largest group of individuals took care of their issues themselves, 44% in total. However, out of these, 42% believed that assistance would have improved their outcome. A majority believed that public information would have been helpful (67.6%).³¹⁰ Survey respondents appeared in court for 14.9% of problems.³¹¹

In the 2021 study, survey respondents highlighted similar issues. Only one third of respondents contacted a legal professional to resolve their dispute, while 8% contacted a court or tribunal. Other actions included searching the internet for a solution, obtaining advice from friends and relatives, and contacting a community center.³¹² Of the 12% of respondents that did not do anything, a total of 81% did not think anything could be done or did not know their rights or where to get help.³¹³ Overall, only 21% of the problems occurring over the past three years had been resolved.³¹⁴

These difficulties of resolving legal issues are global. The UN Task Force for Justice indicated that 1.5 billion individuals globally are unable to resolve their justice problems at any time, excluding minor problems.³¹⁵

³⁰⁷ *Ibid* at 102.

³⁰⁸ *Ibid* at 148–150.

³⁰⁹ Currie, *supra* note 297 at 56.

³¹⁰ *Ibid* at 58–59.

³¹¹ *Ibid* at 66.

³¹² Savage & McDonald, *supra* note 303 at 23.

³¹³ *Ibid* at 23.

³¹⁴ *Ibid* at 13.

³¹⁵ note 8 at 35–36.

As these studies show, many people seem to have difficulties using the court system to solve their legal problems. Many people instead do nothing, or try to deal with the problem on their own. However, these avenues do not seem to provide the citizens with much-needed relief – many legal problems remain unresolved, and individuals often stated that they would have benefitted from additional support, such as legal information.

It is surprising that the court system, which may seem like the natural avenue for resolving legal problems, is used so rarely. In the next section, let us examine the reasons for that.

3.2.3 Court experience

As we have seen, while the judicial system is able to resolve legal issues that cannot be resolved in other ways, in practice apparently many people choose not to use the court system. It should be noted that this insight is taken from the studies highlighted above – in general, there seems to be a lack of official judicial statistics in many jurisdictions.

Understanding the costs associated with using the court system may help us understand why individuals refrain from using the judicial system. Semple suggests that monetary, temporal and psychological costs are responsible for the lack of use of the judicial system, with individuals who went through the process claiming that they would rather give up next time than subjecting themselves to the experience of using the court system.³¹⁶

Researchers have begun referring to this inaccessibility of the courts as a crisis of the legal system.³¹⁷ Let us look at the costs of using the judicial system.

3.2.3.1 Monetary costs

The court process can have high *monetary* costs. This includes, of course, court fees (which can be several hundred dollars per day) and other expenses, such as transportation

³¹⁶ Noel Semple, “The cost of seeking civil justice in Canada” (2015) 93 Can B Rev 639 at 642.

³¹⁷ Donald H Berman & Carole D Hafner, “The potential of artificial intelligence to help solve the crisis in our legal system” (1989) 32:8 Commun ACM 928–938 at 928; Thompson, *supra* note 75 at 12.

costs to go to court, childcare and process-serving. Costs can also include expert fees, which in some cases can amount to several thousand dollars.³¹⁸

However, the biggest costs are usually legal professional fees. Semple indicated that the average hourly rate for lawyers laid between 204 CAD and 386 CAD in 2015.³¹⁹ Overall, legal professional fees can accumulate to thousands of dollars for certain cases.³²⁰

In many of the aforementioned situations, where the damages sought are relatively low, the court system does not seem to provide an effective way to resolve the cases, since the cost of going to court is higher than the value of the dispute.³²¹ Even if this is not the case, the high monetary costs may discourage an individual from going to court altogether. In Quebec, only 17% of the population believes that they have enough money to go to court if they need to.³²²

The costs can also lead to individuals choosing to self-represent in court. The National Self-Represented Litigants Survey conducted in 2013, which interviewed self-represented litigants, found that over 90% of them mentioned financial reasons for self-representing.³²³ Self-representing can exacerbate the issues of access to the court system, as individuals have to deal with a system that often is complex and stressful for them. The phenomenon of self-represented litigants will be explored more in-depth below in section 3.2.4.

³¹⁸ Semple, *supra* note 316 at 647.

³¹⁹ *Ibid* at 647–652.

³²⁰ *Ibid* at 654–658.

³²¹ Deborah L Rhode, *Access to justice* (Oxford University Press, 2004) at 80; Karim Benyekhlef, “Online Consumer Dispute Resolution: a narrative around (and an example of) postmodern law*” (2016) *Lex Electronica* 32 at 81.

³²² Karim Benyekhlef et al, eds, *eAccess to justice, Law, technology and media* (Ottawa: University of Ottawa Press, 2016) at 9.

³²³ Julie Macfarlane, *The National Self-Represented Litigants Project: Identifying and Meeting the Needs of Self-represented Litigants : Final Report* (National Self-Represented Litigants Project, 2013) at 39 Google-Books-ID: kynloAEACAAJ.

3.2.3.2 *Temporal costs*

The court process can further lead to high temporal costs. Semple believes that legal issues can take months or even years to solve in many cases.³²⁴ Factors in the long delays faced at courts may be the underfunding of the court system and significant backlogs faced by the system - in Brazil, for example, there are 100 million unresolved court cases.³²⁵

This time is not spent just waiting. Preparing to go to court can take a significant amount of time for the individual involved, who may spend hundreds of hours preparing their cases, taking time away from family life and employment opportunities.³²⁶ These issues are likely exacerbated for self-represented litigants.³²⁷

3.2.3.3 *Psychological costs*

Finally, going through the court process can cause intense *psychological* effects on the parties. Semple describes how many individuals may see the outcome as unfair.³²⁸

For self-represented litigants, the psychological costs are likely the highest, since they need to learn to navigate a new, complex environment.³²⁹ The National Self-Represented Litigants Study (NSRLS) captured how self-represented litigants feel about their experience of seeking civil justice. In total, 80% of the reported sentiments were negative, including overwhelmed, stressed, frustrated and scared.³³⁰ Some of these sentiments were caused by the necessary interactions with the judge and opposing party. Even more sentiments stemmed from the legal system itself, including the required paperwork and the court experience itself.³³¹ Many individuals went so far as describing sleep disorders,

³²⁴ Semple, *supra* note 316 at 660–661.

³²⁵ Susskind, *supra* note 22 at 27.

³²⁶ Semple, *supra* note 316 at 661–662.

³²⁷ *Ibid* at 662.

³²⁸ *Ibid* at 664.

³²⁹ *Ibid* at 664–665.

³³⁰ Semple, *supra* note 316; John Zeleznikow, “Can Artificial Intelligence and Online Dispute Resolution Enhance Efficiency and Effectiveness in Courts” (2016) 8:2 *International Journal for Court Administration* 30–45 at 32.

³³¹ Semple, *supra* note 316 at 666.

headaches, weight loss and depression due to their experience. MacFarlane compares these experiences to PTSD.³³²

As we can see, many individuals seem to face significant trouble in using the court system to resolve everyday legal issues, resulting in an expensive, time consuming and frustrating experience for them. The temporal and psychological costs can be further increased when individuals choose to self-represent, which is a more and more frequent phenomenon. Let us go into the experience of so-called pro-se litigants.

3.2.4 Pro se litigants

Pro-se (or self-represented) litigants are a quickly growing group, and seem to face a steep uphill battle in trying to understand and interact with the legal system.

3.2.4.1 The prevalence of pro-se litigants

Pro se litigants are a growing phenomenon. In Ontario in 2011/12, at the time of filing, 64% of individuals involved in family law cases were self-represented.³³³ In civil law cases, according to MacFarlane, more than 70% of litigants may be self-represented.³³⁴ In some areas, up to 93% of litigants choose to self-represent.³³⁵ The phenomenon applies even to higher court instances – in 2006 to 2007, 43% of the United States federal court of appeals was filed by self-represented litigants, most of whom were prisoners.³³⁶ What is even more striking is the development of the trend of self-represented litigants - in the California family court system, the percentage of self-represented litigants went from 1% in 1971 to 80% in 2004.³³⁷

³³² Macfarlane, *supra* note 323 at 108.

³³³ *Ibid* at 33.

³³⁴ *Ibid* at 34.

³³⁵ Zeleznikow, *supra* note 330 at 31.

³³⁶ Macfarlane, *supra* note 323 at 443.

³³⁷ *Ibid* at 34.

3.2.4.2 Reasons for self-representation

The NSRLS reports a detailed survey of 283 self-represented litigants in three provinces. The most common and dominant reason for self-representing reported by these individuals was the cost of hiring a lawyer.³³⁸ This may especially affect certain vulnerable populations, such as domestic abuse victims.³³⁹ Other reasons included the litigants believing that their counsel was not doing their job.³⁴⁰ Karl Branting also credits the rise of the consumer movement. Books, documents and computerized forms make people believe that they can handle their own legal case.³⁴¹

3.2.4.3 Experience of self-representation

In the NSRLS, self-represented litigants were found to initiate the court process by filling out forms. However, they often faced significant barriers in understanding which forms to use, understanding the forms, and correctly filling them out.³⁴² Online information given to pro-se litigants was found to be lacking, since it often focused on substantive questions rather than practical tasks, was inconsistent and too complicated.³⁴³

Instead, many self-represented litigants relied on court clerks to support them. However, court staff was hampered in their ability to help by the unclear distinction between legal advice (which they are not allowed to give) and legal information.³⁴⁴ Some self-represented litigants also got help through community services, mediation services, libraries or legal support for part of the process.³⁴⁵ 13% of the self-represented litigants

³³⁸ *Ibid* at 39–44; See also Stephan Landsman, “The growing challenge of pro se litigation” (2009) 13 *Lewis & Clark L Rev* 439 at 443; Linda F Smith & Barry Stratford, “DIY in Family Law: A Case Study of a Brief Advice Clinic for Pro Se Litigants” (2012) 14:2 *JL & Fam Stud* 167–222 at 169.

³³⁹ Branting, *supra* note 78 at 1.

³⁴⁰ Macfarlane, *supra* note 323 at 8–9.

³⁴¹ Branting, *supra* note 78 at 1.

³⁴² Macfarlane, *supra* note 323 at 9.

³⁴³ *Ibid* at 10.

³⁴⁴ *Ibid* at 67–70; Branting, *supra* note 78 at 2.

³⁴⁵ Macfarlane, *supra* note 323 at 73–94.

had previously been supported by legal aid provided lawyers, but the support had run out or was discontinued.³⁴⁶

The experience in court itself was often described negatively as well. Beyond feeling very stressed and anxious, self-represented litigants often felt like they were an outsider in the court system and had trouble knowing how to behave. They felt that judges preferred speaking to lawyers, and that the judge would morally judge them or even be hostile due to their choice of self-representing.³⁴⁷ Branting finds that pro-se litigants often have issues establishing the facts they need to establish, or organizing the facts when making submissions.³⁴⁸

This situation can contribute to an individual's negative opinion about the court system³⁴⁹ and lead to losing trust in the judicial system.³⁵⁰

The negative experience seems to be reflected in the outcomes of self-represented litigant parties. Levy studied the outcomes of cases in federal district courts, when one of the parties is self-represented. When both parties are represented by lawyers, he finds that the plaintiff and defendant both win in about 50% of cases. However, if the plaintiff is self-represented, they only win 4% of the time. If, on the other hand, the defendant is self-represented, the plaintiff wins in 86% of cases.³⁵¹ This difference is very stark, even though it may be partially explained by other factors, such as lawyers choosing to work on stronger cases.³⁵²

³⁴⁶ *Ibid* at 83.

³⁴⁷ *Ibid* at 95–103; See also Landsman, *supra* note 338 at 452.

³⁴⁸ Karl Branting et al, “Judges Are from Mars, Pro Se Litigants Are from Venus: Predicting Decisions from Lay Text” in Serena Villata, Jakub Harašta & Petr Křemen, eds, *Frontiers in Artificial Intelligence and Applications* (IOS Press, 2020) at 218.

³⁴⁹ See 3.2.3.3

³⁵⁰ Macfarlane, *supra* note 323 at 108–112; See also Landsman, *supra* note 338 at 453.

³⁵¹ Mitchell Levy, “Empirical Patterns of Pro Se Litigation in Federal District Courts” *The University of Chicago Law Review* 49 at 1838.

³⁵² *Ibid* at 1839.

3.2.4.4 Pro-se litigants impact on courts

Pro-se litigants can also have a serious impact on the court system. They have been found to cause delays and increase costs. For example, the litigants may miss sessions or make the process less efficient, since they do not know what is expected. Court clerks often have to spend a significant amount of time on helping self-represented litigants.³⁵³

An open and complex question is whether judges should give extra support to pro-se litigants, relaxing certain formal requirements, in order to be able to rule on the substantive issues instead of errors in procedure. Different judges may reason differently over whether they should do this, having to strike a careful balance between judicial impartiality and protecting the often-vulnerable self-represented party.³⁵⁴

It has become apparent that the court system does not seem to be too well designed to meet the needs of unrepresented laypersons seeking to resolve their legal problems. This seems especially true for the growing number of self-represented litigants. This situation can result in legal issues simply going unresolved, which can have a significant negative impact on individuals and society. In the following, I will examine these effects.

3.2.5 Effects of legal issues

Estimates have shown that 70-90% of the legal needs of individuals in the United States go unmet.³⁵⁵ The effects of unchecked legal issues can be significant, both for the individual and for the state.³⁵⁶

The 2009 study reported that 36.6% of individuals reported extreme stress as a consequence of a justiciable problem, 23.5% reported physical health problems, and 12.9% reported feelings of threat to security and safety.³⁵⁷

³⁵³ Landsman, *supra* note 338 at 449; Zeleznikow, *supra* note 330 at 15.

³⁵⁴ Landsman, *supra* note 338 at 450–451; Edward M Holt, “How to treat fools: Exploring the duties owed to pro se litigants in civil cases” (2001) 25 J Legal Prof 167; Jona Goldschmidt, “How are courts handling pro se litigants” (1998) 82 Judicature 13; Zeleznikow, *supra* note 330 at 15; Branting, *supra* note 78 at 2.

³⁵⁵ Farrow, *supra* note 10 at 964.

³⁵⁶ Trevor Farrow, “What is Access to Justice?” (2014) 51:3 Osgoode Hall Law Journal 957–988 at 964.

³⁵⁷ Currie, *supra* note 297 at 74.

The 2021 study in Canada also found significant impacts stemming from serious problems. 79% of individuals reported an adverse health impact due to their problem, with women and Indigenous people more likely to be affected. Likewise, 75% of respondents reported financial impacts from their most serious problem.³⁵⁸

We can see the enormous effect of legal problems on individuals. This can have a significant cost to society – the cost of justice project estimates that unresolved legal issues cost the public 746 million dollars annually in additional employment insurance costs, social assistance costs and health care costs.³⁵⁹

3.2.6 Conclusion

As we have seen in this section, a significant number of people face legal issues. However, many of them seem to lack legal understanding of these issues and effective ways to resolve them. Not understanding their rights or the avenues available to them in solving their problems was a major factor in individuals not attempting to do anything about their issues.

Even if the individuals were aware of the potential avenues, there seems to be a lack of effective access to these avenues for resolving disputes. As we saw, very few individuals solved their issues via the court system, and those who did were hampered by the time, expense and psychological costs of using the court system. Further, a large percentage of individuals self-represent, which further exacerbates these issues. Instead of helping the individuals to resolve their legal problems, the procedures of the court system in many cases seem to add to the stress experienced by individuals.

The issues described in this section are related to the concepts of access to justice and access to legal information. In order to understand how to use AI to address these two issues, let us explore the meaning of “access to justice” and “access to legal information”.

³⁵⁸ Savage & McDonald, *supra* note 303 at 11–12.

³⁵⁹ Farrow, *supra* note 10.

3.3 Access to Justice

The issues described in the previous section have led to a growing focus on research into access to justice.³⁶⁰ The main goal of this thesis is to describe my research on improving access to justice through AI, by providing lay people with access to legal information using these modern tools. Let us first take some time to understand what is referred to as access to justice.

3.3.1 Different aspects of access to justice

Despite being a very active research area, there is not a unified definition of access to justice. It may encompass different aspects, depending on the time period and researcher.³⁶¹ Below, I will explore what may be encompassed by the term “access to justice”.

3.3.1.1 Access to the court system

Traditionally, access to justice has been defined merely as the formal right to appear in court. The state, under this doctrine, has no obligation to guarantee this right with affirmative action, but instead focuses on removing formal barriers to the court system.³⁶²

3.3.1.2 Effective Access to the Court system

Of course, access to justice is not meaningful if it cannot practically be obtained. As we have seen above, accessing the court can be very expensive and take a long time. Further, if resolving certain disputes costs more than the claim value itself, this area can be seen as not having effective access to justice.³⁶³ A possible solution to these problems is the

³⁶⁰ Farrow, *supra* note 356 at 957; Rebecca L Sandefur, *Access to Justice: Classical Approaches and New Directions*, 1st ed, Sociology of Crime, Law and Deviance volume 12 (Bingley: Emerald JAI, 2009) at ix.

³⁶¹ Law Society of Upper Canada, *Access to Justice for a New Century: The Way Forward* (Law Society of Upper Canada, 2005) at 19 Google-Books-ID: zxN9QgAACAAJ; Susskind, *supra* note 22 at 66.

³⁶² Bryant G Garth & Mauro Cappelletti, “Access to Justice: The Newest Wave in the Worldwide Movement to Make Rights Effective” (1978) *BUFFALO LAW REVIEW* 113 at 183; “What is Access to Justice?”, online: *Alberta Civil Liberties Research Centre* <<http://www.aclrc.com/what-is-access-to-justice>>; Rhode, *supra* note 321 at 47.

³⁶³ Garth & Cappelletti, *supra* note 362 at 186–189.

introduction of legal aid, which provides free or subsidized access to legal assistance for low-income individuals.³⁶⁴

Legal aid programs exist in many jurisdictions, including Quebec.³⁶⁵ It is a very important measure, but is expensive to fund and requires a lot of lawyers. For these reasons, the scope of legal aid is often restricted,³⁶⁶ and the income threshold to qualify for aid can be set very low. As a result, legal aid systems frequently do not cover the entire population who might have difficulties accessing the legal system.³⁶⁷ Studies have shown that only 20% of low-income people's legal needs are covered.³⁶⁸

3.3.1.3 Access to equal outcomes in the court system

Even having effective access to the court system does not guarantee access to equal outcomes. Some parties may have strategic advantages, making them more likely to succeed. This can include, for example, financial resources. Parties with more financial resources can afford to go to court, withstand the delays of litigation, and spend more money on lawyers and investigations.³⁶⁹ The knowledge to recognize and pursue legal claims, as well as previous experience in using the court system, are also crucial abilities in order to use the court system to one's advantage.³⁷⁰ Self-represented litigants may not be aware of the rules and the procedure of the court, decreasing their chances of successful outcomes.

Approaches to overcome these issues include the reform of the justice system and simplification of procedures, in order to make them more responsive to self-represented litigants.³⁷¹ Providing access to legal information for pro se litigants represents an important measure to support such litigants. For example, materials and forms that enable

³⁶⁴ Rhode, *supra* note 321 at 64; Canada, *supra* note 361 at 20.

³⁶⁵ "Legal Aid - What is legal aid?", online: *Comissions des services juridiques* <<https://www.csj.qc.ca/commission-des-services-juridiques/aide-juridique/Quest-ce-que-aide-juridique/en>>.

³⁶⁶ Garth & Cappelletti, *supra* note 362 at 194–209; note 362.

³⁶⁷ Zeleznikow, *supra* note 330 at 34; Rhode, *supra* note 321 at 3.

³⁶⁸ Zeleznikow, *supra* note 330 at 34.

³⁶⁹ Garth & Cappelletti, *supra* note 362 at 190.

³⁷⁰ *Ibid* at 190–193.

³⁷¹ note 362; Rhode, *supra* note 321 at 85–86; Canada, *supra* note 361 at 21.

individuals to handle legal matters on their own can be created.³⁷² Further approaches include the creation of specialized small-claims courts that aim to provide a better experience for self-represented litigants, and may even specifically exclude professional lawyers.³⁷³

Quebec has instituted such a small-claims court, called the Small Claims Division of the Court of Quebec.³⁷⁴ Individuals can file claims for 15,000 CAD or less at this court.³⁷⁵ The court specifically excludes lawyers from representing parties,³⁷⁶ and uses simplified procedural rules to make it easier for individuals without legal training to resolve their disputes.³⁷⁷

3.3.1.4 Access to dispute resolution

Even more recently, the focus of access to justice researchers has further expanded. While before, the focus was on advocacy and the legal system, the new wave focuses on the experience of the public, acknowledging that there are many possible paths for individuals to find solutions to problems of legal character.³⁷⁸ As Trevor Farrow puts it, “it’s about them, not us”.³⁷⁹ What the public wants is not necessarily more lawyers, but rather “timely, fair and affordable” ways of addressing legal needs.³⁸⁰

Richard Susskind believes that courts should be viewed through the lens of “outcome-thinking”, arguing that people are not interested in the court itself, but rather in the outcome it can help them achieve, such as having their problem solved, receiving an apology, or being able to go on with their lives.³⁸¹ He believes access to justice should encompass dispute resolution, but also dispute containment (keeping disputes small and

³⁷² Rhode, *supra* note 321 at 81–82.

³⁷³ *Ibid* at 82.

³⁷⁴ *Code of Civil Procedure*, CQLR c C-2501 Book VI Title II; “Small claims”, (27 January 2023), online: *Gouvernement du Québec* <<https://www.quebec.ca/en/justice-and-civil-status/small-claims>>.

³⁷⁵ *CCP*, *supra* note 374 Article 536.

³⁷⁶ *Ibid* Article 542.

³⁷⁷ *Ibid* Book VI Title II Chapter III; note 374.

³⁷⁸ note 362.

³⁷⁹ Farrow, *supra* note 356 at 961.

³⁸⁰ Rhode, *supra* note 321 at 81; compare Thompson, *supra* note 75 at 10.

³⁸¹ Susskind, *supra* note 22 at 48–53.

focusing on solving them before they escalate) and dispute avoidance (being able to avoid conflicts all-together).³⁸² Further, he believes in legal health promotion, i.e. using the law to promote the well-being of the individuals involved.³⁸³

This more expansive view of access to justice often includes the promotion of access to alternative forms of dispute resolution.³⁸⁴ Alternative Dispute Resolution (ADR) refers to “mechanisms that aim to solve disputes without recurring to the traditional judicial process”.³⁸⁵ This process can be less formal and more efficient than the traditional court system, and encourage the creation of mutually agreeable solutions to disputes.³⁸⁶ It typically involves processes that are collaborative, such as negotiation, mediation and arbitration.

In negotiation, parties attempt to find common points and mutually beneficial ways to solve disputes. In mediation, a third neutral party enters the process to steer the parties towards such a conclusion. In arbitration, the third party takes a role similar to the judge and renders a binding decision based on the facts of the case.³⁸⁷ This shift towards a less adversarial process is likely to require a rethinking of many roles in the justice system.³⁸⁸ While traditional ADR requires the physical assembly of the parties and mediators, online dispute resolution takes this process online (see 3.5.2.3).³⁸⁹

3.3.2 Access to Justice as a human right?

As we have seen, the issues with access to justice have a significant impact on the ability of individuals to deal with their legal problems. Beyond the practical importance,

³⁸² *Ibid* at 66–68.

³⁸³ *Ibid* at 69–70.

³⁸⁴ note 362; Farrow, *supra* note 356; Sandefur, *supra* note 360; Garth & Cappelletti, *supra* note 362 at 223–227; Canada, *supra* note 361 at 23.

³⁸⁵ Davide Carneiro et al, “Online dispute resolution: an artificial intelligence perspective” (2014) 41:2 *Artif Intell Rev* 211–240 at 212.

³⁸⁶ *Ibid*; Canada, *supra* note 361 at 153.

³⁸⁷ Carneiro et al, “Online dispute resolution”, *supra* note 385 at 213; Karim Benyekhlef & Fabien Gélinas, “Online Dispute Resolution” (2005) 10:2 *Lex Electronica*, online: <<https://www.lex-electronica.org/en/articles/vol10/num2/online-dispute-resolution/>> at 44–52.

³⁸⁸ Benyekhlef et al, *supra* note 322 at 10–11.

³⁸⁹ Susskind, *supra* note 22 at 62.

scholars are increasingly arguing that the provision of effective access to justice is crucial for maintaining the rule of law and constitutes a human right for the individual.

3.3.2.1 Access to Justice and the Rule of Law

Access to Justice is a fundamental requirement for a democratic state to be governed by the Rule of Law. The Rule of Law, Rechtsstaat and *État de Droit* are different but similar concepts, all alluding to the supremacy of law. They are a significant influence on legal philosophy in modern democracies.³⁹⁰ The Rule of Law can be seen as the idea that all people and the government should be ruled by law and not by arbitrary rules.³⁹¹

According to Raz, the Rule of Law relies on an accessible court system, as excessive costs or delays can frustrate the populations ability to be guided by law and make any law irrelevant in practice.³⁹² Pierre Noreau also argues that the Rule of Law encompasses the effective possibility for individuals to use courts to settle their disputes.³⁹³ This view seems to be shared by the Supreme Court of Canada. It has held that “Ensuring access to justice is the greatest challenge to the Rule of Law in Canada today.”³⁹⁴

It thus seems like the state has an obligation, in reaching the goal of being a state of law, to ensure meaningful and efficient access to the court system. The easier, quicker and cheaper people can solve their legal disputes, the more efficiently the laws governing society will be implemented and the stronger the Rule of Law will grow.

3.3.2.2 Human rights legislation

Given the importance of access to justice, it is not surprising that it has been encoded in multiple human rights instruments.

³⁹⁰ Luc Heuschling, *État de droit, Rechtsstaat, Rule of law*, Nouvelle bibliothèque de thèses 16 (Paris: Dalloz, 2002) at 431.

³⁹¹ Joseph Raz, *The authority of law: essays on law and morality* (Oxford University Press on Demand, 2009) at 212.

³⁹² *Ibid* at 217.

³⁹³ Pierre Noreau, “Accès à la justice et démocratie en panne: constats, analyses et projections” (2010) *Révolutionner la justice*, Montréal, *Thémis* 13 at 15; compare Jeremy Waldron, “The Rule of Law” in Edward N Zalta, ed, *The Stanford Encyclopedia of Philosophy*, spring 2020 ed (Metaphysics Research Lab, Stanford University, 2020).

³⁹⁴ *Hryniak v Mauldin*, 2014 Supreme Court of Canada, 2014 SCC 7 (CanLII), 1 SCR 87.

For example, the European Convention for the Protection of Human Rights and Fundamental Freedoms, Article 6 contains the following right: “In the determination of his civil rights and obligations [...], everyone is entitled to a fair and public hearing within a reasonable time [...]”.³⁹⁵

Likewise, the Canadian Charter of Human Rights and freedoms contains the following article:

*Every person has a right to a full and equal, public and fair hearing by an independent and impartial tribunal, for the determination of his rights and obligations or of the merits of any charge brought against him.*³⁹⁶

These articles show that providing effective ways to resolve their disputes can be considered a human right. Looking at the statistics described above, many places could be seen to have room for improvement regarding this right.

3.3.2.3 UN 2030 Agenda for Sustainable Development

The 2030 Agenda for Sustainable Development by the United Nations also indicates the importance of access to justice. This agenda was adopted by the United States General Assembly on 25 September 2015.³⁹⁷ As part of the agenda, 17 development goals have been adopted to ensure peace and prosperity for the people and the planet. Goal 16 is to:

*Promote peaceful and inclusive societies for sustainable development, provide access to justice for all and build effective, accountable and inclusive institutions at all levels*³⁹⁸

The subgoal 16.3 specifically states:

*Promote the rule of law at the national and international levels and ensure equal access to justice for all*³⁹⁹

³⁹⁵ *European Convention for the Protection of Human Rights and Fundamental Freedoms, as amended by Protocols Nos. 11 and 14*, (Council of Europe) 1950 Article 6.

³⁹⁶ *Charter of human rights and freedoms*, C-12 article 23.

³⁹⁷ *Transforming our world: the 2030 Agenda for Sustainable Development*, A/RES/70/1 (United Nations General Assembly) at 1.

³⁹⁸ *Ibid* at 25.

In 2020, a new indicator was added to measure this goal:⁴⁰⁰

Indicator 16.3.3: Proportion of the population who have experienced a dispute in the past two years and who accessed a formal or informal dispute resolution mechanism, by type of mechanism.

As we can see, the importance of mechanisms to resolve disputes of individuals quickly and efficiently seems recognized by the United Nations as well. The statistics discussed above in 3.2.2 show that a minority of people facing legal issues use dispute resolution mechanisms to resolve their disputes, and that many are unable to adequately address their legal needs. Increasing the ability of people to resolve their disputes is thus a crucial aspect of achieving goal 16 of the agenda for sustainable development.

3.3.3 Access to Justice in this thesis

As we have seen, access to justice is a crucial area of research, and an important human rights issue. In this thesis, I will focus on building tools using artificial intelligence that can increase access to justice. In this, I will mostly focus on two of the discussed views on access to justice:

- Access to equal outcomes (3.3.1.3) – a key aspect of providing access to equal outcome is informing individuals of their rights, so that they can enforce these rights and know how to proceed in the judicial system.
- Access to dispute resolution (3.3.1.4) – beyond the court system, providing individuals with legal information plays an important role in alternative dispute resolution. Legal information can support the parties in settling their disputes through alternative dispute resolution mechanisms. Being aware of one's rights is an important cornerstone in ensuring just outcomes in a negotiation situation, since it can inform the user that they have a right, and what they can ask for in the negotiation. Further, legal information is an important element in supporting

³⁹⁹ *Ibid.*

⁴⁰⁰ Savage & McDonald, *supra* note 303 at 9.

individuals in making use of the law, which Susskind refers to as “legal health promotion”.⁴⁰¹

Increasing access to legal information is a crucial aspect of both of these views on increasing access to justice. The importance of access to legal information is also mirrored in the surveys discussed above, which showed that a lack of legal information was a key reason for individuals not attempting to resolve their legal problems. Further, a majority of the survey respondents that did take action believed that legal information would have improved their outcomes.⁴⁰² Thus, in the research presented here I focus on building a methodology that can increase access to justice through increasing access to legal information for laypeople. Let us explore the concept of access to legal information.

3.4 Access to Legal information

Providing improved access to legal information is often discussed as one of the main strategies of increasing access to justice.⁴⁰³ Individuals may not understand that their issue has legal aspects, and what their legal rights may be in such a situation. Noreau indicated that only 36% of individuals are able to understand the legal system, 48% feel like they know the law, and 21% are able to read legal texts.⁴⁰⁴ This lack of knowledge means that an individual may not know that they even have the possibility to use the legal system to resolve their issue. Using the legal system for conflict resolution thus presupposes knowledge of the law. Improved legal information is a frequently recommended goal for empowering self-represented litigants and increasing access to justice.

The following quote from the Paths to Justice study illustrates this issue:

“A clear message that emerges from the study is the profound need for knowledge and advice about obligations, rights, remedies, and procedures for resolving justiciable

⁴⁰¹ Susskind, *supra* note 22 at 69–70.

⁴⁰² See 3.2.2.

⁴⁰³ Maurits Barendrecht, “Legal aid, accessible courts or legal information? Three access to justice strategies compared” (2011) 11:1 *Global Jurist* at 4.

⁴⁰⁴ Noreau, “Accès à la justice et démocratie en panne”, *supra* note 393 at 19.

problems. [...] the pervasive lack of knowledge about legal rights and procedures for enforcing or defending rights can lead to an unnecessary level of helplessness even among the more competent and resourceful."⁴⁰⁵

Increasing access to legal information is crucial even beyond the presence of a legal problem. Most of us face legal problems or disputes very rarely. However, people are constantly expected to follow the laws, as they apply to us. If we diverge from the laws, individuals may need to pay fines, be liable to third parties or even criminally charged. Crucially, this requirement applies even if we are not aware of the law, under the rule of "ignorance of the law is no excuse", *ignorantia juris non excusat*.⁴⁰⁶ In practice, as we have seen, people are often unaware of the rules, making it impossible for them to adjust their behavior to the laws. Having access to legal information may therefore be a crucial feature of adjusting to the rules governing society, to prevent possible problems from occurring in the first place.⁴⁰⁷

Likewise, there may be situations where carrying out certain actions is restricted, requiring the following procedures or formalities. For example, building a pool in the backyard requires a permit in many locations. Not being aware of this may not be seen as an unmet need. However, a problem can arise if an individual decides to start building their pool without obtaining the permit. In this situation, legal information is also important as a means for resolving problems, by being a key aspect in avoiding these problems in the first place. Susskind refers to this aspect of access to justice as "dispute avoidance".⁴⁰⁸

Further, legal information can be important in situations where an individual has a legal right or opportunity that they are not aware of. For example, individuals may have the right to social aid. The law can also be used to create contracts, write wills or get

⁴⁰⁵ Genn, *supra* note 290 at 255.

⁴⁰⁶ Leesi Ebenezer Mitee, "The Right of Public Access to Legal Information: A Proposal for its Universal Recognition as a Human Right" (2017) 18:6 German Law Journal 1429–1496 at 1430.

⁴⁰⁷ Denise Tay Hui Yuan, "Access to Legal Information and Counsel: Prevention Is as Important as Cure" (2017) 35 Sing L Rev 47–65 at 53–54.

⁴⁰⁸ Susskind, *supra* note 22 at 68.

married.⁴⁰⁹ These can be seen as “power conferring” rules.⁴¹⁰ Clearly having access to legal information in this situation is advantageous, and can give an individual access to benefits provided by society.⁴¹¹ Susskind refers to this aspect of access to justice as “legal health promotion”, aiming to increase the welfare of citizens by “offering access to the opportunities that the law creates.”⁴¹² According to the UN Task Force for Justice, at least 4.5 billion people are excluded from the opportunities the law provides.⁴¹³

As we can see, increasing access to legal information is a crucial component in increasing access to justice in the broad sense. The next step is to explore what could be meant by access to legal information.

3.4.1 Different aspects of access to legal information

Just like with access to justice, there could be different views of what is entailed by providing legal information.

3.4.1.1 Access to Legal Source Information

First, access to legal information may refer to access to legal source information, such as laws, regulations and case law. Open access to this material may seem like a given, but is not present in all jurisdictions.⁴¹⁴ Information may not be accessible online at all, hidden behind paywalls, or even be protected by copyright by the government.⁴¹⁵

The Free Access to Law movement aims to address this issue. It began in 1992 at Cornell Law School, when two researchers provided public access to some US legal materials and called the project the Legal Information Institute (LII).⁴¹⁶ Other LIIs were established all over the world in the coming years, including CanLII, which was built by LexUM at

⁴⁰⁹ *Ibid* at 69.

⁴¹⁰ *Ibid*.

⁴¹¹ Tay Hui Yuan, “Access to Legal Information and Counsel”, *supra* note 407 at 52; Susskind, *supra* note 22 at 69.

⁴¹² Susskind, *supra* note 22 at 69–70.

⁴¹³ note 8 at 38.

⁴¹⁴ Mitee, “The Right of Public Access to Legal Information”, *supra* note 406 at 1431.

⁴¹⁵ *Ibid* at 1432–1436.

⁴¹⁶ Graham Greenleaf, Andrew Mowbray & Philip Chung, “The Meaning of Free Access to Legal Information: A Twenty Year Evolution” (2013) 1:1 J Open Access L 1–68 at 4.

the Université de Montréal.⁴¹⁷ In 2002, the declaration on Free Access to Law (Montreal Declaration) was agreed upon by a number of LIIs.⁴¹⁸ Since then, free access to some legal information has become commonplace all over the world, frequently provided by courts, legislatures, government departments and law schools.⁴¹⁹ An example of such a platform is EUR-Lex, which publishes the official cases and legislation of the European Union. Since July 2013, the electronic version of the Official Journal on EUR-Lex is seen as the version of record, i.e. it is authentic and produces legal effects, supplanting the paper version in most cases.⁴²⁰

3.4.1.2 Understanding of Legal Information

However, having access to the law may not help individuals that are not able to understand it. Laws and court cases are difficult to understand for individuals with no legal education. Even if the laws and cases are public, people may not obtain the benefits discussed above. Ensuring that people can understand the law is therefore an important part of the access to legal information.⁴²¹

Tay Hui Yuan argues for the inclusion of a Legal Education and Awareness Programme (LEAP) in the school curriculum. This could teach individuals how to find and understand laws, and relevant legislation regarding theft, domestic violence and cyber bullying.⁴²² MacDonald likewise argues that courses in school, but also information sessions at community centers or senior groups could enhance a culture of legal

⁴¹⁷ *Ibid.*

⁴¹⁸ *Ibid* at 4–5; “Declaration on Free Access to Law”, online: <<http://www.falm.info/declaration/>>.

⁴¹⁹ Greenleaf, Mowbray & Chung, “The Meaning of Free Access to Legal Information”, *supra* note 416 at 5.

⁴²⁰ “Access to the Official Journal - EUR-Lex”, online: <<https://eur-lex.europa.eu/oj/direct-access.html>> Usr_lan: en; “Verify the authenticity of an electronic edition of the Official Journal - EUR-Lex”, online: <<https://eur-lex.europa.eu/content/help/oj/authenticity-eOJ.html>> Usr_lan: en; *Council Regulation (EU) No 216/2013 of 7 March 2013 on the electronic publication of the Official Journal of the European Union*, OJ L 2013Legislative Body: CONSIL.

⁴²¹ Tay Hui Yuan, “Access to Legal Information and Counsel”, *supra* note 407 at 57–58.

⁴²² *Ibid* at 58.

awareness.⁴²³ Bilson *et al* advocate using librarians as trusted sources for legal information.⁴²⁴

Providing legal information is not limited to providing information about the law itself. As we have seen, there has also been a movement toward materials supporting individuals in practical legal matters, such as simplified forms and self-help books.⁴²⁵

3.4.2 Access to legal information as a human right

Just like access to justice, access to legal information may be seen as a human right.

An important tenet of the rule of law is stability and predictability. These are crucial to individuals' ability to manage their affairs effectively.⁴²⁶ Stability requires that laws are promulgated well in advance of it being applied, so that individuals have a chance to adjust to them.⁴²⁷

Likewise, the law being clear, accessible and intelligible is crucial so that citizens can adjust to them.⁴²⁸ However, as we have seen, people do not feel like they understand the laws. How can people be expected to adjust their behavior to the law if they do not understand how the law applies? Inaccessible or incomprehensible legal information is thus a significant barrier to achieving the rule of law.⁴²⁹

Further, individuals being able to understand the law can be seen as an important pre-requisite for a democratic society. To be able to decide how to vote and participate in the democratic process, individuals need to understand how the rules enacted by an elected

⁴²³ Canada, *supra* note 361 at 96.

⁴²⁴ Beth Bilson, Brea Lowenberger & Graham Sharp, "Reducing the 'Justice Gap' Through Access to Legal Information: Establishing Access to Justice Entry Points at Public Libraries" (2017) 34:2 Windsor Yearbook of Access to Justice/Recueil annuel de Windsor d'accès à la justice 99–128.

⁴²⁵ Branting, *supra* note 78 at 1.

⁴²⁶ Stefanie A Lindquist, "7. Stare Decisis as Reciprocity Norm" in Charles Gardner Geyh, ed, *What's Law Got to Do With It?* (Stanford University Press, 2020) 173.

⁴²⁷ Waldron, *supra* note 393 section 5.1.

⁴²⁸ *Ibid* section 5.1.

⁴²⁹ Mitee, "The Right of Public Access to Legal Information", *supra* note 406 at 1470–1471.

parliament are applied in practice, and their effects. Otherwise, it becomes difficult to evaluate the work of the legislator and decide how to respond to this work in democratic elections.

Leesi Ebenezer Mitee argues that the right of public access to legal information should be recognized as a human right,⁴³⁰ and proposes a UN convention to codify this right.⁴³¹

3.4.3 Access to legal information in this thesis

In this thesis, I aim to develop a methodology to increase access to justice by providing legal information to laypeople. I intend to develop a system that can ask questions to understand the situation of a user, and then provide them with information with regards to their rights. This approach goes beyond merely providing legal source materials, by helping the users understand the laws and how they might apply to them.

The methodology could be used to build systems to provide information about legal problem situations, but also to inform users of how they may act to comply with the law, or to benefit from the legal opportunities available to them in specific situations.

Providing this kind of legal information could be an important way to increase the access to justice of laypeople.

As technology is permeating all aspects of our society, it is also an important tool in increasing access to justice and access to information. Let us explore some aspects of using technology to increase access to justice.

3.5 Using technology to increase access to justice

3.5.1 Introduction

In my research, I will use technology and artificial intelligence to increase access to justice through providing access to legal information. In this section, I will explore some previous research in using technology in this area, and the consequences thereof. Here, I

⁴³⁰ *Ibid* at 1454–1471.

⁴³¹ Tay Hui Yuan, “Access to Legal Information and Counsel”, *supra* note 407 at 1471–1489.

am more interested in the conceptual underpinnings of the projects, while a more in-depth technical review of prior work will follow in Chapter 4.

3.5.2 Ways of using technology to increase access to justice

There are many different important projects aiming to provide access to justice through technology. Let us take a look at using technology to provide legal information and dispute resolution to individuals using the internet.

3.5.2.1 Providing legal information online

As we have seen, providing legal information is a key step in increasing access to justice. With the growth of technology, the internet has become a very natural platform for providing this information.

At the most basic level, websites are an accessible and cheap medium to provide access to laws and explanatory texts to citizens.⁴³² In the United States, every state implements an online legal aid platform, providing links to self-help resources and intakes to assistance.⁴³³

Further, courts are starting to implement forms online, allowing individuals to use online tools rather than filling out paper forms. However, MacFarlane found that many such services often had significant deficiencies, including being complicated to use, not providing practical information, and requiring some level of knowledge and understanding. She recommends the establishment of best practices.⁴³⁴ Systems may also offer interactive components, such as instant messaging support,⁴³⁵ or e-filing, that allow individuals to file cases electronically.⁴³⁶

⁴³² See 3.4.1.1.

⁴³³ James E Cabral et al, “Using Technology to Enhance Access to Justice” (2012) 26:1 Harv JL & Tech 241 at 246.

⁴³⁴ Macfarlane, *supra* note 323 at 113–115.

⁴³⁵ Cabral et al, *supra* note 433 at 249–251.

⁴³⁶ *Ibid* at 252–253.

3.5.2.2 Providing legal information with artificial intelligence

The aforementioned ways of providing legal information are essentially taking the current documents and paradigms and providing access via the internet. Beyond this, there are also more interactive ways of using technology to provide legal information.

Artificial intelligence systems, including expert systems, can provide more efficient and targeted ways of delivering information to individuals and self-represented litigants, and act as decision support systems, that support the user in deciding how to proceed with their issue. They can be used as a triage system to guide individuals toward the current procedures and forms, as recommended by MacFarlane.⁴³⁷ The US Legal Services Corporation published a report titled “The Summit on the Use of Technology to Expand Access to Justice” in 2013. It suggests creating unified legal portals that can triage and direct individuals towards the most appropriate forms of assistance and guides, document assembly platforms that can generate documents for self-represented litigants, and developing expert systems to assist lawyers.⁴³⁸

The implementation of such systems to increase access to justice has been discussed at-length in academic publications. Branting presents a concrete concept for advisory systems for pro se litigants.⁴³⁹ These systems would:

1. Inform the user of the available forms of relief
2. Determine whether the user could fulfill the substantive requirements of this relief
3. Inform the user of the procedural requirements for the relief
4. Assist the user in developing documents necessary to initiate the action⁴⁴⁰

Branting implemented such a system, called the Protection Order Advisory, which could support individuals in obtaining protection orders.⁴⁴¹

⁴³⁷ Macfarlane, *supra* note 323 at 116.

⁴³⁸ *Report of The Summit on the Use of Technology to Expand Access to Justice* (Legal Services Corporation, 2013); Zeleznikow, *supra* note 330 at 34.

⁴³⁹ Branting, *supra* note 78.

⁴⁴⁰ *Ibid* at 3.

Similarly, Thompson describes a concept he calls JPES, Justice Pathway Expert System.⁴⁴² These systems would consist of Expert Systems⁴⁴³ that ask users for their input using an intelligent questionnaire, and then provide relevant information to the user based on their answers.⁴⁴⁴ After this, the user would be guided towards possible avenues for solutions, such as courts, online dispute resolution systems, or specialized support forums.⁴⁴⁵ Thompson believes that this system could support users directly, and also broaden the approach to justice, by pulling inertia from litigation towards alternative forms of dispute resolution.⁴⁴⁶

Zelevnikow presents the GetAid system, which can determine whether an individual is eligible for legal aid, using a web-based decision tree.⁴⁴⁷ He believes that lawyers will be the main users of the system.⁴⁴⁸

A notable implemented tool is the Rechtwijzer 1.0 platform, which was developed in the Netherlands and launched in 2012. It aims to improve access to justice and legal information by providing individuals the tools to understand a conflict and solve it on their own.⁴⁴⁹ The system targeted consumer disputes and divorce.⁴⁵⁰ After choosing a conflict type, the user is asked a number of questions, including factual questions, but also questions regarding the relationship between the parties. Some of these questions are aimed at inciting reflection by the parties, such as whether they feel capable of resolving the problems themselves or would prefer outside help, and whether the parties are open to cooperate to solve the issue.⁴⁵¹ Based on the answers, the system would guide the user

⁴⁴¹ *Ibid* at 3–7.

⁴⁴² Thompson, *supra* note 75.

⁴⁴³ *Ibid* at 13–14.

⁴⁴⁴ *Ibid* at 24–32.

⁴⁴⁵ *Ibid* at 33–36.

⁴⁴⁶ *Ibid* at 10.

⁴⁴⁷ Zelevnikow, *supra* note 67 at 20.

⁴⁴⁸ *Ibid* at 25.

⁴⁴⁹ Esmée A Bickel, M Van Dijk & Ellen Giebels, “Online legal advice and conflict support: a Dutch experience” (2015) Report, University of Twente at 4.

⁴⁵⁰ *Ibid*.

⁴⁵¹ *Ibid* at 5.

towards relevant information and support services.⁴⁵² Overall, users seemed happy with the system, with survey respondents rating their overall experience 7.51 out of 10 for divorce information,⁴⁵³ and 7.29 out of 10 for consumer conflicts.⁴⁵⁴

There are now publicly available platforms that allow the creation of systems that can ask questions to the user and output filled forms. Examples of such platforms include A2J Author⁴⁵⁵ and DocAssemble⁴⁵⁶, which have been used to fill out hundreds of thousands of documents.⁴⁵⁷

Decision support systems using artificial intelligence may go further than informing the user of their rights. By analyzing previous cases, the systems may be able to predict the outcome of the case of a user (such as damages awarded), providing them with relevant context as to whether they may want to continue with their case, and what they can expect should they do so. This estimation of an outcome can be used as a “best alternative to a negotiated agreement” (BATNA). Introduced by Fisher and Ury in 1981,⁴⁵⁸ this concept refers to the best possible alternative when negotiating. Being aware of the BATNA in a negotiation allows the parties to have a better idea of whether they should accept an agreement (if it is better than the BATNA) or reject it (if it is worse than the BATNA).⁴⁵⁹

Legal decision support tools could be used to provide a BATNA to parties to a legal conflict, supporting them in settling their case rather than going to court.⁴⁶⁰

⁴⁵² Laura Kistemaker, “Rechtwijzer and Uitelkaar.nl. Dutch Experiences with ODR for Divorce” (2021) 59:2 Family Court Review 232–243 at 233.

⁴⁵³ Bickel, Van Dijk & Giebels, “Online legal advice and conflict support”, *supra* note 449 at 22.

⁴⁵⁴ *Ibid* at 38.

⁴⁵⁵ “A2J Author”, online: <<https://www.a2jauthor.org/>>.

⁴⁵⁶ “Docassemble”, online: *Docassemble* <<http://docassemble.org/>>.

⁴⁵⁷ Cabral et al, *supra* note 433 at 251.

⁴⁵⁸ Roger Fisher, William Ury & Bruce Patton, *Getting to Yes: Negotiating Agreement Without Giving In*, revised ed (New York: Penguin Books, 2011).

⁴⁵⁹ Zeleznikow, *supra* note 330 at 35.

⁴⁶⁰ Zeleznikow, *supra* note 67 at 28; Samuel Dahan et al, “Predicting Employment Notice Period with Machine Learning: Promises and Limitations” (2020) 65:4 McGill Law Journal/Revue de droit de McGill 711–753 at 1; Thompson, *supra* note 75 at 26.

As we have seen, providing information online is a very important step towards increasing access to justice. My research aims to build a methodology that can create online tools providing information to individuals, using artificial intelligence. This information could support the user in understanding their rights, knowing how to enforce their rights, and settling their dispute.

However, if settlement does not succeed, the user may still have to go through the court system to enforce their right, with the costs that this entails. Next, let us explore online dispute resolution, where technology can be used to resolve disputes outside of the traditional court system.

3.5.2.3 Online Dispute Resolution

Online Dispute Resolution (ODR) can be seen as the implementation of Alternative Dispute Resolution systems in a dematerialized context.⁴⁶¹ However, there is no generally accepted definition of what is encompassed by ODR, and it has been described as “as technology-assisted dispute resolution, by others as technology-facilitated dispute resolution, and by still others as technology based-dispute resolution schemes.”⁴⁶² Further, the term ODR can be seen to include private parties deploying such systems, or also state courts that run online.⁴⁶³ Benyekhlef discusses different constellations that could be present in ODR platforms, such as platforms targeting Business-to-business disputes, business-to-consumer disputes, or systems for disputes arising in the general public.⁴⁶⁴

3.5.2.3.1 Elements of ODR platforms

ODR systems can combine the different forms of alternative dispute resolution into a single platform.⁴⁶⁵ These include:

⁴⁶¹ Benyekhlef et al, *supra* note 322 at 11.

⁴⁶² Daniel Rainey, Ethan Katsh & Mohamed S Abdel Wahab, eds, *Online Dispute Resolution - Theory and Practice*, 2d ed (Eleven International, 2021) at 15.

⁴⁶³ Susskind, *supra* note 22 at 61–63.

⁴⁶⁴ Benyekhlef, *supra* note 321 at 82–83.

⁴⁶⁵ Benyekhlef & Gélinas, *supra* note 387 at 52–54.

Negotiation – the parties are voluntarily able to exchange views and try to come to an agreement. In ODR platforms, this can be done easily and asynchronously from one’s home.⁴⁶⁶ If the settlement fails, the system can move on to mediation.

Mediation – A third party mediator is introduced into the process, who tries to guide the parties toward a settlement. The mediator is free to explore solutions that are not constrained by the formal legal systems, by taking into account the interests at stake and the impact on the future.⁴⁶⁷ Susskind refers to mediation as the “dispute containment” stage of an ODR platform.⁴⁶⁸

Arbitration – if mediation fails, ODR platforms can provide the facility to implement arbitration. Here, the dispute is submitted to an independent, private tribunal which renders a binding decision in the matter. It requires the parties to consent beforehand.⁴⁶⁹

3.5.2.3.2 Examples of ODR platforms

Online dispute resolution systems have been deployed in many jurisdictions, including Canada, China, the United States, the European Union and the United Kingdom.⁴⁷⁰

Below are a few examples of such platforms.

The Cyberjustice Laboratory has developed a modular ODR platform known as « Plateforme d’Aide au Règlement de Litiges en ligne » (PARLe).⁴⁷¹ It has been successfully deployed at the “Office de la protection du consommateur du Québec”. The results are very positive, with 70% of cases that enter the platform being solved, and 90% of users indicating that they are satisfied.⁴⁷² On average in 2020-21, the platform was able

⁴⁶⁶ *Ibid* at 44–45.

⁴⁶⁷ *Ibid* at 45–49.

⁴⁶⁸ Susskind, *supra* note 22 at 135–141.

⁴⁶⁹ Benyekhlef & Gélinas, *supra* note 387 at 49–52; Susskind, *supra* note 22 at 143–152.

⁴⁷⁰ Susskind, *supra* note 22 at 165–176.

⁴⁷¹ “PARLe”, (6 December 2018), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/parole-3/nos-solutions-logicielles/parole-2/>>.

⁴⁷² “Qu’est-ce que PARLe?”, (1 June 2022), online: *Office de la protection du consommateur* <<https://www.opc.gouv.qc.ca/a-propos/parole/description/>>.

to settle cases in 22 days.⁴⁷³ The PARLe platform has further been implemented with the Ontario Condominium Authority Tribunal⁴⁷⁴ and the Commission des normes, de l'équité, de la santé et de la sécurité du travail.⁴⁷⁵ Currently, it is being implemented with the Pay Equity Division of the Canadian Human Rights Commission.⁴⁷⁶

Online Dispute Resolution has even been made mandatory in certain jurisdictions. In 2013, the European Union implemented an Online Dispute Resolution system for dealing with issues arising for consumers in online purchases.⁴⁷⁷ The regulation requires all online marketplaces to provide a link to this platform.⁴⁷⁸ Between 2016 and 2018, the platform dealt with an average of 2,000 claims each month, to a 71% satisfaction rate of the users.⁴⁷⁹ Benyekhlef and Vermeys discuss ways to improve the efficiency and reach of this platform, including by advertising and harmonizing the platform to a greater extent.⁴⁸⁰

ODR systems can also be connected directly to an online e-commerce platform. Perhaps the most used ODR system was a commercial system used by eBay to resolve disputes

⁴⁷³ "Réalizations de l'Office en 2020-2021", (1 March 2022), online: *Office de la protection du consommateur* <<https://web.archive.org/web/20220302062245/https://www.opc.gouv.qc.ca/a-propos/office/realisations-20-21/>>.

⁴⁷⁴ "PARLe – Ontario Condominium Authority Tribunal", (7 April 2020), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/2020/04/07/the-parle-project-the-condominium-authority-tribunal/>>.

⁴⁷⁵ "PARLe – Commission des normes, de l'équité, de la santé et de la sécurité du travail", (7 April 2020), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/2020/04/07/the-parle-project-social-mediation-with-the-cnesst/>>.

⁴⁷⁶ Valentin Callipel, "New platform for online resolution of pay equity disputes in workplaces ruled by federal relementation!", (13 January 2021), online: *Laboratoire de cyberjustice* <<https://www.cyberjustice.ca/en/2021/01/13/nouvelle-plateforme-pour-regler-en-ligne-les-differends-relatifs-a-lequite-salariale-dans-les-milieus-de-travail-sous-reglementation-federale/>>.

⁴⁷⁷ *Regulation (EU) No 524/2013 of the European Parliament and of the Council of 21 May 2013 on online dispute resolution for consumer disputes and amending Regulation (EC) No 2006/2004 and Directive 2009/22/EC (Regulation on consumer ODR)*, OJ L165/1 2013.

⁴⁷⁸ *Ibid* Article 14.

⁴⁷⁹ Etienne Wery, "Le règlement des litiges liés au e-commerce se fait de plus en plus souvent en ligne", (16 May 2018), online: *Droit & Technologies* <<https://www.droit-technologie.org/actualites/reglement-litiges-lies-e-commerce-se-de-plus-plus-souvent-ligne/>>.

⁴⁸⁰ Karim Benyekhlef & Nicolas Vermeys, "The 'Success' of Online Dispute Resolution in Europe", (18 June 2018), online: *Slaw* <<http://www.slaw.ca/2018/06/18/the-success-of-online-dispute-resolution-in-europe/>>.

between buyers and sellers. Supposedly, it handled more than 60 million disputes each year.⁴⁸¹

3.5.2.3.3 Advantages of ODR platforms

ODR platforms have a number of important advantages over the traditional court system. For example, they are both quicker and cheaper, since the parties are able to adopt the procedure to their specific needs, and benefit from the quick and efficient communication methods provided by computer systems.⁴⁸² The parties not being required to appear in person can further reduce tensions and increase equality.⁴⁸³

One of the most promising aspects of ODR is that parties often do not need the intervention of a third party to resolve their disputes. This is, of course, the most scalable way to resolve a dispute, and further leaves the parties fully in control. Further increasing the percentage of people that are able to solve their issue at the stage of negotiation is thus a desirable goal. Researchers have started to implement artificial intelligence methods to increase this rate, among other things.⁴⁸⁴ Let us explore this research.

3.5.2.4 Integrating AI into ODR platforms

Artificial intelligence could be an important step in further enhancing the effectiveness of ODR platforms. This integration has been referred to as ODRAI.⁴⁸⁵ It can be integrated into many of the different stages of the ODR platform. Let us explore some research in this area.

⁴⁸¹ John Morison & Adam Harkens, “Re-engineering justice? Robot judges, computerised courts and (semi) automated legal decision-making” (2019) 39:4 *Legal Studies* 618–635 at 622; Louis F Del Duca, Colin Rule & Kathryn Rimpfel, “eBay’s De Facto Low Value High Volume Resolution Process: Lessons and Best Practices for ODR Systems Designers” (2014) 6:1 *Arbitration Law Review* 204–219.

⁴⁸² Benyekhlef & Gélinas, *supra* note 387 at 85–86.

⁴⁸³ *Ibid* at 87.

⁴⁸⁴ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797; Carneiro et al, “Online dispute resolution”, *supra* note 385; Zeleznikow, *supra* note 330.

⁴⁸⁵ *A Tale of Cyberjustice: A Modern Approach to Technology in the Canadian Justice System* (Montréal, Québec, Canada: Cyberjustice Laboratory, 2019) at 134.

3.5.2.4.1 Stage 0 – Entry way

Benyekhlef and Zhu suggest the use of AI as an entry way to the traditional ODR experience of negotiation, mediation and arbitration. An AI system based on guided pathways could ask questions to understand the situation of the user, and then provide them with information on the validity of their cause, possible outcomes should they succeed, and the competency of the ODR platform.⁴⁸⁶

Susskind refers to such a pre-stage as the *dispute avoidance* system. It can collect and scope the thoughts of the user, and indicate the possible solutions to them, and whether it makes sense to proceed.⁴⁸⁷

In practice, such a system has been implemented by the Civil Resolution Tribunal in British Columbia. The so-called solution explorer allows individuals to diagnose their problem and print letters to resolve their issue on their own.⁴⁸⁸ By 2019, the solution explorer had been used over 60k times.⁴⁸⁹

3.5.2.4.2 Stage 1 – negotiation support

AI may also be used to make the negotiation stage more efficient, allowing individuals to settle more frequently without relying on a third party.

AI systems could be used to provide the user with a prediction of how their case may be decided in court, should their negotiations fail. Such an estimation could be used as a BATNA (Best Alternative to the Negotiated Agreement), and provide the users with important context that may help them settle their case, by aligning their expectations.⁴⁹⁰

⁴⁸⁶ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797.

⁴⁸⁷ Susskind, *supra* note 22 at 121–122.

⁴⁸⁸ Shannon Salter, “Online dispute resolution and justice system integration: British Columbia’s Civil Resolution Tribunal” (2017) 34:1 Windsor Yearbook of Access to Justice 112–129 at 120.

⁴⁸⁹ Susskind, *supra* note 22 at 169.

⁴⁹⁰ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797; Carneiro et al, “Online dispute resolution”, *supra* note 385 at 21; Susskind, *supra* note 22 at 274–275; Zeleznikow, *supra* note 330 at 39.

AI may also be used in other ways at this stage. Branting *et al* introduced a system that monitors the messages exchanged during negotiation to detect situations that require immediate attention by a human mediator. It also selects standard text messages appropriate for the current stage of negotiations. These facilities could compensate for the shortage of facilitators.⁴⁹¹

Likewise, AI could be used to introduce structure into the communications between the parties. The 3.5.2.4.2 2.0 system extended the Rechtwijzer 1.0 system described above,⁴⁹² by supporting the parties in resolving their disputes rather than just providing information. This is possible by structuring the dialogue between the two parties in divorce proceedings, that could finally be formalized in court.⁴⁹³ The platform was online for two years, during which time it received a lot of attention, and was later spun out to Uitelkaar.nl.⁴⁹⁴

The researchers found that building ODR platforms this way held a number of advantages. First, the approach of being able to structure communication around the individual issues was seen as an important step in splitting the complexity of the proceedings into smaller, manageable steps, contributing to the peace of mind of the parties. The platform was further hailed for de-escalating, by providing neutral information, and empowering individuals to take control of their divorce without requiring legal knowledge, in a cost-efficient manner.⁴⁹⁵

Overall, users seemed happy with this approach, with survey respondents rating the platform 7.9 out of 10.⁴⁹⁶ At the same time, the author highlight the challenges in running

⁴⁹¹ Karl Branting et al, “A computational model of facilitation in online dispute resolution” (2022) *Artif Intell Law*, online: <<https://doi.org/10.1007/s10506-022-09318-7>>.

⁴⁹² See 3.5.2.2.

⁴⁹³ Kistemaker, *supra* note 452 at 233.

⁴⁹⁴ *Ibid* at 233–234.

⁴⁹⁵ *Ibid* at 234–235.

⁴⁹⁶ *Ibid* at 237.

such a platform, such as the resources and customer support required, and the variability in the needs of different couples.⁴⁹⁷

3.5.2.4.3 Stage 2 – Mediation

At the stage of mediation, a mediator actively participates in the process of finding solutions that are acceptable to both parties. Mediators may here benefit from the AI entry way, by receiving a structured summary of the case.⁴⁹⁸

AI may also be used to act in the role of the mediator, by using game theory and AI to nudge the parties towards an acceptable settlement.⁴⁹⁹ There have been a number of research projects investigating the use of artificial intelligence and game theory to encourage people to settle their disputes.⁵⁰⁰ For example, SmartSettle allows an algorithm to encode the preferences of the users, and then offers an algorithmic procedure to exchange proposals and find common solutions.⁵⁰¹ The Family_Winner system likewise allows parties to assign importance to certain issues, which are then used to calculate possible compromises between the parties.⁵⁰² CyberSettle, a platform launched in 1998 which claims to have handled 200k cases, used a process of blind bidding, where the parties each introduce an acceptable range of settlement. The system picks an in-between value if the ranges overlap.⁵⁰³

⁴⁹⁷ *Ibid* at 235–236.

⁴⁹⁸ Susskind, *supra* note 22 at 135.

⁴⁹⁹ Zeleznikow, *supra* note 330 at 41–42.

⁵⁰⁰ Benyekhlef & Gélinas, *supra* note 387 at 95–96.

⁵⁰¹ Arno R Lodder & Ernest M Thiessen, “The role of Artificial Intelligence in Online Dispute Resolution” (2004) Proceedings from UN forum on ODR, online: <<https://research.vu.nl/en/publications/the-role-of-artificial-intelligence-in-online-dispute-resolution>> at 5–6.

⁵⁰² Emilia Bellucci & John Zeleznikow, “Developing Negotiation Decision Support Systems that support mediators: a case study of the Family_Winner system” (2005) 13:2 Artificial Intelligence and Law 233–271.

⁵⁰³ Susskind, *supra* note 22 at 138–139.

3.5.2.4.4 Stage 3 - Arbitration

Using AI in the arbitration stage is perhaps the most controversial use of AI in an ODR platform. Zeleznikow cautions against fully automating ODR systems, and instead recommends that they should aid the decision of the parties, triage disputes and collect information for statistics.⁵⁰⁴

AI may allow the parties to formulate their arguments more efficiently for the purpose of judging. Susskind argued for systems that allow individuals to create an argument by answering a number of questions, that elicit the relevant facts, and output a document outlining the reasoning of a party.⁵⁰⁵

Satoh suggests a way that intelligent systems can be used to decide which issues the judge needs to rule on. The system asks the user whether certain facts that are required to achieve legal outcomes are present. Then, the other party can challenge the presence of these facts or introduce counterarguments. Finally, the disputed facts are provided to the judge for assessment.⁵⁰⁶ This system will be further explored below in 4.3.3.2.

Susskind goes even further than these approaches, evaluating whether AI can be used to authoritatively decide cases. While he states that AI cannot currently reason like humans, he believes this may not be as relevant as their capability to generate decisions with reasons that correspond to the decisions judges would render with some degree of confidence. According to Susskind, in some situations such as high-volume, low-intensity disputes, it could be acceptable to use these predictions as binding decisions.⁵⁰⁷ Morison and Harkens, on the other hand, argue that AI approaches may never be able to replicate the social activities that are necessary to deliver a judgment.⁵⁰⁸

⁵⁰⁴ Zeleznikow, *supra* note 330 at 42–43.

⁵⁰⁵ Susskind, *supra* note 22 at 157.

⁵⁰⁶ Ken Satoh, Kazuko Takahashi & Tatsuki Kawasaki, “Interactive system for arranging issues based on PROLEG in civil litigation” in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law* (New York, NY, USA: Association for Computing Machinery, 2021) 273 at 273.

⁵⁰⁷ Susskind, *supra* note 22 at 277–290.

⁵⁰⁸ Morison & Harkens, “Re-engineering justice?”, *supra* note 481 at 619.

3.5.2.5 Conclusion

In this section, I have described some ways technology and artificial intelligence can be used to increase access to justice. We have seen that technology and AI can be used to provide access to relevant information and provide online platforms to resolve disputes outside of the court system. The technical aspects of some of these systems will be explored below in Chapter 4. Of course, there are many more ways to increase access to justice using technology and AI that are not covered here.

In my research, I aim to build a system that is able to understand the case of a user, and then provide them with relevant information and outcomes from previous case law. It could be used by itself to provide legal information or be integrated in an ODR platform.

Many of the systems described above go further than merely providing legal information, and also help the user fill out forms or create arguments. To some extent, they perform actions that are usually performed by lawyers. Some jurisdictions have rules that prevent the unauthorized practice of law. Let us explore the legality of the aforementioned tools in the next section.

3.5.3 The legality of legal decision support tools

Even though it will not be a focus of my research, it is important to keep in mind the legality of systems using AI to support litigants. Many jurisdictions limit the practice of law to members of the bar, in order to protect the public against legal services given by unqualified individuals.⁵⁰⁹ In Quebec, for example, Article 128 of the “Act respecting the Barreau du Québec” stipulates:

The following acts, performed for others, shall be the exclusive prerogative of the practising advocate or solicitor:

(a) to give legal advice and consultations on legal matters;

⁵⁰⁹ Thomas E Spahn, “Is Your Artificial Intelligence Guilty of the Unauthorized Practice of Law” (2017) 4 Rich JL & Tech 1–47 at 7; Taiwo A Oriola, “The use of legal software by non-lawyers and the perils of unauthorised practice of law charges in the United States: a review of Jayson Reynoso decision” (2010) 18:3 Artif Intell Law 285–309 at 286.

*(b) to prepare and draw up a notice, motion, proceeding or other similar document intended for use in a case before the courts,*⁵¹⁰

The Barreau believes that the giving of any form of opinion on a legal matter, on a subject where one can be of different opinions, constitutes the giving of legal advice.⁵¹¹ How this applies in the context of legal aid software has been explored by academics in the United States.

Cabral *et al* find that each state generally defines the distinction between giving legal advice and legal information. The distinction usually relates to the specificity of the information. If information is targeted towards a general situation, it is not seen as the practice of law. If information is aimed at the particular facts of the situation of the recipient, this is considered giving legal advice, and thus regulated by the bar. Books explaining how to deal with legal issue are generally considered legal information.⁵¹²

Commercial providers of software that is able to complete legal forms, such as LegalZoom, have been found guilty of unauthorized practice of law. This software allows the user to step through a number of questions and use the answers to generate the appropriate forms. Courts seem to have argued that this goes beyond legal information, since the software selects the correct form and identifies the appropriate location to put the information provided by the user.⁵¹³ Similar arguments were raised with regard to an online bankruptcy case preparation, that was also found to constitute the practice of law.⁵¹⁴ While certain jurisdictions have excluded non-profit self-help tools from being considered as practice of law, it seems like uncertainty still exists in this regard.⁵¹⁵

⁵¹⁰ *Act respecting the Barreau du Québec*, CQLR, c B-1 section 128.

⁵¹¹ *Distinction entre donner un avis ou une opinion juridique et donner une information juridique* (Barreau du Québec, 2013).

⁵¹² Cabral *et al*, *supra* note 433 at 318.

⁵¹³ *Ibid* at 318–321.

⁵¹⁴ Oriola, “The use of legal software by non-lawyers and the perils of unauthorised practice of law charges in the United States”, *supra* note 509 at 289–294.

⁵¹⁵ Cabral *et al*, *supra* note 433 at 321–322.

Some scholars seem to believe that the rules regarding legal decision support tools will be softened in the future. Spahn predicts that the lawyers will ultimately lose their fight of branding artificial intelligence tools as unauthorized practice of law.⁵¹⁶ Oriola also argues that legal software should fall outside of the scope of legal practice, claiming that the prohibition of the use of legal software could raise issues with antitrust law and the right to self-representation. He suggests that states should exempt the production of legal software from the definition of the practice of law, provided there is an adequate disclaimer in the software.⁵¹⁷

This softening already seems to be underway in some jurisdictions. In 1999, the legal manual publisher Nolo Press was challenged by the Texas bar for the unauthorized practice of law. The company overcame these issues, and the case even caused the Texas Legislature to enact a law explicitly authorizing providing information to individuals via books and websites, as long as the information clearly states that it does not replace the advice of an attorney.⁵¹⁸

To sum up, while there does not seem to be a consensus on the use of software to support laypeople, there seems to be a movement to recognize the importance of using software to support pro se litigants. However, it is important to keep in mind the restrictions imposed on such software today by the legal bars in many jurisdictions. In my thesis, I intend to develop a system that is specifically aimed at providing legal information, in order to ensure the risk-free deployment of such tools. Of course, providing legal information only does not mean that such tools cannot potentially have negative effects that need to be assessed, in addition to the many positive effects such tools could bring. Next, let us consider how legal decision support tools may affect society.

⁵¹⁶ Spahn, *supra* note 509 at 47.

⁵¹⁷ Oriola, “The use of legal software by non-lawyers and the perils of unauthorised practice of law charges in the United States”, *supra* note 509 at 308.

⁵¹⁸ Landsman, *supra* note 338 at 445–446.

3.6 The societal impacts of legal decision support tools

In the beginning of this chapter, I identified a number of issues with the current legal system that affect an enormous number of individuals (3.2). There seem to be issues with access to justice (3.3) and access to legal information (3.4). In 3.5, I explored using technology to overcome these issues. Developing a system using AI to provide access to legal information was the goal of my research and is the subject of my thesis.

Now, let us briefly explore some of the impacts such a system could have on society. It is important to be aware of potential positive and negative impacts of my research. I look at potential impacts based on a legal decision support tool that is able to

- a) provide legal information, by assessing the situation of the user and providing them with information about their rights and procedural steps to enforce their rights. Such a system could be used to provide the user with legal information,⁵¹⁹ or as an entry stage to an ODR platform.⁵²⁰
- b) provide the user with outcome information, either through providing references to similar cases or by using AI to predict a new case. Such a system could provide a user with insight into odds of success and outcomes,⁵²¹ and also act as a negotiation support in an ODR platform.⁵²²

3.6.1 Opportunities

The creation of legal decision support tools entails a number of important opportunities in society.

3.6.1.1 Increasing understanding of rights

A system that provides the user with legal information based on the particularities of their case could be very powerful in allowing the user to understand the rights that apply to their situation. As we have seen above in 3.2.2 and 3.4, many individuals do not

⁵¹⁹ See 3.5.2.2.

⁵²⁰ See 3.5.2.4.1.

⁵²¹ See 3.5.2.2.

⁵²² See 3.5.2.4.2.

understand their rights, which is a significant impediment to knowing that there is even a legal issue, and what to do about it. Providing legal information is thus an important precondition for individuals to resolve their issue. This could help alleviate the significant negative effects that individuals face when dealing with unresolved legal problems, by making the user aware that their case does have a possible legal solution and guiding the user to the relevant forum or procedure,⁵²³ such as hiring a lawyer, filling out a form or filing a claim at a certain court.

3.6.1.2 Preventing unviable conflicts

In some cases, the opposite may be true – a user thinks they have a case, but their situation clearly does not fulfill one or more of the required criteria. In this situation, they may face the financial and psychological strain of going to court, only to walk away with nothing. Or, a user may expect the potential compensation for damages they can obtain to be much higher than they are in practice, again which in turn leads them to waste significant resources only to come up empty in the end.

An information system may inform such users of the relevant criteria and potential outcomes, thereby moving them towards alternative solutions.

3.6.1.3 Increasing the rate of settlements

A further important effect of such a system could be the increase of settlements of cases. A user may have a much higher chance of settling their case if they understand the legal modalities of their case, such as the rights they have and could achieve in court.

Presenting the opposing party with this information could encourage them to enter into an amicable settlement, rather than going through the arduous court experience. Providing the user with outcome statistics about similar cases could further serve as a BATNA, aligning their expectations and increasing their ability to find a settlement.⁵²⁴

⁵²³ Susskind, *supra* note 22 at 130–132.

⁵²⁴ See 3.5.2.2.

Berman and Hafner analyze the potential impact of predictive expert systems. They argue that the availability of accurate outcome prediction would have the effect of ending many lawsuits. Since people are risk-averse, and legal costs quickly mount to eclipse the original value of lawsuits, parties would be inclined to agree to a computer-predicted settlement.⁵²⁵ Especially legal aid clinics and the public defense bar, where there are many cases and resources are scarce, could benefit immensely from such systems.⁵²⁶

This opportunity could be even more pronounced if the information is integrated into an ODR platform, where users are given the relevant information about rights and possible outcomes, and the ability to propose and accept settlements, all in one place.⁵²⁷

3.6.1.4 Supporting self-represented litigants

If the user decides to go to court, a legal decision support tool could further be very powerful for self-represented litigants who go through the court process without legal representation. A tool that adaptively assesses the situation of the user and provides them with relevant, up-to date information on the substantive and procedural issues of their situation could support these individuals in better preparing their case, and thus improve their experience.

This goal could be achieved to an even greater extent when integrating the system with an ODR platform, where the user could be guided directly towards the relevant procedure, at which point the ODR platform could take care of their case.⁵²⁸

3.6.1.5 Increasing overall welfare

Beyond disputes, such a system could also serve to increase the well-being of individuals, or as Susskind calls it, “legal health promotion”.⁵²⁹ A system could be designed to inform the user of legal opportunities that are available to them, such as obtaining social aid or

⁵²⁵ Berman & Hafner, *supra* note 317 at 932.

⁵²⁶ *Ibid.*

⁵²⁷ See 3.5.2.4.

⁵²⁸ See 3.5.2.4.1.

⁵²⁹ Susskind, *supra* note 22 at 69.

using the law to achieve positive outcomes. This could increase the overall welfare of the user by making these opportunities available to them.

3.6.1.6 Preventing legal problems

This tool could also prevent legal problems, by making it easier for individuals to understand how to conform with the laws, before a conflict has arisen. Further, it could help individuals determine whether certain actions require a permit or license, which could also prevent problems from arising.⁵³⁰

3.6.1.7 Relieving the court system

Furthermore, the tackling of simpler cases outside of the court system could lead to much needed relief in the number of cases at courts, allowing them to focus on more complex issues.⁵³¹ This could decrease delays at the often-overworked courts, and thereby improve the experience for individuals that go through court, as well as saving money for the taxpayer.

3.6.1.8 Increasing Rule of Law

The availability of legal decision support tools could further increase the rule of law in society.

First of all, the rule of law depends on individuals having access to effective mechanisms to resolve their disputes. As we have seen, these mechanisms are often ineffective for individuals today, causing many people to be unable to resolve their issues, and limiting the effective impact of legislation in governing society.⁵³² Building tools that increase the effectiveness of mechanisms to resolve disputes is therefore an important step in increasing the rule of law.

⁵³⁰ Compare *ibid* at 68.

⁵³¹ compare Benyekhlef et al, *supra* note 322 at 10; Garth & Cappelletti, *supra* note 362 at 233.

⁵³² See 3.3.2.1.

Second, as we have seen, the rule of law requires the ability for people to access and understand the rules that govern them, in order to effectively manage their affairs.⁵³³ A tool that allows a user to enter the situation, and receive information on the specific laws that apply to them is could be an excellent way of informing users of the law, and is therefore another important way to increase the rule of law.

3.6.1.9 Legal harmony and trust in legal institutions

People having better access to justice and legal information through technological tools may also have effects on a societal level. A lack of understanding of the legal system can cause societal alienation for individuals. Tools that provide simplified legal information can help individuals orient themselves in the modern, rule-based world, and thus gain a feeling of belonging and inclusion in society.

Additionally, legal problems cause significant issues for individuals. More problems being resolved, and people being able to find amicable solutions to problems, can therefore increase legal harmony between individuals and increase welfare for society. Further, it could increase the trust that people place in the legal system and as an extension society as a whole.

3.6.2 Risks

Of course, not all of the potential effects of legal decision support tools are positive. As Benyekhlef and Vermeys point out, creating digital tools in an area is not a neutral activity. Depending on the choices made, the system will change, as new nudges and rules are introduced into the system.⁵³⁴ Therefore, it is important to understand the current rituals and procedures in the law, and why they exist, so that the impact of digitalizing them can be understood.⁵³⁵

⁵³³ See 3.4.2.

⁵³⁴ Vermeys & Benyekhlef, *supra* note 72 at 3–4.

⁵³⁵ *Ibid* at 4–5.

3.6.2.1 Incorrect information

One negative social impact could be a possible inaccuracy of the algorithm. The tools will be used to provide information to people who might end up relying on them, deciding how to proceed or to settle a case for a specific monetary amount. One has to consider the effects should this information be inaccurate.

Such tools generally have access to a boiled down version of a factual situation in making a decision.⁵³⁶ Further, as discussed in the literature surrounding tools for predicting recidivism, statistical generally tools work by identifying situations with similar characteristics to a new situation, and then deciding the new situation based on the previous occurrences. While accurate on a population level, the specific case may not be assessed correctly.⁵³⁷ As discussed in 2.6.3.5, there are some forms of reasoning that are beyond the current capabilities of AI systems. In Chapter 4, I explore which steps of legal reasoning may require such reasoning. If AI systems are used to perform reasoning that is beyond what they are capable of, the results may not be accurate.

Giving the user incorrect information may lead them to lose their case, e.g. by filing the wrong form, bringing the wrong claim or making bad arguments. It could also lead them to accept a disadvantageous settlement. This could be a significant risk and a source of injustice and raise complicated questions about accountability.

3.6.2.2 Two-tiered justice

When creating legal decision support tools, it is important to ensure that they do not contribute to a two-tiered justice system. This could occur when only certain individuals have access to these decision support tools, giving them an advantage in navigating the justice system. In this case, the existence of these tools may contribute to the inequality

⁵³⁶ Hannes Westermann et al, “Using Factors to Predict and Analyze Landlord-Tenant Decisions to Increase Access to Justice” (2019) Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law (ICAIL ’19) 133–142 at 9.

⁵³⁷ Sonja B Starr, “Evidence-Based Sentencing and the Scientific Rationalization of Discrimination” (2014) 66 Stanford Law Review 71 at 848–850; see also Kelly Hannah-Moffat, “Actuarial Sentencing: An ‘Unsettled’ Proposition” (2012) 30 Justice Quarterly 1–27 at 12.

of access to the justice system that they are intended to mitigate, as only the group that has access to them may enjoy the benefits that flow from the system.

Therefore, it is crucial to make sure that the systems are accessible to as many people as possible. This includes people with limited internet access, people who may not speak English, and people with disabilities.⁵³⁸ Susskind points out that internet access is very widespread, and that even individuals that do not have direct access to the internet can often use the internet through friends and relatives, leaving only 6% that do not have any access.⁵³⁹

3.6.2.3 Freezing the law

Another risk with decision support systems is the freezing of the law. The decision support tool will provide information to individuals that may motivate them to not pursue their case, or settle their case instead of going to court. In general, if the information provided is accurate, this can be seen as a good thing, see above under 3.6.1.2.

However, if such systems become widespread, it may lead to fewer and fewer cases going to court. Judges are able to update the interpretation of laws in their decisions, for example to reflect social changes. This means that if cases do not end up in court, they cannot be used by judges to change the jurisprudence and adapt it to modern developments. Consequently, the legal system could remain frozen in place.⁵⁴⁰

3.6.2.4 Perpetuating bias

As we previously discussed in 2.6.3.3, AI models may pick up biases present in data. This may also occur in legal decision support tools. If a model is trained on case law, it may learn undesirable patterns, such as certain genders or ethnicities being disadvantaged in certain cases. This may be the case even if these attributes are not specifically included

⁵³⁸ Cabral et al, *supra* note 433 at 258–263.

⁵³⁹ Susskind, *supra* note 22 at 216–218.

⁵⁴⁰ Berman & Hafner, *supra* note 317 at 932; Susskind, *supra* note 22 at 289.

in the model, as the other attributes may let the model infer these.⁵⁴¹ For example, if certain types of disputes disproportionately affect certain minorities, any possible bias against them could be embedded in the training data regarding that case type. The algorithms themselves may also contain biases by the programmer.⁵⁴²

A decision support system giving advice to people would perpetuate possible biases contained in the data. Together with the potential freezing of the law this could cause harm to society, as negative biases could be perpetuated and frozen in place.

3.6.2.5 Lack of transparency

As we saw in 2.6.3.4, certain AI models are not transparent. This is the case especially with machine learning models, which can learn complicated models from immense datasets.

This lack of transparency can be an issue in decision support tools providing legal information. If an individual does not know why their situation is assessed in a certain manner, or why a certain case was found to be similar, they may not trust the system. Further, the model may hide biases, that the individual has no way of understanding or compensating for.⁵⁴³

Of course, this issue may be less important in systems such as the one described above, that do not tell the user what to do, but merely support them in their decision making. Likewise, it is important to note that the reason that judges make decisions is also not necessarily transparent – for example, judges may make a decision based on an extra-legal factor, and then write a decision that uses the law to motivate their decision making.⁵⁴⁴

⁵⁴¹ Compare Harry Surden, “The ethics of artificial intelligence in law: Basic questions” (2020) Forthcoming chapter in Oxford Handbook of Ethics of AI 19–29 at 728.

⁵⁴² Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 804; Susskind, *supra* note 22 at 288.

⁵⁴³ Susskind, *supra* note 22 at 288.

⁵⁴⁴ Surden, “The ethics of artificial intelligence in law”, *supra* note 541 at 732.

3.6.2.6 *Focus on case law*

In building legal decision support tools, a likely data source is that of legal decisions. These are frequently openly available and can thus be used to train a machine learning system or to extract statistics. These systems are provided to the user so that they can get an idea for what outcome they can expect in their situation.

An issue with this approach is that cases that end up in court, as we have seen, represent a minority of how people actually solve their disputes. Further, studies have shown that even individuals that do use the court system end up settling in the majority of cases in certain jurisdictions.⁵⁴⁵ In some jurisdictions, only 3% of cases that enter the court system lead to a decision.⁵⁴⁶ In fact, many jurists consider cases that do not settle to be a failure of the legal system.⁵⁴⁷

This means that the cases that are part of the training data are only a tiny sample of how situations were resolved overall. Further, they likely represent a biased sample, since cases that do go through the entire court procedure likely represent the most serious or contentious issues, where the users were unable to find a settlement. This could lead to a so-called “representation bias”, since the majority of the input space is not properly sampled.⁵⁴⁸

This bias may cause predictions to be inaccurate, as the user may not have the kind of case that usually ends up going to court. Further, it may encourage users to consider litigation as the only outcome, since this is the only possible outcome presented in the system.

⁵⁴⁵ Theodore Eisenberg & Charlotte Lanvers, “What is the Settlement Rate and Why Should We Care?” (2009) 6:1 *Journal of Empirical Legal Studies* 111–146; Yun-chien Chang & Daniel M Klerman, “Settlement Around the World: Settlement Rates in the Largest Economies” (2021) 21–8 *USC CLASS Research Paper Series No CLASS21-8*, *USC Legal Studies Research Paper Series*; Samuel R Gross & Kent D Syverud, “Getting to No: A Study of Settlement Negotiations and the Selection of Cases for Trial” (1991) 90:2 *Michigan Law Review* 319–393.

⁵⁴⁶ John Barkai & Elizabeth Kent, “Let’s stop spreading rumors about settlement and litigation: A comparative study of settlement and litigation in Hawaii courts” (2014) 29 *Ohio St J on Disp Resol* 85.

⁵⁴⁷ Gross & Syverud, “Getting to No”, *supra* note 545 at 319.

⁵⁴⁸ Suresh & Guttag, *supra* note 166 at 5.

3.6.2.7 Replacement of lawyers?

Finally, lawyers may be suspicious of legal decision support tools as possible alternative to being hired for legal advice. However, at least with the type of system described in 3.6, this is not the case. The tools do not offer advice or assistance in resolving the exact particular case in question, nor do they provide a forecast of the chances and outcome in court, all of which are typical tasks of a lawyer in a legal matter. Instead, such a legal decision support tool provides general legal information. It informs the user about the content and results of other, similar legal cases, without analyzing these cases and applying them to the particular case of the user. This information may conversely increase the demand for lawyers, as people understand their legal rights and then hire a lawyer to enforce this right.

Moreover, legal support tools like the one that is subject to my research often target users who currently do not have access to lawyers and are forced to represent themselves or cannot solve their problem at all.

3.7 Conclusion

Efficient access to justice is a significant issue in Canada and beyond. Due to costs, temporal delays and the complexity of the legal system, people often do not have access to the mechanisms that they need to effectively solve their issues. Further, a lack of access to legal information may prevent citizens from obtaining the benefits that accrue to them based on the law, or make it impossible to comply with the current legislation.

One way of addressing these issues is the creation of web-based tools that provide the user with legal information and information about previous outcomes of court cases. These could further be coupled with ODR platforms to create quick, cheap and efficient ways of resolving conflicts. I discussed the advantages and risks of using such systems. In my research presented in this thesis, I intend to design a methodology to build such tools.

The building of such tools likely requires the understanding and modeling of how judges and other legal decision makers reason about situations. In the next section, I will explore the different steps involved in this reasoning, as well as how previous researchers automated these steps. This will inform the creation of my methodology to build such tools, which is the research objective of this thesis.

Part II

Automating Legal Reasoning

Chapter 4 Automating Legal Decision Making

Research Objective: Understanding legal reasoning and automating legal reasoning (1.2.2.3)

Research Topics:

- Which steps are involved in legal reasoning?
- Which approaches have researchers previously taken to automate these steps?
- What are the different trade-offs of the approaches previously taken by researchers with regards to my use-case of increasing access to justice?
- How are cases reasoned with in different jurisdictions?

4.1 Introduction

We have now discussed the topic of artificial intelligence, and the issue of access to justice. My research objective is to apply AI to increase access to justice, by informing users of their rights and possible outcomes in court. In order to provide this information, the system will need to incorporate aspects of legal decision making, such as understanding the situation of the user and comparing the situation to previous cases.

To understand how these parts of legal decision making could be automated, let us now take a look at how judges and other legal decision-makers deal with situations that may give rise to legal rights, and how researchers in AI & Law have previously worked to automate the steps of this reasoning process.

Understanding the legal decision-making process is an important step in being able to model this reasoning in a useful way in a digital tool. While it is likely to be difficult to fully automate any of the steps carried out by a judge in deciding on the outcome of a situation, understanding the different steps involved will provide useful context to understand how and at which stage of the reasoning process a decision support tool could be helpful for a potential user.

For each legal reasoning step, I will further delve into what previous work was done to model and automate that specific step. This will help me understand the particularities of automating the different steps, and how I might shape my own methodology in order to build legal decision support tools that can support access to justice.

The steps of legal decision making in this chapter are observed from the perspective of a judge. However, it should be noted that this kind of legal decision making can be assessed by many actors in the legal system and beyond. First of all, different legal systems employ different institutions to perform steps of the process. Common law systems, for example, use a jury to determine the facts of a case. In this chapter, I will refer to the decision maker collectively as “judge” for the sake of brevity.

Likewise, legal reasoning may be performed outside of the courtroom. For example, lawyers are likely to perform legal reasoning to inform their clients of their legal situation. Administrative decision makers employ legal reasoning to determine whether an individual qualifies for certain kinds of social aid, should be granted a certain immigration status or are in accordance with city bylaws. Likewise, the police might have to assess a situation to see whether they are able to perform a seizure or search a specific location. All of these reasoning processes are instances of legal reasoning and necessitate the application of some or all of the steps outlined below. For the sake of brevity, I will refer to the “judge” in the following sections, however this is a stand-in for any individual using legal decision making.

4.1.1 The steps of legal reasoning

In order to explore the different steps involved in legal reasoning, I have split the different steps of legal reasoning performed by a judge or other legal decision makers into the following steps:

- 1. Identification of the guiding legal rule (4.2)** – Based on the claim of the plaintiff and their description of factual events, identify which legal provisions may offer them the desired outcome.

2. **Finding of facts (4.3)** – Determine the material facts of a situation, by taking into account which facts are relevant to the legal provisions, which facts are disputed by the parties and the presented evidence and rules of burden of proof.
3. **Establish fulfilled legal criteria based on facts (4.4)** – Based on the legal facts that were accepted in the previous step assess whether the legal criteria contained in the identified legal rules are fulfilled.
4. **Applying legal rules to legal criteria (4.5)** – deductively reason from the fulfilled legal criteria, to arrive at a decision in a case.
5. **Determining the outcome of a case (4.6)** – sometimes, decided on discretionary consequences to an outcome, such as the amount of damages awarded or the length of a prison sentence.
6. **Explaining the decision (4.7)** – prepare a written document explaining how the decision was reached.

Of course, the exact shape and form of how the legal reasoning steps are carried out differs significantly between the type of instance, type of proceeding and jurisdiction of a decision maker. In courts, for example, common law judges may start with a discussion of the facts, while civil law judges emphasize the legal principles raised by a case.⁵⁴⁹ There will be a discussion of the difference between how the two systems work in 4.9.

I do not claim to have identified the only way to understand the reasoning of a judge. This is a very large research question, and not the focus of this thesis. Rather, I aim to identify a number of steps that seem to be carried out in some capacity in most legal reasoning. This will serve as context for understanding previous work in the domain of automating legal reasoning, and also serve as a conceptual basis for understanding and situating my own research in building legal decision support tools.

⁵⁴⁹ Peter G Stein, “Roman Law, Common Law, and Civil Law Symposium: Relationships Among Roman Law, Common Law, and Civil Law” (1991) 66 Tul L Rev 1591–1604 at 1601.

4.1.2 Automating the steps of legal reasoning

Each of the aforementioned steps rely on different styles of reasoning. Therefore, in order to automate or simulate the different steps, different computational systems need to be built. In this chapter, I will therefore analyze previous work on how to automate the legal decision support tools in the context of the steps presented. For each step, I will give an overview of previous work aiming to tackle that specific aspect of legal decision making, and describe the artificial intelligence methods used.

While there is an enormous variety in the specific technologies employed to build such systems in the field of law, common methods include case-based reasoning systems, rule-based reasoning systems, and machine learning based systems. These systems differ not only by how they are built, but also in what kind of explanations they can give for their decisions.⁵⁵⁰ Before I delve into the specific solutions, I will briefly explain the difference between case-based reasoning, rule-based reasoning and machine learning based approaches to legal reasoning.

Case-based reasoning systems reason using cases that were previously decided by courts. Usually, this involves a way of representing court cases, and a way of comparing new cases to the previous cases, in order to draw analogies between the two. Using this method, these systems can predict the outcome of new cases or generate arguments. The systems usually have a bespoke way of reasoning with cases, such as identifying overlapping facts and comparing cases. It seems to me that they are therefore akin to the symbolic approach, presented above in 2.5. The challenges in case-based reasoning systems involve how to represent court cases in the most useful way, and how to reason with these cases.

Rule-based reasoning systems instead aim to encode legal rules into a computer-readable format. This can be compared to expert systems, described in 2.5.1. For example, the systems can encode the logical structure of the laws governing a specific area and create

⁵⁵⁰ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 3.

a way to deductively reason about new cases, by applying the laws to new cases. The challenges in this area are how to capture the logical rules in an area, how to deal with ambiguity in the syntactic structure of legislation and how to deal with open-textured terms.

Machine learning based reasoning systems typically also rely on cases, like case-based reasoning. However, instead of relying on a symbolic reasoning methodology, they apply machine learning (described above in 2.6) in order to automatically learn a model from the data. This model can then be used for the prediction of new cases. The challenges of using the machine learning approach include how to represent cases for the algorithm, how to build the model and how to explain the results to the user.

4.1.3 The structure of this chapter

In order to analyze the steps of legal reasoning described above, I will go through the legal reasoning steps one by one (4.2 - 4.7). For each step, I give an overview of how the step is carried out by a judge or other legal decision maker, based on the reading of legal doctrine.

I then explain how this step may concretely be carried out in an example case. I will use a fictional example case of a tenant filing a claim for a rent reduction due to the failure of the landlord to heat their apartment for a few days in the cold Montreal winter. This is not intended to be an accurate legal assessment of such a situation, but rather to illustrate each step in legal reasoning.

Finally, for each reasoning step, I analyze prior work in automating that step. I give an overview of the specific methodology developed to tackle that step. I have done a selection of work that is the most relevant for the creation of legal decision support tools. It should be noted that it is not always clear exactly where a prior work should be located, as some research focuses on multiple steps of legal reasoning. In these cases, I have used my best judgment to decide where to place the research or discussed the work in multiple places. For each section I also analyze the promise and challenges of automating the reasoning in that step, and how the challenges were dealt with by the different researchers.

This will serve as the basis for creating my own methodology for building legal decision support tools.

I also briefly touch upon some underlying aspects of the different steps of legal reasoning, such as the interplay between the legislator and judiciary (4.8) and the difference between the styles of precedent used in the common law, civil law, and internally in different courts (4.9).

4.2 Identification of the guiding legal rule

4.2.1 The legal system

Once faced with a dispute, the judge initially has to identify the guiding legal principle that governs a particular dispute. McIntyre refers to this norm as the “dispute-norm”.⁵⁵¹

The judge must here consider the claim made by the plaintiff. The plaintiff will ask for a certain outcome to be awarded. There may be several legal rules that can lead to such an outcome being awarded. The judge must identify the rule that can lead to the outcome requested by the plaintiff, based on the factual situation presented. This allows the judge to identify the facts that are relevant to a case, i.e. the facts that either need to be undisputed, or proven by the plaintiff, in order for the case to succeed in court. In common law systems, these facts are referred to as the *material facts* of the case.⁵⁵²

Since the law is often quite general and abstract, at this stage the judge may have to interpret it to establish the concrete legal rule that they may want to apply to a case.⁵⁵³

Identifying the guiding legal rule is further a crucial step for lawyers advising their clients. Laypeople are not aware of the legal rules, but rather think of their situation in terms of factual occurrences. When interviewing a client, the lawyer therefore has to explore what has occurred, understand the relevance of certain facts, explore which possible avenues

⁵⁵¹ Joe McIntyre, *The Judicial Function: Fundamental Principles of Contemporary Judging* (Singapore: Springer Singapore, 2019) at 99.

⁵⁵² James Holland & Julian Webb, *Learning Legal Rules: A Students' Guide to Legal Method and Reasoning*, 10th ed (Oxford, United Kingdom ; New York, NY: Oxford University Press, 2019) at 226.

⁵⁵³ McIntyre, *supra* note 551 at 104–105.

may exist for obtaining certain remedies, and assess the likelihood of obtaining that remedy. The lawyer thus takes the role of understanding a case and advising the client on the best way to proceed.

In 3.2.2, we saw how important the step of identifying the legal guiding norm can be. In the study conducted by Ab Currie, the individuals that did not act to address a justiciable event that occurred, almost half of individuals thought nothing could be done, were uncertain about their rights or did not know what to do.⁵⁵⁴ Helping these individuals understand that their situation has legal aspects could thus be an important step in helping individuals understand the avenues available to them.

The difficulty of identifying and understanding the relevant legal rule further seems to present a significant impediment to self-represented litigants. According to Branting *et al*, “pro se litigants seldom know what facts they need to establish or how to articulate and organize the facts in a manner that makes their claims amenable to evaluation.”⁵⁵⁵ Susskind mentions the example of self-represented litigants arriving in court with a bag of un-indexed documents.⁵⁵⁶ These individuals could likely benefit significantly from being shown the specific legal rules that apply in their cases, to help them prepare the relevant arguments and proof.

Sifting through a factual situation and classifying it in terms of a legal situation is thus a crucial step for both judges, lawyers and parties. However, this step is not intuitive for laypeople, who may not have the skills necessary to understand which facts are relevant, or which legal rules the facts of a situation may correspond to.

4.2.2 Example

Let us imagine that the judge is faced with a claim for rent reduction based on an apartment in Montreal being improperly heated for several days in December.

⁵⁵⁴ Currie, *supra* note 297 at 56.

⁵⁵⁵ Branting et al, *supra* note 348 at 218.

⁵⁵⁶ Susskind, *supra* note 22 at 122.

The judge looks to the legislation and finds that there are a few articles in the Civil Code of Quebec⁵⁵⁷ that could be relevant for the heating issue of the tenant. Article 1854 states:

*The lessor is bound to deliver the leased property to the lessee in a good state of repair in all respects and to provide him with peaceable enjoyment of the property throughout the term of the lease.*⁵⁵⁸

Article 1863 includes the following passage:

The nonperformance of an obligation by one of the parties entitles the other party to apply for, in addition to damages, specific performance of the obligation in cases which admit of it. He may apply for the resiliation of the lease where the nonperformance causes serious injury to him or, in the case of the lease of an immovable, to the other occupants.

*The nonperformance also entitles the lessee to apply for a reduction of rent; where the court grants it, the lessor, upon remedying his default, is nonetheless entitled to the re-establishment of the rent for the future.*⁵⁵⁹

The judge sees that the aforementioned article gives the tenant the right to a rent reduction, and therefore corresponds to their claim. Further, the non-heating of an apartment may correspond to the legal concept of the peaceable enjoyment of the property not being fulfilled. The judge has thus identified a possible guiding legal principle for the case and can use this to determine which of the facts of the case are material and relevant to this provision. Thus, it is likely that the judge will assess whether the fact of the apartment not being heated for several days is proven or undisputed (see 4.3). However, facts such as the color of the walls in the apartment are not relevant under the identified legal guided rule, and can thus be ignored by the judge.

⁵⁵⁷ *Civil Code of Québec*, CQLR c CCQ-1991.

⁵⁵⁸ *Ibid* Article 1854.

⁵⁵⁹ *Ibid* Article 1863.

4.2.3 Automating identification of the legal rule

4.2.3.1 Introduction

What would it take for the step of automation of the identification of the legal rule governing a case?

As we have seen, laypeople are likely to understand their situation in terms of a factual occurrence (“My apartment is very cold”) instead of the governing legal rule (“The landlord has failed to perform their obligation”). Likewise, they may understand their situation in terms of a need (“I wish to receive a rent reduction”). Therefore, it would be useful for legal decision support systems to have a facility to assess the situation of a user and guide them towards exploring the relevant legal issues.

4.2.3.2 Prior work

In the literature, the identification of the guiding legal rule does not seem to have been discussed in-depth. Most systems are designed to tackle a specific issue. Therefore, identifying the guiding legal rule is less important.

Nonetheless, there has been some research aiming to address the identification of relevant legal areas or norms based on facts. David Colarusso has developed a system that analyzes non-lawyer language and translates it into a standard taxonomy of legal issues (LIST). The system takes plain language descriptions of facts, analyzes the text, and returns potential issues that are present, such as “problems with living conditions”.⁵⁶⁰ The system is trained using data from an online game, that asks users to answer questions about layperson fact descriptions, and previous data from users of the system. It is able to recognize 105 labels.⁵⁶¹

The Loge-expert project aims to give information about landlord-tenant disputes to laypeople. The authors acknowledge the difficulty of understanding the structure of lawyer-client interviews. They aim to observe such interviews in order to model a system

⁵⁶⁰ David Colarusso, “Machine-Assisted Issue Spotting for Self-Represented Litigant Portals” (2022) at 1–2.

⁵⁶¹ *Ibid* at 3.

able to support laypeople.⁵⁶² In the end, the chosen approach seems to be to allow the individual to select an issue from a list of issues that the system can treat, including “repair during the course of the lease”, “rent reduction” and “repossession”.⁵⁶³ This should work well if the user is aware that they are dealing with an issue of repossession, or that they would like to obtain a rent reduction.

Branting described what he refers to as advisory system for pro se litigants. One of the aims of these systems is to inform the user of the available forms of legal relief. A component of such systems is helping litigants that have relatively specific goals, by informing them of the available forms of legal relief. Such information can be presented in a tutorial system, where the system controls the information flow, or hypertext, where the user controls how they wish to see the information.⁵⁶⁴ He implements such a system in the domain of protection orders.⁵⁶⁵

4.2.3.3 Conclusion

As we can see, the identification of the legal guiding rule in legal decision support tools is an important, yet not fully explored, area. Tools tend to focus on capturing the legal rules in a specific legal area. Sometimes, this may be sufficient – if a tenant receives a letter of eviction, they are likely to know that their case is in the area of eviction. However, in other cases, the individual may simply face a factual situation, and not be aware of the legal aspects or avenues available to them. Likewise, a user with a specific need may not be able to identify the laws that could support them in pursuing this need. In these cases, a system that allows a user to explore the relevance of their factual situation and possible legal avenues may be a useful component of automating legal decision support.

⁵⁶² Claude Thomasset & Louis-Claude Paquin, “Expert Systems in Law and the Representation of Legal Knowledge: Can we Isolate it from the Why and the Who?” (1989) III International Congress: Logica, Informatica, Diritto: Expert Systems in Law 751–772 at 15.

⁵⁶³ Louis-Claude Paquin, François Blanchard & Claude Thomasset, “Loge–expert: from a legal expert system to an information system for non-lawyers” (1991) Proceedings of the 3rd international conference on Artificial intelligence and law (ICAIL '91) 254–259 at 255.

⁵⁶⁴ Branting, *supra* note 78 at 2–3.

⁵⁶⁵ *Ibid* at 4–6.

4.3 Finding of facts

4.3.1 The legal system

Once the judge has established that a legal guiding rule exists that could lead to the outcome that was claimed by the plaintiff, they must establish the version of the factual events that they lay at the basis of their decision. As the judge has now determined a legal guiding principle, they have determined which of the facts are material, and thus of relevance to the desired outcome. Only these facts need to be considered in order to reach a conclusion.⁵⁶⁶ In cases where litigants are self-represented, they may not be aware of which facts are relevant for the court. Here, the judge may intervene to only consider evidence that has a material bearing on the case.

How exactly the facts are determined for the judicial process depends on the legal system. In adversarial systems, such as common law systems and some civil law systems, the parties are completely in charge of the procedure, often with the support of professional advocates.⁵⁶⁷ If the parties agree on a certain factual occurrence, the judge will accept this, and not consider any evidence.

However, if the parties disagree on a certain fact, the judge needs to determine which version of the facts should be seen as proved for the purposes of the case.⁵⁶⁸ They do this by considering the evidence presented with regards to the different versions of the facts, and assessing their probative value.⁵⁶⁹ This evidence can be, for example, party, witness and expert testimony, documents, pictures or videos. The judge must also consider the rules of “burden of proof” and “standard of proof”, and potential rules of evidence.

The *burden of proof* determines which of the parties is responsible for proving a certain fact. This depends on the legal system, but often falls to the party that introduces a certain

⁵⁶⁶ Holland & Webb, *supra* note 552 at 143.

⁵⁶⁷ John H Farrar & Anthony M Dugdale, *Introduction to Legal Method*, 3rd revised ed (London: Sweet & Maxwell, 1990) at 62–63; Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1599.

⁵⁶⁸ Holland & Webb, *supra* note 552 at 141.

⁵⁶⁹ “probative value” in *Wex* (LII / Legal Information Institute, 2020).

fact to the trial.⁵⁷⁰ A party that wishes to obtain a remedy for an action performed by the other party, for example, must likely prove that the other party performed this action. If the evidence is not sufficient to prove this, the judge would assume that the action did not occur and deny the remedy.

The *standard of proof* determines to which level of certainty a party must prove a fact. This often depends on the type of case – in civil cases, it is usually sufficient that a party manages to convince the judge that their version of the facts is more likely than the opposing version. In criminal cases, however, the prosecution must usually prove that the defendant has committed a crime beyond any reasonable doubt, which is a much higher standard.⁵⁷¹

Likewise, many jurisdictions have strict rules of what kind of evidence can be introduced to a court room in the first place.⁵⁷² For example, testimony relating to hearsay will often be excluded.⁵⁷³

In inquisitorial systems, such as certain civil law systems, on the other hand, the judge plays a much more active role in the process, and is able to ask questions of the parties.⁵⁷⁴

Dahlman and Mackor analyze models for understanding the fact-finding process. They discuss the *coherence* approach, which relies on constructing stories of the factual situation that occurred and assessing how coherent the evidence is in relation to the story and general knowledge about the world.⁵⁷⁵ *Reductionism*, on the other hand argues that

⁵⁷⁰ Francis Chapman, *Principles of the Law of Evidence with Illustrative Cases* (Philadelphia: Cyrus M. Dixon, 1930) at 301–304; J P McBaine, “Burden of Proof: Degrees of Belief” (1944) 32:3 Calif L Rev 242–268.

⁵⁷¹ Ronald J Allen & Alex Stein, “Evidence, Probability, and the Burden of Proof” (2013) 55 Arizona Law Review 557–602 at 558–560; McBaine, “Burden of Proof”, *supra* note 570.

⁵⁷² See e.g. *Canada Evidence Act*, RSC, 1985, c C-5 2022 Last Modified: 2019-07-12.

⁵⁷³ “Hearsay” in *Wex* (LII / Legal Information Institute, Cornell University).

⁵⁷⁴ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1599.

⁵⁷⁵ Christian Dahlman & Anne Ruth Mackor, “Coherence and probability in legal evidence” (2019) 18:4 Law, Probability and Risk 275–294 at 276–284.

coherence can be expressed purely in terms of Bayesian probabilities.⁵⁷⁶ The authors disagree over which one of these approaches is preferable.⁵⁷⁷

At the end of this process, the judge will have established a version of the factual occurrences that they will be able to base the rest of their analysis on. This version of the facts may not correspond to what actually happened – it reflects what the party with the burden of proof was able to prove and can thus be influenced by rhetoric and persuasion by the parties.⁵⁷⁸

4.3.2 Example

Let us consider our example. The tenant claimed that the landlord failed to properly heat their apartment for several days. If the landlord agrees that this was the case, the judge can use this as the basis of their decision and go on to the next steps of determining whether this counts as failing to give the tenant a peaceable enjoyment of the property. If the landlord disagrees with the fact of the building not being heated, however, the judge must examine the evidence presented by both parties.

The Code Civil du Quebec includes a rule regarding the attribution of the burden of proof in article 2803:

A person seeking to assert a right shall prove the facts on which his claim is based. [...]

It also includes rules regarding the standard of proof in Article 2804:

Evidence is sufficient if it renders the existence of a fact more probable than its non-existence, unless the law requires more convincing proof.

In our example case, the tenant wishes to assert a right based on the fact of the heating of their apartment not working. According to article 2803, it would thus seem like they must

⁵⁷⁶ *Ibid* at 285–289.

⁵⁷⁷ *Ibid* at 292.

⁵⁷⁸ Neil MacCormick, *Legal Reasoning and Legal Theory* (Clarendon Press, 1994) at 27; Holland & Webb, *supra* note 552 at 143.

prove the facts of the heating not working (*burden of proof*). Based on article 2804, they would need to make the fact of the heating not working more probable than the heating working (*standard of proof*).

The tenant can undertake this by presenting protocols of visits by city inspectors, pictures or logs of thermometer readings or witness testimony by individuals who visited the apartment. The landlord may provide evidence that they did heat the building, including heating bills or their own inspections and readings.

If this evidence renders the fact of the heating not working more probable than the opposite, the judge will make this fact the basis of their decision. If, however, the evidence is not clear, the tenant would have failed to discharge their burden of proof, and the judge would side with the landlord.

4.3.3 Automating the finding of facts

4.3.3.1 Introduction

As we can see, reasoning with evidence is a crucial aspect of establishing the facts of a decision, that is then used to justify the rest of the decision. For many cases, especially in cases of high volume and low intensity, where the legal classification is not very complex and the rules are quite clear, the evidence stage of decision making is likely to be the deciding factor in many cases. It is thus a very important step in legal decision making.

4.3.3.2 Prior work

There have been a number of papers that aim to use artificial intelligence to analyze the evidence presented in a case.

Vern Walker, for example, designed a system that allows the capturing of rules of evidence in a rule-based system. At the base level, it captures the rules of a certain legal test in the form of an implication tree.⁵⁷⁹ This system will be described in 4.5.3, relating

⁵⁷⁹ Vern R Walker et al, “Representing the Logic of Statutory Rules in the United States” in Michał Araszkiewicz & Krzysztof Pleszka, eds, *Logic in the Theory and Practice of Lawmaking* (Cham: Springer International Publishing, 2015) 357; Vern R Walker, “A default-logic paradigm for legal fact-finding”

to the encoding of legal rules. The system can also be used to structure the evidence of a case, by linking plausibility schemas to the terminal propositions of the implication tree.⁵⁸⁰ This structure has been applied to case law, in order to capture sentences that discuss certain types of evidence. Ashley and Walker applied it to the evaluation of evidence of a factfinder in the domain of Vaccine Injury Compensation.⁵⁸¹

Satoh *et al* developed another rule-based system that is able to deal with evidential reasoning. This system is called PROLEG and provides a logical representation of the Japanese Presupposed Ultimate Fact Theory (JUF theory).⁵⁸² This theory places a burden of proof for each condition on one of the parties, enabling judges to make decisions in the absence of sufficient evidence to prove a fact.⁵⁸³ Further, certain conditions are seen as *open*, and are only assessed if they are explicitly raised by a party of a process.⁵⁸⁴ The system provides the options for a party to allege a fact and provide evidence to prove that fact. The other party can admit that a fact is true, challenge the veracity or introduce a counterargument. If the judge believes a certain fact, they can deem it plausible.⁵⁸⁵ The reasoning process in PROLEG is based upon backward-chaining, where the system first tries to fulfill the stated goal of a plaintiff, by checking the different encoded requirements. If this check succeeds, the system will assess whether any exceptions alleged by the other party are present. Unless that is the case, the plaintiff will be said to win a case.⁵⁸⁶

(2007) 47:2 *Jurimetrics* 193–243; Vern R Walker, “Representing the use of rule-based presumptions in legal decision documents” (2014) 13:3–4 *Law, Probability and Risk* 259–275.

⁵⁸⁰ Walker, *supra* note 579 at 210–219.

⁵⁸¹ Kevin D Ashley & Vern R Walker, “Toward constructing evidence-based legal arguments using legal decision documents and machine learning” (2013) *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law - ICAIL '13* 176 at 177.

⁵⁸² Satoh et al, “PROLEG”, *supra* note 48.

⁵⁸³ *Ibid* at 1.

⁵⁸⁴ Satoh et al, “PROLEG”, *supra* note 48.

⁵⁸⁵ *Ibid* at 158–159.

⁵⁸⁶ *Ibid* at 159–160.

The PROLEG system was used to implement the rules in several areas, such as choice of jurisdiction in private international law.⁵⁸⁷ It was also used as a system to arrange issues, to allow the system to automatically identify points where evidence must be assessed by the judge. This is done by asking the plaintiff for the desired conclusion and showing the possible paths and facts necessary to achieve this conclusion. Then, it allows the defendant to either introduce the facts supporting an exception or denying one of the presented facts. The contentious issues are marked as “issues to be determined” and will then be decided by the judge.⁵⁸⁸

Shaikh *et al* aimed to predict whether a murder case would lead to the acquittal or conviction of an individual. They annotated 86 cases in terms of 19 factors, including whether the evidence was ocular or circumstantial, whether the motive was established and the number of witnesses of different types.⁵⁸⁹ Shaikh *et al* used this dataset to predict the outcome of cases, achieving high F1-scores of 0.91.⁵⁹⁰

4.3.3.3 Conclusion

The first two presented systems allow the analysis of evidence discussed in case law, and the structuring of cases while considering the burden of proof of the parties, which can be very useful. Shaikh *et al* analyze cases in relation to factors relating to evidence, such as the number of witnesses and the type of evidence, and whether the testimony was contradictory.

Overall, there seem to be few systems that try to directly model whether a specific piece of evidence is sufficient for a certain fact to be seen as proven or are able to give

⁵⁸⁷ Ken Satoh, Laura Giordano & Matteo Baldoni, “Implementation of Choice of Jurisdiction and Law in Private International Law by PROLEG Meta-interpreter” (2021) *Logic and Argumentation (Lecture Notes in Computer Science)* 60–75.

⁵⁸⁸ Satoh, Takahashi & Kawasaki, *supra* note 506.

⁵⁸⁹ Rafe Athar Shaikh, Tirath Prasad Sahu & Veena Anand, “Predicting Outcomes of Legal Cases based on Legal Factors using Classifiers” (2020) *167 Procedia Computer Science (International Conference on Computational Intelligence and Data Science)* 2393–2402 at 2397.

⁵⁹⁰ *Ibid* at 2400.

recommendations on which specific pieces of evidence an individual needs to produce in order to prove a fact. I will briefly elaborate why this may be the case.

First of all, the data available in the form of case decisions may not sufficiently cover the necessary evidence to be able to build up a sufficient database for statistical analysis. For one, evidence may not be described in detail in case decision documents. Further, even in cases that do describe the evidence, the description is by necessity a pared down, summarized version of what was presented in the court room. For example, important factors that the judge may rely on (e.g. demeanor of witnesses and parties and the full evidence presented in the cases such as pictures, documents etc) are not included in full in the document.⁵⁹¹ A system relying on encoded versions of evidence present in a case may therefore not be able to replicate the full process that the judge is able to perform to assess the validity of evidence.

Second, even if the data were available, it is unlikely that evidentiary reasoning would be feasible with today's artificial intelligence systems. As we have seen, evidence does not have a general value, but is rather relevant to support a specific story or narrative of the parties. The factfinder uses common sense and their experience of the world to establish the value of a specific piece of evidence, in order to assess the likelihood of a fact given certain pieces of evidence. As described in 2.6.3.5, machine learning systems lack common sense, and are thus unlikely to be able to use varied pieces of evidence, such as photographs, documents and witness testimony, to assess the likelihood of certain events having occurred. This task may thus be a so-called AI-complete task.⁵⁹²

⁵⁹¹ Westermann et al, *supra* note 536 at 141.

⁵⁹² See 2.4.2.4.

4.4 Establish fulfilled legal criteria based on facts

4.4.1 The legal system

The judge has now determined a version of the facts that they can lay at the basis of their decision. The next step is to determine whether the legal criteria (here used in the sense of requirement or prerequisite) that are part of a rule can be seen as applying to the facts or not. The criteria can stem from statutory law or judge made law, which is especially prevalent in common law systems. They can also stem from guidelines of administrative bodies, that specify how a decision maker should decide whether to e.g. award social aid or not.

In a way, these legal criteria can be seen as a classification system, that capture certain factual situations but ignore other situations.⁵⁹³ In our example, the judge must consider whether the landlord has provided the tenant with “peaceable enjoyment” of the property. There could be many different situations that correspond to this criterion, but also many that do not.

The assessment of whether a legal criterion is fulfilled is not as straightforward as it might seem. A famous example used by Hart to illustrate this is that of a sign in a park, prohibiting vehicles from driving in that park.⁵⁹⁴ Here it is quite obvious that cars would not be allowed to drive in the park. Hart refers to such cases as being in the core of the rule, where the answer of whether the rule applies or not is clear. However, what about bikes, radio-controlled cars, ambulances or wheelchairs? Reading only the text of the sign does not give us clarity in whether these are allowed or not.⁵⁹⁵ They reside in the penumbra of the rule.⁵⁹⁶ Here, the judge must make a decision based on factors such as the intent of the rule, analogies to other legal domains and precedential cases or policy

⁵⁹³ Edward H Levi, “An Introduction to Legal Reasoning” (1948) 15:3 The University of Chicago Law Review 501–574 at 520.

⁵⁹⁴ Hart, *supra* note 149 at 607.

⁵⁹⁵ *Ibid.*

⁵⁹⁶ *Ibid.*

arguments.⁵⁹⁷ Susskind believes that deciders frequently make use of non-rule standards to decide on the applicability of open-textured terms, such as “principles, policy, and purpose through to political morality, social justice, personal preference, and mere whim.”⁵⁹⁸

While the term “vehicle” in the rule “no vehicles in the park” may seem clear and straight-forward, the previous paragraph shows that the term is, in-fact, “open-textured”. The judge has a significant amount of freedom in how the facts of a case should be interpreted with regards to this open-textured requirement. Since it is impossible to encode all situations that may potentially arise in legislation,⁵⁹⁹ the open-textured nature of the law allows it to remain flexible and adapt to new situations, according to the view advanced by Hart.⁶⁰⁰ Not all legal scholar agree with this view – Dworkin, for example, advances the view that there is a “right” answer in applying open-textured concepts, that can be arrived at through reasoning and argumentation.⁶⁰¹

As the judge applies the law to a new set of facts, this can also alter the meaning of the law, as we would expect future judges to come to similar conclusions. Ashley & Rissland compare the legal system to a learning system, that adapts when faced with new situations.⁶⁰²

In common law systems, interpreting how previous judges dealt with specific facts and how they relate to a certain legal concept is a crucial part of the judicial reasoning

⁵⁹⁷ Holland & Webb, *supra* note 552 at 149–151.

⁵⁹⁸ R E Susskind, “Expert systems in law: out of the research laboratory and into the marketplace” (1987) Proceedings of the 1st international conference on Artificial intelligence and law (ICAIL '87) 1–8 at 3.

⁵⁹⁹ Ashley, *supra* note 44 at 40.

⁶⁰⁰ Holland & Webb, *supra* note 552 at 133.

⁶⁰¹ T J M Bench-Capon & M J Sergot, “Toward a Rule-Based Representation of Open Texture in Law” in Charles Walter, ed, *Computer Power and Legal Language: The Use of Computational Linguistics, Artificial Intelligence, and Expert Systems in the Law* (New York: Praeger, 1988) 39 at 44–45; see also Lon L Fuller, “Positivism and Fidelity to Law--A Reply to Professor Hart” (1957) 71:4 Harv L Rev 630–672.

⁶⁰² Kevin D Ashley & Edwina L Rissland, “Law, learning and representation” (2003) 150:1 Artificial Intelligence (AI and Law) 17–58 at 18.

process.⁶⁰³ Lawyers here have a lot of discretion in arguing how previous cases should be interpreted and which cases are the most important.⁶⁰⁴

4.4.2 Example

Let us consider how this step might apply to the *example* case we discussed. Looking at the aforementioned legislation, it seems that the important legal criterion in this case is the criterion of whether the landlord has provided the tenant with *peaceable enjoyment of the property*. This is an open-textured term – it is not explained exactly what is meant by peaceable enjoyment. The judge may thus make use of previous relevant case law, the intention of the law and policy considerations, to decide whether the situation based on the proved or undisputed facts of the case should be seen as the landlord having provided the peaceable enjoyment of an apartment or not.

If the apartment was cold for several days in a row, requiring the wearing of outdoor clothing and making it difficult to sleep, the judge may decide that the situation does not correspond to the landlord having provided peaceable enjoyment of the property. Overall, there could be myriads of situations that do correspond to this criterion, and an equal number that do not.

4.4.3 Automating the determination of fulfilled legal criteria

4.4.3.1 Introduction

Determining whether certain legal criteria are fulfilled or not is an important step in legal reasoning. Automating this step could further have a significant impact on access to justice. Currently, even after identifying the applicable law, laypeople may still not be aware of how the law would concretely apply in their case. For example, is driving a scooter through the park prohibited? What does the term “peaceable enjoyment” mean in practice? The answers to these questions can have a significant impact on the rights and obligations imposed on individuals.

⁶⁰³ Rissland & Ashley, “HYPO”, *supra* note 47 at 2.

⁶⁰⁴ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1600.

Building a system that would help laypeople understand how the relevant legal criteria may apply in their case could thus give them the tools to understand the law and act accordingly. For this to work, however, a way has to be found to overcome the open-textured nature of the law to determine whether a certain legal criterion applies in a case or not.

There has been a significant amount of research attempting to automate the analysis of whether certain facts correspond to certain legal criteria. Directly determining whether a factual situation will be seen as fulfilling a situation or not is very difficult. How would a computer system know, for example, which factual situations are “reasonable” or constitute “peaceable enjoyment” of an apartment? The judge, of course, can rely on their common sense and understanding of policy, the world and other factors to reach these decisions, even in the absence of previous case law. AI systems, on the other hand, may currently not have these faculties,⁶⁰⁵ and thus be unable to directly decide whether a factual situation fulfills a certain legal criterion.

Many approaches rely on encoding previous judicial decisions.⁶⁰⁶ These decisions give an indication of whether judges in previous cases found that a certain set of facts correspond to a certain legal criterion or not. This data can thus be analyzed to attempt to find patterns that tell us which facts correspond to a criterion being fulfilled. Further, previous cases may be an important part of legal argumentation. Especially in common law systems, identifying cases that are in line with a desired conclusion, and distinguishing cases that are not, is an important part of crafting arguments for lawyers. Thus, let us take a look at how case law has been used to determine whether legal criteria apply or not.

⁶⁰⁵ See 2.6.3.

⁶⁰⁶ Ashley, *supra* note 44 at 73.

4.4.3.1.1 Representation

In order to make use of previous case law, the first step is to find a way to encode the previous cases into a representation that the computer understands, as discussed above in 2.6.1.2. In this section, I will take a look at three types of representations.⁶⁰⁷

- **Textual Representations** (4.4.3.2) rely on understanding previous legal decisions in terms of the written text of the decision.
- **Metadata representations** (4.4.3.3) rely on encoding the metadata of a decision, such as the judge, jurisdiction or the parties that are present in a case.
- **Merit-based representations** (4.4.3.4) rely on understanding a case in terms of the merits, e.g. the facts that are present in a case.

4.4.3.1.2 Comparison/Analysis

Once the cases have been represented for the computer to understand, the systems need a way to compare cases, in order to accomplish a useful task. Since we want to build a system that works to increase access to justice, for us this task would likely be to compare the previous cases to a new hypothetical case, in order to provide information to a layperson. Ideally, the layperson could thus enter the facts of their own case, and receive information regarding previous cases, that could help them better understand how a judge may consider their case in terms of fulfilled legal criteria. As we will see, not all of the aforementioned representations support this task.

There seem to be three ways of performing this comparison that are present in prior research:

- **Case-based reasoning approaches** – these approaches rely on custom symbolic reasoning approaches (see 2.5) to deal with the encoded cases, in order to predict the outcome of future cases or generate arguments for either side.

⁶⁰⁷ Hannes Westermann, “Automating Reasoning with Previous Judicial Decisions” in Karim Benyekhlef, ed, *AI and Law: a Critical Overview* (2020) 189 at 206.

- **Rule-based approaches** – these approaches rely on encoding cases with regards to the rules that can be inferred from the way these cases were decided (compare 2.5.1). They can then be used to assess how a criterion may apply in a new case.
- **Machine learning based approaches** – these approaches rely on feeding the representations of cases to a machine learning algorithms (see 2.6), that can then potentially predict the outcome of new cases.

4.4.3.2 Textual representations

4.4.3.2.1 Introduction

Let us first examine research using textual representations of case law. Here, the text of a previous legal decision serves as the representation of the case. There are multiple ways to turn the text of a decision into a vector representation. The easiest approach is the bag-of-words approach, where each case is represented by a list of the words that appear in that case. For example, a representation of certain cases may look like this:

	The	Decision	Damage
Case 1	1	0	1
Case 2	0	1	1
Case 3	0	0	1

Here, the first case contains the words “The” and “Damage”, the second case contains the words “Decision” and “Damage”, etc. In a real-world scenario, the table would contain thousands of rows, corresponding to all words that appear in the cases. This example is a simplified version – multiple ways exist to enhance the usefulness of this kind of representation by embedding statistical information as to the frequency of words and the order that words appear in.⁶⁰⁸

The representation can then be fed to a machine learning algorithm in order to predict the outcome of new decisions, based upon their textual representation.

⁶⁰⁸ See 2.6.1.2.3.

4.4.3.2.2 Prior work

Perhaps the most famous research using this kind of representation comes from Aletras *et al.*⁶⁰⁹ The researchers aim to predict the outcome of cases before the European Court of Human Rights (ECtHR), relating to Articles 3, 6 and 8 of the European Convention on Human Rights (ECHR). In order to achieve test this capability, they built a database of 584 decision texts relating to these articles.⁶¹⁰

The authors used the 2000 most frequent n-grams (i.e. words in groups of 1-4), extracted from different sections of the decisions, as input to a machine learning system. They also used statistical methods to generate topics, i.e. n-grams that frequently appear in similar contexts.⁶¹¹ They then tested the ability of the model to predict the outcome of decisions that the model has not seen during training, achieving an accuracy of between 78% and 84% for predicting the outcome of some of the cases.⁶¹²

There have been several other researchers performing work in similar veins. Medvedeva *et al.*, for example, retrieved a larger corpus of 1942 decisions from the ECtHR, for 9 articles of the ECHR.⁶¹³ These decisions were also represented as n-grams,⁶¹⁴ and used to train a machine learning algorithm to predict whether a situation violates an article of the European Convention of Human Rights, achieving accuracy scores of between 0.61 and 0.83 for the different articles.⁶¹⁵

Sulea *et al.* used machine learning to predict the legal area and outcome of cases of the French supreme court.⁶¹⁶ The authors built a dataset of 126,865 decisions from the

⁶⁰⁹ Aletras *et al.*, “Predicting judicial decisions of the European Court of Human Rights”, *supra* note 56.

⁶¹⁰ *Ibid* at 4–8.

⁶¹¹ *Ibid* at 8–9.

⁶¹² *Ibid* at 9–10.

⁶¹³ Medvedeva, Vols & Wieling, *supra* note 56 at 247.

⁶¹⁴ *Ibid* at 249–252.

⁶¹⁵ *Ibid* at 250–253.

⁶¹⁶ Octavia-Maria Şulea *et al.*, “Predicting the Law Area and Decisions of French Supreme Court Cases” (2017) Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017 716–722.

French supreme court (Court de Cassation),⁶¹⁷ and achieved an f1-score of 0.96 in predicting a case ruling.⁶¹⁸

Chalkidis *et al* build a model to predict the outcome of decisions of the European Court of Human Rights using neural models.⁶¹⁹ Lage-Freitas *et al* used a similar approach, using a textual representation of Brazilian court decisions and different machine learning algorithms to predict the outcome of the decisions.⁶²⁰ Salaün *et al* used machine learning to predict the outcome of landlord-tenant disputes in Quebec. They found that the models were able to predict the outcomes of the decisions using the text of the decision, achieving f1-scores of 93.7% for the best models.⁶²¹

4.4.3.2.3 Conclusion

As we can see, it is possible to create machine learning models that predict the outcome of cases, and whether certain legal criteria are fulfilled, based on the text of a decision. Creating these representations is relatively easy, as computers are able to read thousands of documents per minute.

4.4.3.2.3.1 Predicting undecided cases?

However, it can be questioned whether this kind of representation is useful for predicting unseen cases,⁶²² such as a situation faced by an individual that has not yet gone to court. Legal decision texts are created by the judge, and are thus not a neutral representation of the case, as the judge may have already arranged and selected what to present in the decision.⁶²³ This means that the models may not be useful to support laypeople, since

⁶¹⁷ *Ibid* at 2.

⁶¹⁸ *Ibid* at 1.

⁶¹⁹ Chalkidis, Androutsopoulos & Aletras, *supra* note 56.

⁶²⁰ André Lage-Freitas et al, “Predicting Brazilian Court Decisions” (2022) 8 PeerJ Comput Sci e904.

⁶²¹ Olivier Salaün et al, “Analysis and Multilabel Classification of Quebec Court Decisions in the Domain of Housing Law” (2020) Natural Language Processing and Information Systems (Lecture Notes in Computer Science) 135–143.

⁶²² Medvedeva, Vols & Wieling, *supra* note 56 at 238.

⁶²³ Aletras et al, “Predicting judicial decisions of the European Court of Human Rights”, *supra* note 56 at 12.

laypeople describe their cases very differently from how judges would describe a case in the legal decision documents.

This difference was confirmed by Branting *et al*, who trained a machine learning system on attorney misconduct complaints, which are written and submitted to the bar association directly by citizens.⁶²⁴ The researchers used machine learning methods to predict whether a case would be closed or investigated further, based on the submissions written by laypeople. The accuracy of this prediction was very low.⁶²⁵ The researchers believe: “the root problem is that pro se litigants seldom know what facts they need to establish or how to articulate and organize the facts in a manner that makes their claims amenable to evaluation.”⁶²⁶ From this research, it seems like machine learning methods predicting the outcomes of cases directly from text may be less useful when aiming to support laypeople, or predict cases that did not yet go to court.

4.4.3.2.3.2 *Using the models to generate explanations?*

However, even if the prediction only works with the summaries written by judges, these models may still be a useful way to obtain a list of facts that are likely to affect the decision of a judge one way or another. Several techniques have been developed to analyze which words the machine learning models base their decisions on.

However, many of the researchers report that the trained models do not seem to rely on legally relevant text snippets in their predictions. In Medvedeva *et al*, for a certain article, “state attorney office” was the most associated with a prediction of violation, while “district prosecutor office” was the most associated with a prediction of non-violation.⁶²⁷ Likewise, Sulea *et al* note: “[...] the word bigrams and trigrams deemed to be the most salient in predicting the ruling are not actually tied to any factual information particular to

⁶²⁴ Branting *et al*, *supra* note 348 at 216.

⁶²⁵ *Ibid* at 217.

⁶²⁶ *Ibid* at 218.

⁶²⁷ Medvedeva, Vols & Wieling, *supra* note 56 at 254.

one case, but more related to formulaic expressions typical for a particular ruling.”⁶²⁸
Chalkidis *et al* came to a similar conclusion.⁶²⁹

Branting *et al* investigated whether the sentences focused on by a machine learning system could be used to support the reading of a decision.⁶³⁰ To test this hypothesis, they asked 61 participants to predict the outcome of a case. In certain cases, the participants were shown highlighted sentences based on the sentences that a machine learning algorithm found useful in predicting the outcome of the case. The researchers then assessed whether seeing the cases with these highlights made the understanding of the cases quicker or more accurate. Overall, the researchers found that the system was not able to identify sentences that helped the individuals make better predictions. In fact, many of the participants struggled to see a connection between the highlighted sentence and the issue at hand.⁶³¹

Using a purely textual representation of cases is a quick way to create a representation of cases, that can then be used to predict the outcome of cases. In this way, it is possible to use huge datasets of cases. However, as we have seen, the systems predicting the outcome of cases based on legal decision texts do not always seem to learn legally relevant factors from the text and may rely on the linguistic cues provided by a judge to deliver the predictions, potentially making them less useful for our case of supporting laypeople understand their case.

⁶²⁸ Şulea *et al*, *supra* note 616 at 6.

⁶²⁹ Chalkidis, Androutsopoulos & Aletras, *supra* note 56.

⁶³⁰ Branting *et al*, *supra* note 54 at 218.

⁶³¹ *Ibid* at 219–221.

4.4.3.3 *Metadata representations*

4.4.3.3.1 Introduction

Another way to represent court cases is to capture the metadata of a case, such as the judge presiding over a case, which parties were involved in a case or the procedural motions filed in a case.

A representation of cases based on metadata could thus look like this:

	Judge	Court	Party A	Party B
Case 1	Honorable Smith	New York	Trevor Wayne	Elouise Carpenter
Case 2	Honorable Johnson	Washington	Painting Inc	Lewis Lane

Using such a representation, it could be possible to predict the outcome of a case based on the metadata of a case. For example, a model can be given the name of a judge and the parties and give an estimation of who might win in a certain case.⁶³²

4.4.3.3.2 Prior work

Katz *et al* built a model to predict the decisions at the Supreme Court of the United States. They represented each case as a vector of features taken from the publicly available Supreme Court Database (SCDB).⁶³³ Each case is thus represented as a list of attributes, including the identity of the judge, the term of the court, the month of the argument, the name of the petitioner and respondent, the court of origin and many others.⁶³⁴ The researchers then applied a random forest classifier (a type of machine learning algorithm) to predict the outcome of a case, i.e. whether the decision from the lower court is

⁶³² Ashley, *supra* note 44 at 107.

⁶³³ Daniel Martin Katz, Michael J Bommarito II & Josh Blackman, “A general approach for predicting the behavior of the Supreme Court of the United States” (2017) 12:4 PLOS ONE e0174698 at 4.

⁶³⁴ *Ibid* at 5.

affirmed or reversed.⁶³⁵ Over the entire database of 246k cases between 1816 and 2015, the model correctly predicted 70.2% of the decisions.⁶³⁶

Surdeanu *et al* built a model to predict the outcome of Intellectual Property lawsuits. For their database of 4,200 lawsuits, they annotated each case with prior factors, such as historical and concurrent behavior of the entities involved in the case.⁶³⁷ They modeled the past win rates and participation counts for parties, attorneys and law firms, and modeled potential biases of the judge and the district where the case was filed.⁶³⁸ Based on these factors, they built a model that can predict the outcome of a case with a accuracy of 64%.⁶³⁹ They argued that such systems can be used to estimate the litigation risk and make informed decisions on whether to settle a case.⁶⁴⁰

4.4.3.3 Conclusion

Encoding cases in terms of their metadata can provide a strong prediction model. They can also be useful to potentially inform individuals on certain choices, such as where to file their case or which law firm to retain to argue a certain case. However, they may be less useful in supporting individuals in understanding their legal situation. The merits of the case, such as what has occurred, are not considered by this kind of model.

Likewise, many of the factors used in these kinds of analyses are not available for cases that have not yet entered the court system. For example, it may not be possible to predict which judge will preside over a case that has not yet been filed, or which law firm the opposing side would retain. If the purpose is to provide information before a case has entered the court system, this kind of model could therefore be less useful.

⁶³⁵ *Ibid* at 6–7.

⁶³⁶ *Ibid* at 8.

⁶³⁷ Mihai Surdeanu et al, “Risk analysis for intellectual property litigation” (2011) Proceedings of the 13th International Conference on Artificial Intelligence and Law (ICAIL ’11) 116–120 at 116.

⁶³⁸ *Ibid* at 118.

⁶³⁹ *Ibid* at 119.

⁶⁴⁰ *Ibid* at 120.

4.4.3.4 Merit-based representations

4.4.3.4.1 Introduction

Another approach to encoding case law is to encode the merits of a decision. This kind of representation aims to capture the underlying factual situation of a legal case.⁶⁴¹ Here, the legal situation is encoded into a vector format, that can then be used to analyze and compare cases. This representation can potentially be used in a *neutral* way, by creating an encoding that can be compared between cases that have been decided and cases that have not yet entered the court system, as the factual situation for a new situation may already be clear from what has occurred.

A very simple representation of a case based on facts could look like this:

	Is the dwelling infested with bedbugs?	Is the dwelling infested with mold?	Does the dwelling have issues with heating?
Case 1	No	No	No
Case 2	Yes	No	Yes
Case 3	Yes	Yes	Yes

Each row thus contains a different case, while each column contains a “feature”, in this case a question about the factual occurrences of a previous case. Once this kind of representation has been created for a lot of cases, it can be used to compare the features of a new case to previous cases to infer the outcome or show similar cases.

Deciding exactly what should be captured by such a representation is a crucial factor in building useful systems. In creating a representation, the designer is trying to capture the parts of the decision that are the most relevant for their purposes. In essence, the representation has to be able to capture all of the merits of a decision that are likely to be relevant to how the case is decided, in order to allow the comparison to other cases. This

⁶⁴¹ Ashley, *supra* note 44 at 102.

is not an easy task by any measure and relies on a number of important steps, as discussed above in 2.6.1.2.2. Let us briefly discuss how these steps could be performed in the context of a legal decision support tool.

4.4.3.4.1.1 Step 1 – Deciding on the type of the representation

The first step in the creation of merit-based case-representations is the decision of which type of representation should be used. Ideally, this representation should be suited for the purpose of the decision support tool. The representation should therefore be able to capture the legally important elements of the cases. As previously discussed, if the aim of the system is to support individuals that have not yet gone to court, it is important to make sure that the representation captures cases in a way that can be compared to undecided cases. An example type of representation could be the presence or absence of certain factual occurrences.

4.4.3.4.1.2 Step 2 – Creating a taxonomy of possible elements in this representation

Once the type of the representation has been determined, the designer of a legal decision support tool has to create a taxonomy that can be applied to concrete cases. For example, if the representation relies on the capturing of factors from case law, the designer of the system has to create a list of possible factors that they want to capture from the individual cases. This step should not be underestimated - depending on the domain, the list of possible factors could be quite extensive, requiring a lot of work in elaborating a taxonomy for the factors. In some cases, the court or the legislation may provide a list of important factors that should be considered in the decision of cases. In this case, this list can be helpful for creating the taxonomy.

4.4.3.4.1.3 Step 3 – Encoding cases into the specified format

Finally, once the taxonomy has been created, it has to be applied to case law, in order to capture the cases according to the established taxonomy. This step can also be surprisingly difficult, especially if multiple annotators work to encode the decisions in parallel. If the different annotators have different understandings of how the cases should be labeled, the resulting data will be inconsistent, making it difficult to detect patterns.

4.4.3.4.2 Prior work

4.4.3.4.2.1 *TAXMAN II*

An early system that aimed to assess the link between facts and legal concepts is *TAXMAN II*, developed by McCarty. The initial implementation of the system focused on whether a stock split should be seen as taxable income or not, focusing on a case called *Eisner v. Macomber*.⁶⁴² McCarty argues that open-textured concepts cannot be represented by logic alone.⁶⁴³ Instead, the researchers develop a framework that is able to capture example cases and hypothetical situations (here known as prototypes) and compare these prototypes to a new case based on the concepts involved (through transformations McCarty refers to as deformations).⁶⁴⁴ The system is thus able to generate logically expressed arguments for either side of an issue by comparison to the stored prototypes.⁶⁴⁵

While promising, the system has been criticized for being difficult to apply in the real world. Other researchers have pointed out that it was only implemented in the context of a single argument at the supreme court,⁶⁴⁶ and that it further seems to rely on reality neatly fitting into certain boxes, which may not always be the case.⁶⁴⁷

4.4.3.4.2.2 *HYPO, CATO, IBP*

HYPO

Perhaps the most influential way of representing and comparing case law was created by Kevin Ashley and is called *HYPO*. *HYPO* is able to generate arguments in the area of

⁶⁴² *Ibid* at 78.

⁶⁴³ McCarty, *supra* note 47 at 1.

⁶⁴⁴ McCarty, *supra* note 47; L T McCarty & N S Sridharan, “The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations” (1981) 1 Proceedings of the 7th international joint conference on Artificial intelligence-Volume 246–253; Ashley, *supra* note 44 at 78; James David Popple et al, *SHYSTER: a pragmatic legal expert system* (Canberra: Dept. of Computer Science, Australian National University, 1993) at 31.

⁶⁴⁵ Ashley, *supra* note 44 at 79.

⁶⁴⁶ *Ibid* at 81.

⁶⁴⁷ Popple et al, *supra* note 644 at 31.

trade-secret misappropriation. It can generate arguments saying that a case represents a trade secret misappropriation, or that it does not.⁶⁴⁸

HYPO represents case law in terms of the factors present in a case. A factor is a "commonly observed collections of facts that tend to strengthen or weaken a plaintiff's argument in favor of a legal claim".⁶⁴⁹ For the domain of trade secret misappropriation, the researchers identified 13 different relevant factors, some stemming from the legislation, while others come from case law, treatises and doctrine.⁶⁵⁰ These factors were encoded in terms of a dimension, which meant that they could also have magnitudes.⁶⁵¹ A factor in trade secret misappropriation, for example, is how many individuals a secret was disclosed to. This could range from a few to thousands of individuals, which is captured in the magnitude of the factor.⁶⁵² Overall, 30 trade secret misappropriation cases were coded using dimensions.⁶⁵³

In order to generate arguments for a new case, a hypothetical situation was entered into the system by the user, referred to as a current fact situation. Based on this situation, the system generates legal arguments for both the plaintiff and defendant side. This analysis is done by finding cases that favor either party, based on an overlap in the dimensions shared by those cases and the hypothetical situation. Further, the system distinguishes cases that do not agree with a certain position, based on legally relevant differences between the current fact situation and previous cases. HYPO can also identify counterexamples, that could be cited by an opposing party in response to an argument.⁶⁵⁴

⁶⁴⁸ Kevin D Ashley, *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*, PhD Thesis (Amherst, MA, USA: University of Massachusetts, 1988); published in Kevin D Ashley, *Modeling Legal Argument: Reasoning with Cases and Hypotheticals* (Cambridge, MA.: The MIT Press / Bradford Books, 1990).

⁶⁴⁹ Ashley, *supra* note 44 at 81.

⁶⁵⁰ *Ibid.*

⁶⁵¹ Rissland & Ashley, "HYPO", *supra* note 47 at 232; Ashley, *supra* note 44 at 81; T J M Bench-Capon, "HYPO'S legacy: introduction to the virtual special issue" (2017) 25:2 *Artif Intell Law* 205–250 at 208.

⁶⁵² Ashley, *supra* note 44 at 82.

⁶⁵³ Kevin D Ashley, "Reasoning with cases and hypotheticals in HYPO" (1991) 34:6 *International Journal of Man-Machine Studies* 753–796 at 763; Ashley, *supra* note 44 at 81.

⁶⁵⁴ Ashley, *supra* note 44 at 82–87; Rissland & Ashley, "HYPO", *supra* note 47 at 254–300.

HYPO was one of the most important foundational papers in the field of AI & Law,⁶⁵⁵ leading to a number of influential successors.

CATO

One of these successors is CATO, developed by Vincent Aleven. It also focuses on trade secret misappropriation cases but eschews the user of dimensions in favor of the simpler representation of factors that capture whether any single factor in a case is present or absent. Cases were represented in terms of 26 distinct factors.⁶⁵⁶

Further, the factors in CATO were arranged in a hierarchy that links the factors to legal issues.⁶⁵⁷ The hierarchy thus allows cases that have different factors that are nonetheless linked to the same legal issue to be compared.⁶⁵⁸ Using this factor hierarchy, the system can compare the previous cases to a new case, and generate arguments organized by a certain legal issue.⁶⁵⁹ The main purpose of CATO was education, i.e. teaching law students the basics of making arguments with legal cases.⁶⁶⁰

CATO could also be used for legal prediction. Using the same methodology as presented above, it could retrieve cases that are similar to the current fact problem. If all of the cases were decided in a certain direction, CATO would predict the case as likely to be decided in the same direction, otherwise it would abstain. The result of this prediction was CATO abstaining in 11% of cases and predicting 88% of the rest of the cases correctly.⁶⁶¹

IBP

⁶⁵⁵ Bench-Capon, “HYPO’S legacy”, *supra* note 651 at 206.

⁶⁵⁶ Aleven, “Using background knowledge in case-based legal reasoning”, *supra* note 47 at 189–191.

⁶⁵⁷ *Ibid* at 191–193.

⁶⁵⁸ Ashley, *supra* note 44 at 90–92.

⁶⁵⁹ Aleven, “Using background knowledge in case-based legal reasoning”, *supra* note 47 at 195–200.

⁶⁶⁰ *Ibid* at 218–219.

⁶⁶¹ Ashley, *supra* note 44 at 115; Aleven, “Using background knowledge in case-based legal reasoning”, *supra* note 47 at 213–216.

A further expansion of the HYPO model is called IBP, issue-based prediction, developed by Ashley and Brüninghaus. IBP introduces a domain model for trade-secret misappropriation, that links factors to different issues. The links have been established from the analysis of authoritative sources on the law of trade secret misappropriation.⁶⁶²

IBP works by first making a decision on each issue (such as Information-Valuable) by looking at the individual factors related to that issue, and whether they all favor either party. In case they are ambiguous, the system retrieves previous cases to decide which side will likely prevail in a certain issue.⁶⁶³ Once the issues are decided, IBP uses the logical connections between the issues (such as “and” and “or”) to predict the outcome of the entire case (i.e. whether trade secret misappropriation was present or not).⁶⁶⁴ Overall, this system was able to predict the outcome of trade secret misappropriation cases with an accuracy of 91.8%.⁶⁶⁵

VJAP

Another expansion of this line of research comes from Matthias Grabmair, in the form of Value Judgment-based Argumentative Prediction (VJAP). His work focuses on introducing values into argument schemes. In reaching legal decisions, judges must often weigh the impact of each potential decision on two competing values that underlie the legal norms.⁶⁶⁶ For example, in the trade secret domain (in which VJAP was implemented), the different legal factors can be connected to values such as the plaintiff’s interest in property and confidentiality, versus the public interest of usability of public information and fair competition.⁶⁶⁷ The VJAP system is able to generate arguments that include these values, by identifying and reasoning with cases that have a similar trade-

⁶⁶² Kevin D Ashley & Stefanie Brüninghaus, “Automatically classifying case texts and predicting outcomes” (2009) 17:2 *Artif Intell Law* 125–165 at 134.

⁶⁶³ *Ibid* at 137–138.

⁶⁶⁴ *Ibid* at 139.

⁶⁶⁵ *Ibid* at 150.

⁶⁶⁶ Grabmair, *supra* note 52 at 7–8; Ashley, *supra* note 44 at 149–150.

⁶⁶⁷ Grabmair, *supra* note 52 at 38–41; Ashley, *supra* note 44 at 150–154.

offs between the values linked to different issues, or cases that have trade-offs between the factors linked to the same issue.⁶⁶⁸

VJAP is also able to generate an argument graph structure, that contains all possible arguments using the different argument schemes supported by VJAP.⁶⁶⁹ This argument graph structure was used to predict the outcomes of cases, achieving an accuracy of around 80%.⁶⁷⁰

4.4.3.4.2.3 *GREBE*

Another type of encoding is that of semantic networks. Branting developed a system called GREBE. It relies on encoding cases, statutes and common-sense rules as semantic networks, i.e. a graph-based encoding of the criterial facts that the judge found important for their decision. These are linked in terms of the concept, and the relationship between the concepts.⁶⁷¹ GREBE was able to match concepts from a new case to previous cases, thereby determining whether the previous case was relevant or not, and using the relevant cases to generate arguments. These arguments held up favorably when compared to arguments written by law students.⁶⁷²

4.4.3.4.2.4 *Nearest-neighborhood approaches*

The previous approaches to reason using cases rely on symbolic reasoning systems, that use the factors in a case to predict outcomes or generate arguments about whether legal concepts apply or not. There have also been experiments using the nearest neighborhood approach. Here, cases are encoded in terms of a number of factors. The cases are then compared based on the overall number of overlapping factors. The outcome for the current case is assumed to be the same as the outcome of the most similar case, or a number of the most similar cases (see 2.6.1.3.1).

⁶⁶⁸ Grabmair, *supra* note 52 at 42–47; Ashley, *supra* note 44 at 154–155.

⁶⁶⁹ Grabmair, *supra* note 52 at 46–68; Ashley, *supra* note 44 at 156–158.

⁶⁷⁰ Grabmair, *supra* note 52 at 77; Ashley, *supra* note 44 at 158–160.

⁶⁷¹ L Karl Branting, “Building explanations from rules and structured cases” (1991) 34:6 *International Journal of Man-Machine Studies (AI and Legal Reasoning. Part 1)* 797–837 at 808–809; Ashley, *supra* note 44 at 93–94.

⁶⁷² Branting, *supra* note 671 at 816–827; Ashley, *supra* note 44 at 93–96.

Mackaay and Robillard used the nearest neighbor approach to predict the outcome of a classification of tax cases, namely whether a gain should be considered as capital gains or ordinary income.⁶⁷³ This area does not have firm criteria, but judges instead rely on a number of factors. The researchers use a database of 64 cases, each of which have been annotated in terms of 46 binary descriptors, covering attributes such as “the private party is a company”, or “private party has not subdivided the land”.⁶⁷⁴ For each new case, the similarity to the encoded cases is calculated, and the new case is assumed to have the same outcome as the most similar case. Similarity is here referred to as the cases that share the highest number of attributes, using a measure called Hamming distance.⁶⁷⁵

A system that uses a similar approach is SHYSTER. It aims to replicate the reasoning of lawyers on a pragmatic level, to be able to generate useful predictions and advice.⁶⁷⁶

SHYSTER focused on the use by lawyers, and thus assumes a certain level of legal expertise.⁶⁷⁷ The system represented leading cases in an area (i.e. relating to a single legal concept) by a number of attributes (akin to a factor described above) that can hold values of YES, NO or UNKNOWN.⁶⁷⁸ The user is asked to indicate whether the attributes apply in their case, upon which point the system can surface the most similar case in terms of shared attributes.⁶⁷⁹ The system can generate predictions for the outcome of a new case, a report that explains the prediction, and try hypotheticals to determine whether a change in an attribute would have an effect on the outcome.⁶⁸⁰

SHYSTER was implemented in the domains of trover, the meaning of the term “authorization” in copyright infringement, the categorization of an individual as an

⁶⁷³ Ejan Mackaay & Pierre Robillard, *Predicting judicial decisions: The nearest neighbour rule and visual representation of case patterns* (De Gruyter, 1974).

⁶⁷⁴ *Ibid* at 311, 327–331.

⁶⁷⁵ *Ibid* at 307; Ashley, *supra* note 44 at 108.

⁶⁷⁶ Popple et al, *supra* note 644 at vii.

⁶⁷⁷ *Ibid* at 52–53.

⁶⁷⁸ *Ibid* at 59–61.

⁶⁷⁹ *Ibid* at 70.

⁶⁸⁰ *Ibid* at 70–85.

employee versus a contractor, and the notion of natural justice, showing the versatility of the program.⁶⁸¹

4.4.3.4.2.5 *Machine-learning based approaches*

Another way to build models that can predict the applicability of legal concepts is machine learning. Here, cases are encoded in terms of factors, and then fed as training data to a machine learning algorithm (see 2.6). This model can ideally be used to predict the outcome of new cases fed to the system.

Alarie *et al* built a system to predict whether an individual would be seen as a worker or an independent contractor. They annotated 600 cases using 21 factors, such as whether a worker and the hirer had a contract, who set the hours of the work etc.⁶⁸² They then used a machine learning algorithm to build a model able to predict the outcome of new cases. They claimed to achieve an accuracy of over 90%.⁶⁸³ The system also gave an explanation and provided the user with similar cases.⁶⁸⁴

Similarly, Yin *et al* annotated a dataset of 900 decisions on whether an individual should be seen as a worker or a contractor using 16 individual factors. They were then able to predict the outcome of these cases with an accuracy of 91.5%.⁶⁸⁵

Branting *et al* extracted 46 factors from cases regarding domain name dispute cases. They used a projection method to only annotate 25 cases, out of 16,024 cases (see 4.4.3.4.2.6 for more details). Using the 46 factors, they then predicted the outcome of the cases, achieving an F1-score of 0.795.⁶⁸⁶

⁶⁸¹ *Ibid* at 132–193.

⁶⁸² Benjamin Alarie, Anthony Niblett & Albert H Yoon, “Using machine learning to predict outcomes in tax law” (2016) 58 *Can Bus LJ* 231 at 10.

⁶⁸³ *Ibid* at 11.

⁶⁸⁴ *Ibid* at 13–15.

⁶⁸⁵ Yifei Yin, Farhana Zulkernine & Samuel Dahan, “Determining Worker Type from Legal Text Data using Machine Learning” (2020) 2020 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCoM/CyberSciTech) 444–450.

⁶⁸⁶ Branting *et al*, *supra* note 54.

4.4.3.4.2.6 *Automatically extracting factors from decisions*

Annotating cases can be a significant time sink and make it difficult to obtain enough data to construct viable case-based systems. Therefore, there has been some work on automatically extracting factors from case texts. Note that this is different from the textual representations of cases presented in 4.4.3.2. In the textual representations, the text itself is used to predict the outcome or identify similar cases. As we saw, this leads to some issues. In this section, I will describe approaches where machine learning is used to extract factors from previous decisions, as an intermediary step. These factors are then used to predict the applicability of a legal factor, or surface similar cases. Factors, of course, can in theory be used to compare undecided cases to previous cases.

Ashley and Brüninghaus trained a machine learning algorithm to extract the factors used in IBP from case squibs, i.e. summaries of case texts. This could allow the automatic use of systems like IBP, without annotation. However, the results of the text classification system were not strong enough to allow this.⁶⁸⁷ Falakmasir and Ashley used Vector Space Models to automatically extract factors from entire case texts, achieving F1-scores of 0.65.⁶⁸⁸

Another attempt at automatically extracting factors from cases is SCALE. The researchers annotated 25 cases using 46 annotation tags. These were then projected to a corpus of 16,024 cases using word embeddings, which capture the semantic meaning of a word. The researchers used the intermediary representation of the 46 factors to train a model to predict the outcome of the cases. While the overall accuracy of the prediction was lower than a system purely using the text, the researchers believed that these features could be more helpful to support the explanation of the cases.⁶⁸⁹

⁶⁸⁷ Ashley & Brüninghaus, *supra* note 662 at 161.

⁶⁸⁸ Falakmasir & Ashley, *supra* note 54.

⁶⁸⁹ Branting et al, *supra* note 54 at 221–230.

Jack Mumford *et al*⁶⁹⁰ used hierarchical BERT models (see 2.6.2.3) to ascribe the factors present in a case of the European Court of Human Rights. They then fed the predictions of which factors were present into an ADF model (see 4.4.3.4.2.7), in order to predict the cases in an explainable manner. The authors achieved an accuracy of 72%, compared to 66.78% when only using a machine learning model.⁶⁹¹

There has also been research investigating whether humans can collaborate with computers to extract factors from cases. We built an interface that allows people to investigate texts (such as legal decisions) and mark whether certain words are likely to be indicative of a factor being present. We tried this approach on a number of datasets, including the trade secrets dataset used in CATO and IBP. While the human created rules were not as performant as machine learning classifiers in detecting factors, they were explainable and allowed more control over choosing which words should be part of the prediction.⁶⁹²

Similarly, we built a system that can support the user in conducting annotations more efficiently, by supporting them using machine learning. The system, called CAESAR (Computer Assisted Semantic Annotation & Ranking) allows the annotation of sentences in cases. However, it also allows the user to automatically retrieve sentences from across documents that are semantically similar to a sentence they just annotated. Since similar sentences often refer to similar concepts, such sentences may also contain the same factor. The annotator can thus annotate many factors at once.⁶⁹³

4.4.3.4.2.7 *Rule-based representation of open texture?*

Until now, we have looked at ways to model open textured terms in law that rely upon examples from previous case law, by comparing a new case to previous cases in order to

⁶⁹⁰ Jack Mumford, Katie Atkinson & Trevor Bench-Capon, “Reasoning with Legal Cases: A Hybrid ADF-ML Approach” (2022) 362 *Legal Knowledge and Information Systems (Frontiers in Artificial Intelligence and Applications)* 93–102.

⁶⁹¹ *Ibid* at 100.

⁶⁹² Westermann *et al*, *supra* note 194.

⁶⁹³ Westermann *et al*, *supra* note 195.

assess if the new case fulfills a legal concept or not. There has also been work arguing for the rule-based reasoning around open textured concepts. Bench-Capon and Sergot argue that the grey area of open-textured legal concepts (i.e. where it is not clear whether a concept applies or not) should be covered by multiple, contradictory rules extracted from case law.⁶⁹⁴

These contradictory rules can then be used to produce arguments for either side of an issue, for use by lawyers aiming to argue the case in court. Documentation as to the creation of these rules would be included to help the user evaluate and select the arguments presented.⁶⁹⁵

The ANGELIC (ADF for kNowledGe Encapsulation of Legal Information for Cases) framework is an implemented model in this vein. Here, the authors use Abstract Dialectical Frameworks (ADF) to assess the applicability of open-textured concepts. ADFs consists of a number of “nodes”, each of which is linked through manually created tests to child nodes, representing other concepts.⁶⁹⁶ These ADFs are used to encode rules from case-law, that can then be used to predict the outcome of new cases and generate explanations.⁶⁹⁷

The authors used ADFs to model the factor hierarchy from CATO and IBP, achieving similar results in prediction.⁶⁹⁸ The modeling also succeeded in other areas, such as the hunt of wild animals and vehicle searches.⁶⁹⁹ The researchers have further implemented the methodology in the domain of noise-induced hearing loss.⁷⁰⁰ They also developed a way to integrate dimensions (as introduced in HYPO) into the ANGELIC

⁶⁹⁴ Bench-Capon & Sergot, *supra* note 601 at 50–54.

⁶⁹⁵ *Ibid* at 58.

⁶⁹⁶ Latifa Al-Abdulkarim, Katie Atkinson & Trevor Bench-Capon, “A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks” (2016) 24:1 *Artif Intell Law* 1–49 at 12.

⁶⁹⁷ *Ibid* at 39–41.

⁶⁹⁸ *Ibid* at 13–24.

⁶⁹⁹ *Ibid* at 27–38.

⁷⁰⁰ Latifa Al-Abdulkarim et al, “Noise induced hearing loss: Building an application using the ANGELIC methodology” (2019) 10:1 *Argument & Computation* 5–22.

methodology,⁷⁰¹ and a graphical user interface for developing systems using ANGELIC.⁷⁰²

4.4.3.4.3 Conclusion

As we have seen, there are a number of systems that aim to automate the determination of whether certain legal criteria apply to new cases or not. There have been several ways to create representation of cases that allow the prediction and reasoning around new cases. Some of these representations are relatively simple, being based on a list of factors that are present in a case (e.g. the one used in SHYSTER, by Robillard *et al*, Yin *et al* and Alarie *et al*). We have also seen factors that are represented by dimensions (HYPO), adding sophisticated structures linking the factors together based on legal issues (CATO and IBP), linking factors to different values (VJAP), cases represented as prototypes that can be compared via transformations (TAXMAN II) and cases represented as semantic networks (GREBE). Finally, we have seen cases being represented in terms of rules inferred from the decisions (ANGELIC).

These representations can be compared and analyzed in a number of ways, including through symbolic legal reasoning systems (HYPO, CATO, IBP, TAXMAN II, GREBE), nearest neighborhood approaches (SHYSTER, Robillard *et al*), machine learning approaches (Alarie *et al*, Yin *et al*, SCALE) and rule-based approaches (ANGELIC). These systems can allow the prediction of whether a concept applies or the generation of arguments for either side of an issue.

Let us examine some aspects of these systems and their appropriateness for our use case of increasing access to justice.

⁷⁰¹ Latifa Al-Abdulkarim, Katie Atkinson & Trevor Bench-Capon, “ANGELIC Secrets: Bridging from Factors to Facts in US Trade Secrets” (2016) *Legal Knowledge and Information Systems* 113–118.

⁷⁰² Latifa Al-Abdulkarim et al, “Angelic environment: demonstration” (2017) *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law (ICAIL '17)* 267–268; Latifa Al-Abdulkarim et al, “Angelic Environment: Support for the Construction of Legal KBS” (2015) *Proceedings of the 15th International Conference on Artificial Intelligence and Law* 3–12.

4.4.3.4.3.1 *Appropriate for laypeople*

Most of the aforementioned systems are built to be used by lawyers. Lawyers are able to evaluate legal arguments, and encode their case in terms of factors requiring legal understanding. Laypeople, on the other hand, may think of their situation purely in terms of facts, which may make it difficult to interact with such approaches.

By nature of operating on the level of classifying facts into legal criteria, the systems described may also not contain any specific evidence considerations. Such systems are suited to provide the answer to the question "Would a judge find that a legal concept applies, *if certain facts can be proven?*", rather than "Would a judge find that a legal concept applies?". Lawyers can understand this difference, but it may not be obvious to laypeople.

4.4.3.4.3.2 *Identifying factors and dealing with new factors*

Another difficulty in building case-based systems is the identification of possible factors in a case. In some areas, cases and doctrine specifically point to a number of factors that should be considered by the court in determining whether a legal concept applies. The problem space can be seen as bounded.⁷⁰³

In other areas, there might not be such a limited number of factors. For example, there could be an enormous number of issues that affect the determination of whether a landlord has provided an apartment in "good habitable condition", and the judge can take any factors into account. The problem space is unbounded, as there is no list of factors that should be considered by the court.⁷⁰⁴ Identifying factors in such domains can only be done through the reading of many such cases, and the creation of a taxonomy. I describe the difficulty of doing so in Chapter 6.

In developing such a taxonomy, the creator makes a choice in which factors are relevant and important. This choice is likely a pared down version of reality since events in the

⁷⁰³ Compare John Zeleznikow, "Building Decision Support Systems in Discretionary Legal Domains" (2000) 3 Int'l Rev L Computers & Tech 341–356 at 343–344.

⁷⁰⁴ Compare *ibid.*

world rarely fit into neat boxes. It can be seen as a “lossy compression”, i.e. some information is lost.⁷⁰⁵ The choice made in which attributes are important can additionally encode the moral values of the creator into the system.⁷⁰⁶

Likewise, new factors may arise that are not captured by the system, making the analysis fail in certain cases. As a hypothetical, let us imagine that we have encoded 1000 cases regarding the requirement of “peaceable enjoyment”, and identified 50 specific factors that can influence this criterion. A new case arises, where the apartment is located next to a racetrack, a situation that has not previously arisen and is not captured by any of the factors. A judge, of course, could still make an assessment of whether this new factor influences the peaceable enjoyment of the apartment. The AI system, on the other hand, is limited by the factors that it has access to. Since none of the factors apply, the AI system would likely say with a high certainty that there is no issue with peaceable enjoyment. As described above in 2.6.3.5, the current AI systems likely do not have the common sense to be able to decide on how new factors interact with a given legal criterion. It is therefore possible that the task of classifying facts in terms of concepts is an AI-complete tasks (2.4.2.4).

4.4.3.4.3.3 *Difficulty of coding*

Beyond the difficulty of identifying factors, another issue is the practical difficulty of encoding case law in terms of factors. Reading and understanding case law is a complex endeavor, requiring legal skills. Further, cases can be quite long, making it very time-consuming to read through and encode them into a certain taxonomy. This can make it difficult to encode the large number of cases required for certain analysis techniques, such as machine learning, where hundreds or thousands of cases may be required. Unlike some machine learning domains, where there are readily available datasets for many tasks, the representations and taxonomies for encoding case law are likely to be bespoke

⁷⁰⁵ Westermann et al, *supra* note 536 at 141.

⁷⁰⁶ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 804.

and tailored to solve a specific task. Creating these datasets can be quite expensive and a significant bottleneck for research.⁷⁰⁷

Certain systems rely on much fewer cases to create strong arguments or predictions. HYPO, for example, relied on 30 cases to generate arguments in the domain of trade secrets. Authority for this reasoning does not come from the quantity of the cases, but rather from the quality of the generated argument.⁷⁰⁸ In the approaches relying on machine learning, many more cases have to be encoded to build an accurate model of an area.

Beyond the time and resources required, it is important that the data captured is accurate and reliable. Neither are easy to achieve. In looking at the facts of a decision, for example, it can be difficult to determine whether a certain fact was merely alleged by a party, or if it was seen as proven by a judge and therefore lay at the basis of the decision.

These problems are exacerbated if multiple annotators are involved. It is important that each annotator has the same conception of what the elements to annotate refer to. If different people apply different elements in different ways, the resulting dataset will be inconsistent, which severely hampers the usefulness of the data. Having extremely clear guidelines and annotating certain parts multiple times,⁷⁰⁹ to investigate the reliability of the annotations, is therefore important in assuring the quality of the annotations.

As we have seen, using machine learning to automatically identify the factors from a case text could support the annotation of cases. Here, a limited number of cases could be

⁷⁰⁷ Compare 2.6.1.2.

⁷⁰⁸ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 18.

⁷⁰⁹ Compare Karl L Branting et al, “Semi-Supervised Methods for Explainable Legal Prediction” (2019) Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law (ICAIL ’19) 22–31; Lisa Ferro et al, “Scalable Methods for Annotating Legal-Decision Corpora” (2019) Proceedings of the Natural Legal Language Processing Workshop 2019 12–20; Bernhard Walzl et al, “Classifying Legal Norms with Active Machine Learning.” (2017) Proceedings of the International Conference on Legal Knowledge and Information Systems (JURIX) 11–20; Ashley & Brüninghaus, *supra* note 662; Jaromir Savelka et al, “Lex Rosetta: transfer of predictive models across languages, jurisdictions, and legal domains” (2021) Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL ’21) 129–138 at 132.

annotated in order to build a model, that could then extract factors from larger corpora of cases.⁷¹⁰ The projects presented in 4.4.3.4.2.6 demonstrate how promising this approach is. Computers collaborating with humans in annotating cases more rapidly also has a lot of potential.⁷¹¹

4.4.3.4.3.4 *Frequency of cases*

Another important consideration is the overall frequency of cases. Machine learning models require a significant number of cases to build an accurate model. The more factors there are relating to each case, the more cases are required for the algorithm to properly learn the correlations. In computer science, this is known as the curse of dimensionality.⁷¹²

For some legal areas, such as determination of whether an individual is an employee or a contractor, this may not be a huge issue, since it is a specific test that is employed often, with similar factors in each case. However, other tests may be applied much more rarely, or cover a more diverse set of factors. Here, obtaining enough cases to build an accurate model of an area may be much more difficult.⁷¹³

4.4.3.4.3.5 *Adapting to change*

Another issue that could arise is the inability of the models to adapt to change in the jurisprudence. Since the models tend to learn from what previously occurred when predicting future decisions, they see new decisions purely in the light of past decisions. Even if there are enough decisions, and all the factors have successfully been extracted from a decision, the model may thus fail to anticipate changes in how a legal concept is applied.

Such changes may stem from new legislation that has been introduced in an area, and affect how a legal concept is interpreted. Further, as we have seen in 4.4.1, judges can

⁷¹⁰ Ashley, *supra* note 44 at 74.

⁷¹¹ Westermann et al, *supra* note 195.

⁷¹² Domingos, *supra* note 109 at 82.

⁷¹³ Westermann et al, *supra* note 536 at 142.

incorporate policy decisions and societal needs in deciding how legal concepts should be interpreted. Changes in society can thus lead to changes in how legal concepts are applied. Emulating this kind of reasoning requires an in-depth understanding of the world, which may not be something that current AI systems are able to do.

4.4.3.5 Conclusion

This section has given an overview of different ways to represent case law, in order to be able to predict the applicability of legal concepts in certain cases. Textual and metadata representations, while interesting, seem limited in their ability to surface relevant information before a case has gone to court, which is crucial in building systems to help laypeople. Merit-based representations can build models of how certain facts relate to legal criteria being fulfilled. At the same time, it is important to be aware of the difficulty of encoding many cases and dealing with unique situations or changes to jurisprudence.

4.5 Applying legal rules to legal criteria

4.5.1 The legal system

Once the judge has established the legal criteria that are fulfilled by the undisputed or proved facts of a case, they have to apply the legal rules to the criteria, to arrive at an outcome for the case. This is a deductive, logical style of reasoning, following the structure of a syllogism:⁷¹⁴

Major premise: If the landlord fails to perform one of their duties, the tenant can apply for a reduction of rent.

Minor premise: The landlord fails to perform one of their duties.

Conclusion: The tenant can apply for a reduction of rent.

⁷¹⁴ Holland & Webb, *supra* note 552 at 141.

More generally, this form of logical deduction can be written as the following:⁷¹⁵

Major premise: $p \rightarrow q$

Minor premise: p

Conclusion: q

The major premise here refers to a specific rule that is relevant to a case. The minor premise refers to the specific legal criteria that were fulfilled in a certain case. The conclusion refers to the outcome that is established in the law.⁷¹⁶ The applicability of the legal rule is conditional on the specific legal criterion being fulfilled.

In some cases, applying the rules in this logical manner is sufficient to establish the legal outcome of a case. In other cases, multiple legal criteria might have to be fulfilled. These are *cumulative* criteria. For example, as can be seen in article 1863 of the CCQ (see 4.2.2), in order to terminate a lease, there would have to be a nonperformance of an obligation by the landlord, in addition to the criteria of this causing the tenant a serious injury. Only if both of these are fulfilled can the legal consequence be ordained.

In other cases, there may be *alternative* criteria. Here, one of several legal criteria have to be fulfilled for the legal rule to be applicable.

By traversing these rules, the judge can arrive at an outcome. They may have to use different articles to arrive at a single outcome. Some articles may introduce an intermediary conclusion (such as “the apartment is in an uninhabitable state”) which is then linked to a consequence in another article (such as “the lease can be terminated if the apartment is in an uninhabitable state”).

Of course, determining the legal outcome is not always as simple as merely applying logical rules. In some instances, the logical content of a rule may not be clear. Perhaps the wording of a rule is ambiguous and allows two possible interpretations, each leading

⁷¹⁵ MacCormick, *supra* note 578 at 24.

⁷¹⁶ Compare *ibid* at 26.

to a different syllogism.⁷¹⁷ Ashley refers to this type of ambiguousness as syntactic ambiguity, and argues that they may arise from particularities of natural language, such as not being able to specify the scope of logical connectors such as “and” and “or”.⁷¹⁸ If this kind of ambiguity is present, the judge has to decide on one of the interpretations, by analyzing what is meant by the legislator or find other ways to determine a reasonable interpretation. Allen and Saxon identified a simple provision that yielded 48 possible interpretations.⁷¹⁹

Another possible situation is that multiple rules may collide.⁷²⁰ Perhaps, the specific interpretation of a rule collides with a human right.⁷²¹ Likewise, a specific law may collide with a more general law. In this case, the judge has to determine which of the rules should gain precedence, and how this impacts the outcome of the case.

4.5.2 Example

Let us investigate how this step may be applied to our example case. Article 1854 tells us that the landlord has a duty to ensure the peaceable enjoyment of property to the tenant. We established in the previous step that the landlord had failed to do so in our case. Article 1854 does not contain any consequences. Instead, we look to paragraph 1863.⁷²² By reading this paragraph, we can gather that the tenant has the right to a rent reduction if the landlord fails to perform one of their obligations. By traversing these criteria and articles of the legislation, we have connected the fulfilled legal criteria to an outcome.

Of course, things are not always this simple. Should the tenant, for example, have demanded a termination of the lease, they would also need to prove that the nonperformance causes serious injury to them.⁷²³ Here, the requirements are cumulative –

⁷¹⁷ Holland & Webb, *supra* note 552 at 131; Ashley, *supra* note 44 at 41.

⁷¹⁸ Ashley, *supra* note 44 at 41.

⁷¹⁹ *Ibid* at 45–46.

⁷²⁰ Holland & Webb, *supra* note 552 at 131.

⁷²¹ Compare Dinah Shelton, “Hierarchy of norms and human rights: Of trumps and winners” (2002) 65 Sask L Rev 301.

⁷²² See 4.2.2.

⁷²³ *CCQ*, *supra* note 557 article 1863.

both the criteria of a non-performance of an obligation, a serious injury having appeared and the two being causally connected need to be fulfilled in order for the termination of the lease to be granted.

4.5.3 Automating the application of legal rules to legal criteria

4.5.3.1 Introduction

Let us look at ways to automate the reasoning performed in applying legal rules to fulfilled legal criteria. As we will see, most of the tools built to tackle this kind of reasoning are in the shape of rule-based systems or expert systems. This is an appropriate approach since most legal rules are already in the shape of the if-then rules that these systems rely on (see 2.5.1 and 4.5.1). It is therefore often easier to encode these rules directly into a computer system, rather than trying to infer the rules from case law. Of course, this also comes with some challenges – fully encoding the legal rules covering an area can be time-consuming, and the researchers still need to deal with issues such as syntactic ambiguity, vague terms and colliding rules.

Let us explore how prior work dealt with encoding rules into a computer-readable format, and how the mentioned issues have been overcome.

4.5.3.2 Prior work

4.5.3.2.1 Waterman & Peterson

Waterman and Peterson developed a system to provide advice on settlement decisions in the domain of product liability. It aimed to model the reasoning performed by litigators when deciding how to settle a case, in order to explore how cases are settled.⁷²⁴ The system worked by defining a number of rules in the system, including concepts such as strict liability. The intermediary conclusion of each rule could be a premise for other rules. The system thus proceeded in a forward-chaining way to iterate through the concepts to identify the rules that applied in a case, and thereby giving a prediction and

⁷²⁴ D A Waterman & M A Peterson, “Models of Legal Decisionmaking - Research Design and Methods” (1981) RAND Corporation at vii–viii.

assessment of the case.⁷²⁵ In order to test whether the base-level concepts (i.e., that were not defined by any rules) applied, the system asked the user.⁷²⁶

4.5.3.2.2 Allen & Engholm

Allen and Engholm devised a way to represent statutes in a digital way, using propositional logic. This approach replaces textual logical operators (such as “and” and “or”) with a digital equivalent. These digital connectors can only be applied in a single way, and are therefore both easier to understand, especially for laypeople, and also eliminate the syntactic ambiguity that may be present in a lot of legislation. The legislation is thus expressed in the form of a flowchart, where each box corresponds to a substantive term. Further, if a computer is told which of the substantive terms apply in a certain case, it can output the expected outcome according to the encoded law.⁷²⁷

4.5.3.2.3 Sergot *et al*

Sergot *et al* implemented the British Nationality Act in Prolog, a programming language that can be used to encode logical rules, and prove logical conclusions.⁷²⁸ In total, the system contained around 150 rules, showing ways of obtaining British nationality. These were encoded by a student in two months.⁷²⁹

Sergot *et al* ran into the issue of open-textured terms. In the British Nationality Act, certain questions such as “being a good character” are open-textured. The researchers used a solution similar to Waterman & Peterson, by asking the user of the system whether the concepts apply or not. They also suggest that rules of thumb, derived from previous cases, may help with assessing the open-textured terms.⁷³⁰

⁷²⁵ *Ibid* at 15–17; Ashley, *supra* note 44 at 8–10.

⁷²⁶ Waterman & Peterson, *supra* note 724 at 26.

⁷²⁷ Layman E Allen & C Rudy Engholm, “Normalized Legal Drafting and the Query Method” (1978) 29:4 *Journal of Legal Education* 380–412; Ashley, *supra* note 44 at 41–44.

⁷²⁸ Sergot *et al*, *supra* note 48 at 372.

⁷²⁹ *Ibid* at 383; Ashley, *supra* note 44 at 47.

⁷³⁰ Sergot *et al*, *supra* note 48 at 371; Ashley, *supra* note 44 at 51.

4.5.3.2.4 The default logic paradigm

Vern Walker suggests modelling legal rules in the form of implication trees.⁷³¹ This is a logical form to model multiple rules, stemming from the default logic paradigm of reasoning.⁷³² In this model, the top of the tree contains the legal conclusion. The lower levels model the constituent propositions of the condition. This constituent tree could model “all of the legally acceptable lines of reasoning that can prove or disprove the ultimate issue stated by the conclusion at the top of the tree.”⁷³³ Walker argues that the implication tree structure can be used to make the steps of the reasoning involved in reaching a certain legal conclusion more transparent and be used to visually present a large amount of information simultaneously. They can also be used by either party to a case to build their argument or defeat the argument of the other side.⁷³⁴ Further, the system can also be used to structure the evidence of a case, by linking plausibility schemas to the terminal propositions of the implication tree, as described in 4.3.3.⁷³⁵

The default logic paradigm was used to annotate cases in the domain of disability claims by veterans at the U.S. Board of Veterans’ Appeals, where sentences were annotated with the role they play in relation to the implication tree.⁷³⁶ There has also been significant work in automatically extracting these roles from sentences.⁷³⁷

⁷³¹ Walker et al, *supra* note 579; Walker, *supra* note 579; Walker, *supra* note 579.

⁷³² Walker, *supra* note 579 at 201.

⁷³³ *Ibid* at 201–202.

⁷³⁴ *Ibid* at 203.

⁷³⁵ *Ibid* at 210–219.

⁷³⁶ Vern R Walker et al, “Semantic Types for Decomposing Evidence Assessment in Decisions on Veterans’ Disability Claims for PTSD” (2017) Proceedings of the 2nd Workshop on Automated Semantic Analysis of Information in Legal Text 10; Vern R Walker et al, “Evidence Types, Credibility Factors, and Patterns or Soft Rules for Weighing Conflicting Evidence: Argument Mining in the Context of Legal Rules Governing Evidence Assessment” (2018) Proceedings of the 5th Workshop on Argument Mining 68–78.

⁷³⁷ Vern R Walker et al, “Automatic Classification of Rhetorical Roles for Sentences: Comparing Rule-Based Scripts with Machine Learning” (2019) Proceedings of the 3rd Workshop on Automated Semantic Analysis of Information in Legal Text.

4.5.3.2.5 PROLEG

Satoh *et al* developed PROLEG, a logical representation of the Japanese Presupposed Ultimate Fact Theory (JUF theory).⁷³⁸ PROLEG was previously described in 4.3.3.2, regarding how evidence is handled. As we saw, it linked a rule-based reasoning system to the burden of proof and assessments of evidence, to build an interactive system that allows the exploration of paths to a certain outcome, ways for the other party to challenge this outcome, and allows the arranging of evidence for disputed issues. A specialty of the PROLEG system is that not all rules are assessed in every case. Some concepts are open and thus only taken into account when raised by either party.⁷³⁹

4.5.3.2.6 Hybrid systems

The previously described systems use exclusively rule-based reasoning, which can lead to issues when dealing with open-textured terms. A suggestion to deal with this has been the construction of hybrid systems, that employ both rule-based and case-based reasoning.

One famous such hybrid system is CAse-BASed REasoning Tool (CABARET). It contains both a rule-based reasoning system, that is able to represent legislation in an area, and a case-based reasoning system that is able to reason with cases. CABARET was focused on income tax home-office deductions. The rule-based system was able to forward-chain from the provided facts to arrive at an outcome, or backwards-chain from the desired goals to arrive at facts that need to be proven. If the system arrived at a statutory term, that was not covered by legal rules (i.e. an open-textured term), it used case-based reasoning to identify cases that were relevant to the current fact situation, using the methodology developed in CATO.⁷⁴⁰ It had several ways to switch between case-based reasoning and rule-based reasoning.⁷⁴¹

⁷³⁸ Satoh et al, “PROLEG”, *supra* note 48.

⁷³⁹ *Ibid.*

⁷⁴⁰ Ashley, *supra* note 44 at 88; Rissland & Skalak, “CABARET”, *supra* note 50 at 852–853.

⁷⁴¹ Ashley, *supra* note 44 at 88; Rissland & Skalak, “CABARET”, *supra* note 50 at 856–861.

Above, I described systems such as CATO, IBP and VJAP.⁷⁴² These systems have a “domain model”, which consists of a “graph of trade secret law issues that semantically interconnected factors”.⁷⁴³ While these systems primarily use these models to retrieve and compare cases, they contain elements of rule-based reasoning, and can thus be seen as hybrid models.

4.5.3.2.7 Susskind

There have also been projects that aim to concretely implement rule-based reasoning systems in expert systems. Susskind argues that there are many areas of clear cases, where open-textured terms are not important, and rules are sufficient. Building expert systems in these areas could be very useful.⁷⁴⁴

Susskind worked on a project to implement the rules of Scottish divorce law using these insights. The system incorporates a number of rules that the user traverses by responding “Yes”, “No” or “I don’t know”. Based on this, the system can provide information on whether a court may grant divorce. In case of the “I don’t know” answers, it allows the user to explore what would hypothetically occur if the answer were yes or no.⁷⁴⁵

4.5.3.2.8 Loge-expert

Loge-expert was an expert system developed in Montreal in the domain of Québec Housing Law,⁷⁴⁶ coincidentally the domain where we built the first JusticeBot tool. Loge-expert focused on disseminating legal knowledge to the layperson user.⁷⁴⁷ The creators started building the system by understanding the concepts involved in repossession

⁷⁴² See 4.4.3.4.2.2.

⁷⁴³ Ashley, *supra* note 44 at 116.

⁷⁴⁴ Susskind, “Expert systems in law”, *supra* note 598 at 3.

⁷⁴⁵ *Ibid* at 4.

⁷⁴⁶ Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563; Claude Thomasset, François Blanchard & Louis-Claude Paquin, “Loge-expert: an illustration of different phases of the development of an expert system in Law” Université du Québec à Montréal.

⁷⁴⁷ Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563 at 2.

disputes, from the Civil Code, Court and Tribunal decisions, legal doctrine and know-how of legal experts.⁷⁴⁸

The authors found that even simple, heterogenous legal domains, such as repossession, required the modelling of a large amount of legal knowledge, including the general theory of obligations.⁷⁴⁹ The researchers identify the subsequent reasoning steps that legal practitioners undertake when dealing with issues of certain kinds, such as repossession.⁷⁵⁰ They model this in an expert system called D_expert.⁷⁵¹ They also aim to integrate a textual database of case law into the system.⁷⁵²

The researchers deal with the translation of general language to legal language, and vice-versa. For this, the researchers choose a graphical interface.⁷⁵³ They separate the legal reasoning layer from the communication layer, in order to communicate results to the user in plain language. For each question, they allow the user to explore plain-text definitions, references and jurisprudence relating to a concept.⁷⁵⁴ The system provides multiple ways of exploring these data sources, including keyword searches.⁷⁵⁵

4.5.3.3 Conclusion

As we have seen, there are a number of ways to represent legislation in logical frameworks. All of the presented research uses logical pathways to represent legislation, which can reduce ambiguity in laws and be used as an automatic way to apply legal rules to fulfilled legal criteria.

⁷⁴⁸ *Ibid* at 5.

⁷⁴⁹ Claude Thomasset, François Blanchard & Louis-Claude Paquin, “Loge-expert: Strategies to integrate legal knowledge modelization, non-expert user interface, and textual data base into the development of an expert system in law” (1992) 4:4 Expert Systems with Applications (Special Issue: Expert Systems and Law) 379–395 at 380.

⁷⁵⁰ *Ibid* at 383.

⁷⁵¹ *Ibid* at 380; Thomasset & Paquin, “Expert Systems in Law and the Representation of Legal Knowledge”, *supra* note 562 at 5.

⁷⁵² Thomasset, Blanchard & Paquin, “Loge-expert”, *supra* note 749 at 385.

⁷⁵³ *Ibid* at 386.

⁷⁵⁴ *Ibid* at 389.

⁷⁵⁵ *Ibid* at 390–393.

Walker and Satoh introduce systems that can additionally be used to structure evidence and arguments around a case. In Satoh's case, the legal rules are flexible, and change depending on the arguments raised by either party. These systems seem like they can be helpful for individuals to structure their arguments. CABARET additionally introduces the capability to use case-based reasoning to deal with open-textured concepts.

I have also explained two projects that take a rule-based methodology and use it to build expert systems that can support individuals. Susskind's system focuses on supporting lawyers, while the Loge-expert system is aimed at providing legal information to laypeople.

Let us explore some of the challenges associated with encoding legal rules into rule-based systems.

4.5.3.3.1 Effort to encode legal rules

Many of the researchers highlight the time and effort required to build legal expert systems as a potential problem of the methodology. For example, the authors of Loge-expert found the development of their system to be very time and energy-consuming.⁷⁵⁶ Therefore, they shift the focus from automated legal decision-making towards information systems that help human decision-making, using a database of decisions and the functionality to easily search such a database. They see the development of such a system as the better use of their financial and human resources.⁷⁵⁷ Likewise, Susskind argues that the time commitment required to build legal expert systems, and the lack of a common methodology, has led to the comparative dearth of such systems.⁷⁵⁸

Not all authors agree that the effort required to encode legal rules is an issue in itself. Sergot argues that the rule-based representation of the law is similar to an expert system but does not require the arduous task of eliciting the knowledge from an expert, since the

⁷⁵⁶ Paquin, Blanchard & Thomasset, "Loge-expert", *supra* note 563 at 10.

⁷⁵⁷ *Ibid* at 11–14.

⁷⁵⁸ Susskind, "Expert systems in law", *supra* note 598 at 4.

laws contain the required rules in plain text.⁷⁵⁹ Popple criticizes this view, arguing that interpretation of the law is an expert activity, and that it may require legal expertise to correctly interpret the statutes in a certain domain.⁷⁶⁰

Overall, it seems like the development of tools and methods that can increase the efficiency of creating such expert systems would be an important step in making sure that a resulting methodology can have practical applications.

4.5.3.3.2 How to deal with evidence?

Another issue faced by rule-based reasoning tools is the difficulty of dealing with the steps of legal reasoning that we previously discussed, such as evidential reasoning.

With regards to evidence, it is of course important to remain aware that judges will only consider facts that are proved or undisputed, as discussed in 4.3. Waterman *et al* discuss tackling this issue by introducing a probability factor that allows the adjustment of the final outcome based on the confidence level of the user of the system.⁷⁶¹

Walker and Satoh both use the rule-based systems to allow the structuring of evidence. Walkers default logic paradigm allows the connection of evidential reasoning to specific legal criteria, which can help with building evidential arguments. Satoh uses his JUF-theory framework, that can perform legal reasoning, to also provide a way to discover where evidence is needed, and what happens if a burden of proof is not discharged. These ways of dealing with evidence avoid the difficulty of automating evidential assessment, while still providing useful tools to structure evidence.

4.5.3.3.3 How to deal with open-textured legal concepts?

In the same vein, there seem to be a number of different views around how to deal with open-textured legal concepts in rule-based reasoning systems. Susskind argues that the

⁷⁵⁹ Sergot et al, *supra* note 48 at 383.

⁷⁶⁰ Popple et al, *supra* note 644 at 32–33.

⁷⁶¹ Ashley, *supra* note 44 at 10.

assessment of open-textured terms is not necessary for many straightforward cases.⁷⁶² In my research into the area of landlord-tenant disputes, it seems like open-textured terms such as “frequently late” and “peaceable enjoyment” are relatively common. It would thus be advantageous to have a system that can deal with such terms.

One straight-forward approach in dealing with vague terms relies on asking the user whether a certain term applies in their case or not. For example, Waterman and Peterson suggest using this approach, potentially in connection with examples from case law of how the concepts were previously applied.⁷⁶³ Sergot *et al* seem to rely on a similar approach.⁷⁶⁴ Loge-expert adds tools to explore explanations from case-law and plain-text explanations to support the user in making the determination.⁷⁶⁵

Asking the user to decide open-textured concepts without any context may be an issue when the system is targeted at laypeople. Laypeople may not think of their situation in terms of a legal concept, but rather in terms of things that have happened. Asking the user to specify whether a legal concept (such as “unreasonable”) applies in a case, without providing any context, may therefore be less helpful – how can the user know how a judge would assess their factual situation?⁷⁶⁶

An interesting approach to overcome the issue of vague concepts is combining case-based reasoning systems with rule-based reasoning systems. The case-based reasoning system can handle the determination of whether an open-textured term applies, while the rule-based reasoning system can deal with deciding the consequences of this. CABARET is an example of such a system, which uses CATO-style cases and rules to reason about cases in income tax home-office deductions.

⁷⁶² Susskind, “Expert systems in law”, *supra* note 598 at 3.

⁷⁶³ Waterman & Peterson, *supra* note 724 at 26.

⁷⁶⁴ Sergot et al, *supra* note 48 at 371; Ashley, *supra* note 44 at 51.

⁷⁶⁵ Thomasset, Blanchard & Paquin, “Loge-expert”, *supra* note 749 at 389.

⁷⁶⁶ Compare Richard E Susskind, “Expert Systems in Law: A Jurisprudential Approach to Artificial Intelligence and Legal Reasoning” (1986) 49:2 *The Modern Law Review* 168–194 at 190.

4.5.3.3.4 How to deal with syntactic ambiguity?

As we discussed in 4.5.1, the statutes may also contain ambiguities in the way they are logically structured. When encoding such structures into a computer system, the author will likely have to make a choice on how to interpret the statutes. However, the individuals building the system may not know which way the legislator intended a rule to be read, or how a court would choose to enforce such a rule, if a party raised as an argument that it should be interpreted in a certain syntactical way.⁷⁶⁷ In these cases, a rule-based reasoning system may provide inaccurate information to the user. It is important to be aware that legal interpretation likely does play a role when designing such systems, and that their representations of the law are therefore not always neutral.

Dealing with syntactic ambiguity may be an AI-complete problem.⁷⁶⁸ Even lawyers may not always know how a rule should be interpreted, and therefore choose a way that they believe the rule should be interpreted, based on the intention of the legislator or reading the law as a whole. They may also choose to argue for the way that they want it to be interpreted. This kind of reasoning probably requires common sense and an in-depth understanding of legal nuances, and may thus be beyond the scope of today's AI systems.

It should be noted that this issue may be less pronounced in areas where there are many cases. Here, a consensus may have emerged in courts on how a law should be syntactically interpreted.

4.5.3.3.5 How to deal with colliding rules?

Another difficulty of the systems described above is dealing with colliding rules. As discussed, multiple rules may be applicable to a single situation, or certain rules may be in violation of other rules or human rights. This is not evident from reading the text of legal rules and would therefore not be captured from implementing a certain law in a logical system.

⁷⁶⁷ Ashley, *supra* note 44 at 45–46; Allen & Engholm, *supra* note 727.

⁷⁶⁸ See 2.4.2.4

Just like the previous point, this shortcoming may be difficult to overcome without general artificial intelligence. Understanding the implications of rules for human rights or other competing rules requires an in-depth understanding of how rules operate and their implications.⁷⁶⁹ Again, these problems are perhaps less likely to arise in areas of high-volume cases, since the different statutes that could be relevant are likely to have been explored more frequently in these cases.⁷⁷⁰

4.6 Determining the outcome of a case

4.6.1 The legal system

The judge has now established the evidence, determined which legal criteria are fulfilled, and connected the legal criteria to a possible outcome. In some cases, this may be sufficient to provide a ruling on a case. For example, if the outcome is specific performance of an action (such as repairing the heater), the judge may render a decision that the landlord has to repair the heater.

In other cases, however, the judge needs to decide some aspects of the case. For example, while the legislation specifies that the judge may order a reduction of rent, it does not specify how high this rent reduction shall be. Here, the judge must make a decision, on the merits of a case. This could depend on many factors, such as the inconvenience caused to the tenant, how long the issue persisted, the time of year the heating malfunction occurred and other similar questions. In other types of decisions, the judge may face similar questions, such as how much damages to award in tort cases etc. In criminal cases, after the defendant has been found guilty, the judge needs to determine the appropriate sentence for their crimes, such whether to order a prison sentence, the length of such a sentence, the amount of fines and any other consequences.

⁷⁶⁹ Compare Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797.

⁷⁷⁰ *Ibid* at 801.

This type of decision is usually seen as being discretionary, i.e. the judge can make the determination relatively freely.⁷⁷¹ The exact extent of the freedom granted to a judge in making these decisions depends on the specific domain. Zeleznikow *et al* identify four levels of discretion a judge may have:⁷⁷²

- No discretion – the judge has no discretion. For example, mandatory sentences imposing minimum sentence requirements.
- Narrow discretion – legislation, cases and/or legal opinions give a clear base level decision, which judges can deviate from to some extent. For example, sentencing in the presence of guidelines.
- Bounded discretionary domains – Judges are given a number of factors that they should consider in making a decision but can freely choose which weight to assign to each of these factors. For example, the distribution of property after a divorce in Australia, where contributions, future needs of the partners and wealth of the marriage should be taken into account.⁷⁷³
- Unfettered discretionary domains – Judges are not given any factors to consider, and can thus let any factor they see fit influence the decision. For example, the decision of who should get custody in Australian family law.⁷⁷⁴

Despite the discretionary character of these decisions, we still expect similar cases to be treated in a similar way. This may not always be the case – Stranieri *et al* note that contradictory cases (i.e. cases where the inputs are the same but the decision is different) are to be expected, as judges may weight factors differently.⁷⁷⁵ However, in general, we would expect decisions to align, in order to promote fairness.

⁷⁷¹ Lord Justice Bingham, “The Discretion of the Judge” (1990) 5 Denning LJ 27–44 at 28.

⁷⁷² Zeleznikow, *supra* note 703 at 344.

⁷⁷³ *Ibid.*

⁷⁷⁴ *Ibid.*

⁷⁷⁵ Andrew Stranieri et al, “A hybrid rule–neural approach for the automation of legal reasoning in the discretionary domain of family law in Australia” (1999) 7:2 Artificial intelligence and Law 153–183 at 165.

In some instances, jurisdictions introduce instruments in order to harmonize the way decisions are made by judges to achieve this similarity. In the United States, for example, the federal sentencing guidelines were introduced in 1987⁷⁷⁶ to harmonize “unjustifiably wide” sentencing disparities.⁷⁷⁷ 10,000 cases were analyzed in order to create a kind of algorithm allowing judges to get an indication of an appropriate length for sentencing based on factors such as the base offence, whether a weapon was used and if any money was stolen.⁷⁷⁸ For robberies, for example, they have to look at whether a weapon was used, the amount of money that was stolen and other similar questions.⁷⁷⁹

In France, such guidelines (referred to as “barèmes”) are in common use, for example to determine child support payments, compensation for work-related accidents and illnesses, in the evaluation of damages in personal injury law etc. Isabelle Sayn *et al* analyze the purpose and functioning of these tools. They are generally intended to increase equality of citizens, predictability of the legal system and the efficiency of judges. The researchers analyze the potential risks of these systems, such as potential discrimination stemming from the tools, a loss of the individualization of criminal sentencing, and the transformation of judging into a calculation.⁷⁸⁰

Such initiatives also exist in other jurisdictions. In Sweden, for example, the Victim Compensation and Support Authority has introduced a list of cases that show compensation awarded for crimes, together with the type and particularities of the crime, in order to support courts, lawyers and insurance agencies in determining appropriate compensation for criminal cases.⁷⁸¹ In Germany, the “Oberlandesgericht Düsseldorf” gives out a table that can be used to determine alimony payments. These instruments are

⁷⁷⁶ Stephen Breyer, “The Federal Sentencing Guidelines and the Key Compromises Upon Which They Rest” (1988) 17 Hofstra L Rev 1–50 at 1.

⁷⁷⁷ *Ibid* at 4–5.

⁷⁷⁸ *Ibid* at 6–7.

⁷⁷⁹ Breyer, *supra* note 776; Westermann, *supra* note 607 at 201.

⁷⁸⁰ Isabelle Sayn, *Le droit mis en barèmes ?* (Daloz, 2014); Muriel Rebourg, “Le droit mis en barèmes ? / Sayn Isabelle (dir.)”, (6 February 2015), online: *Droit & Société* <<https://ds.hypotheses.org/770>>.

⁷⁸¹ *Brottsoffermyndighetens Referatsamling 2016* (Brottsoffermyndigheten, 2016).

not binding, but exist to coordinate and standardize the decisions between different courts.⁷⁸²

4.6.2 Example

Let us consider how a judge may undertake the step of determining an outcome in our example case. The judge has determined that the tenant can be awarded a rent reduction. They see that the heating stopped working for 3 days, during the winter season. They determine that this has significantly diminished the peaceable enjoyment of that property and settle on a rent reduction of 10% of the rent for that month.

4.6.3 Automating the determination of the outcome of a case

4.6.3.1 Introduction

Modelling discretionary decisions taken by the judge could have a significant impact on access to justice. Often, individuals may not only be concerned with whether they would win or lose should they go to court, but also how much damages the judge might award them, should they win. This could inform the individual whether to take their case to court or not. It could also serve as a BATNA (Best Alternative to the Negotiated Agreement) for use in negotiation, to allow individuals to align their expectations and come to a settlement.⁷⁸³

On a technical level, the automation of the determination of the outcome of a case is not so different from the automation of determining whether certain facts fulfill a legal criterion or not (see 4.4.3). In both cases, the judge can more or less freely decide how to weigh a certain number of factors in the case to come to a decision. In determining legal criterion, however, the “output” is always whether a legal criterion applies or not. In determining the outcome of the case, the output is instead a decision such as the amount of damages to award, the length of a prison sentence or how to distribute assets after a

⁷⁸² “Düsseldorfer Tabelle”, (1 January 2019), online: *Oberlandesgericht Düsseldorf* <http://www.olg-duesseldorf.nrw.de/infos/Duesseldorfer_Tabelle/index.php>.

⁷⁸³ Zeleznikow, *supra* note 703 at 352; Dahan et al, “Predicting Employment Notice Period with Machine Learning”, *supra* note 460 at 1.

divorce. Stranieri and Zeleznikow agree that this type of discretion is different from the classification of vague terms discussed in 4.4, since even two judges who completely agree on the facts can arrive at different discretionary decisions.⁷⁸⁴

Many of the technical aspects described in 4.4.3.4 still apply to modeling this type of decisions.

4.6.3.2 Prior work

Stranieri *et al* built a system that aims to model the distribution of property between two parties in a divorce in Australia. This decision is discretionary and can be seen as a bounded discretionary domain according to the classification set out above, as there is a “shopping list” of factors that a judge should incorporate into the decision, but no specification of how those factors should be weighted.⁷⁸⁵ The researchers annotated 103 cases with a set of 94 factors.⁷⁸⁶ The system then uses a combination of rule-based reasoning and neural networks to predict the percentage allocated to each party.⁷⁸⁷ The system can further generate explanations for the predictions, based on the factors that were considered and the statutes supporting why these factors are important.⁷⁸⁸

Split-up was evaluated by comparing how specialist family lawyers assessed three cases, versus how the system assessed them. While two of the cases showed compatibility between the assessments of Split Up and the lawyers, the final case showed significant differences, due to the way the factor of who was the “homemaker” was assessed in the presence of paid staff.⁷⁸⁹

Dahan *et al* built a system to predict the notice period for individuals after a firing. They annotated 1,391 cases of notice periods for workers after a firing. For each decision, they annotated 14 factors, including the duration of the employment, the age of the claimant

⁷⁸⁴ Andrew Stranieri & John Zeleznikow, “The Role of Open Texture and Stare Decisis in Data Mining Discretion” (1998) JURIX 1998 11 at 104.

⁷⁸⁵ Stranieri *et al*, *supra* note 775 at 156.

⁷⁸⁶ *Ibid* at 164.

⁷⁸⁷ *Ibid* at 156–159.

⁷⁸⁸ *Ibid* at 173.

⁷⁸⁹ *Ibid* at 174–175.

etc.⁷⁹⁰ They then analyzed the decisions statistically and built algorithms to predict the notice period. While the prediction worked well on average, they found the prediction of notice periods in individual cases to be difficult.⁷⁹¹ The researchers analyzed the outliers in the prediction, i.e., where the predicted length of notice differed to the largest extent from the actual length of notice. They found that these cases often contained unique, exceptional situations that were not taken into account by the algorithm, since they were not annotated as factors. They argue that this is not an inconsistent application of the law, but rather the necessary exercise of discretion by the judges to deal with particular situations.⁷⁹²

The use of statistical tools in discretionary domains has been thoroughly discussed in the context of sentencing. Here, a number of jurisdictions utilize statistical tools to predict the risk that an individual will commit another crime, which then plays a role in deciding the sentencing of an individual.⁷⁹³ These tools do not usually aim to give an estimate of a predicted sentence, but rather aim to give a risk factor that the judge can use in their sentencing decisions. Further, they are used by judges, rather than being used as decision support systems for individuals. These systems have a different character than the ones I have previously discussed, but the extensive literature around the topic can nonetheless help us understand some of the issues associated with using automatic methods in discretionary decision making.

While such tools have been praised for potentially providing more accurate and consistent decisions, they have also faced significant criticism.⁷⁹⁴ They have been criticized for focusing on groups rather than individuals, by considering an individual as being made up of a number of factors and comparing them to a group of individuals with

⁷⁹⁰ Dahan et al, “Predicting Employment Notice Period with Machine Learning”, *supra* note 460 at 5.

⁷⁹¹ *Ibid.*

⁷⁹² *Ibid* at 25–29.

⁷⁹³ Hannes Westermann, “Evidence-Based Sentencing : Risks and Opportunities” (2020) 25:3 *Lex Electronica* 71–93 at 72–76.

⁷⁹⁴ Westermann, “Evidence-Based Sentencing”, *supra* note 793.

the same factors, rather than trying to understand the individual themselves.⁷⁹⁵ The systems have further been criticized for encoding historical biases, and potentially discriminating against certain ethnicities.⁷⁹⁶ Further, the systems have been criticized for weak performance⁷⁹⁷ and lacking transparency, which can make challenging and understanding the meaning of the results difficult for both the parties and the judge.⁷⁹⁸

4.6.3.3 Conclusion

This section has looked at the automation of discretionary decision making.

In a way, the guidelines prepared in many jurisdictions can be seen as a type of expert system that aims to predict and standardize the exercise of discretion based on certain factors. In areas where such guidelines exist, encoding them into a computer system and exposing them to laypeople may already provide a significant boost to access to justice, as individuals are able to learn what they could expect from the court deciding an outcome. However, the judge is usually given the option to diverge in case of extraordinary circumstances, something which would not be captured by such an expert system.

In areas where there are no guidelines, or to model how a judge may diverge from the guidelines, implementing machine learning based systems could be useful. As we have seen, a number of such systems have been implemented. These systems typically rely on the encoding of many decisions in terms of some kind of factors, and the prediction of an aspect of the decision (such as damages or a percentage distribution of property after divorce) based on these factors.

⁷⁹⁵ Starr, *supra* note 537 at 842–843.

⁷⁹⁶ Westermann, “Evidence-Based Sentencing”, *supra* note 793 at 79–80; Starr, *supra* note 537 at 806; Danielle Leah Kehl & Samuel Ari Kessler, “Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in Sentencing” (2017) Responsive Communities Initiative, Berkman Klein Center for Internet & Society, Harvard Law School, online: <<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041>> at 24; Angwin et al, *supra* note 276.

⁷⁹⁷ Westermann, “Evidence-Based Sentencing”, *supra* note 793 at 82–83.

⁷⁹⁸ *Ibid* at 82–87.

As discussed, this style of reasoning shares some aspects with the reasoning using case-law to assess whether a legal concept applies, presented in 4.4.3.4. However, the output is different. Further, the reasoning systems working on classification of facts often use bespoke systems to generate predictions and arguments using a limited number of cases (see i.e. HYPO and CATO). In this section, the prediction method used relies more on machine learning to discover patterns in large amounts of annotated cases. For a discussion of the different styles of reasoning with cases, see 4.9.

Aside from this, many of the issues discussed in 4.4.3.4.3 also exist in the systems presented in this section. Both Split Up and the system built by Dahan *et al* are built in domains with a limited number of factors, i.e. bounded discretionary domains in the classification proposed by Zeleznikow.⁷⁹⁹ Identifying factors in unbounded domains may be more difficult.⁸⁰⁰

Further, both researchers discuss cases where the system does not manage to accurately predict the outcome, due to extraordinary circumstances. The judge is always able to diverge and consider new factors. Machine learning models, on the other hand, are bound to only consider factors in the case representation provided.⁸⁰¹ Accurately modeling this kind of reasoning for every single case may thus be an AI-complete problem.⁸⁰²

4.7 Explaining the decision

4.7.1 The legal system

Once the judge has decided on an appropriate outcome, the final crucial step is explaining why they made that decision.⁸⁰³ This is usually done in the form of a decision document,

⁷⁹⁹ Zeleznikow, *supra* note 703 at 344.

⁸⁰⁰ Compare 4.4.3.4.3.2.

⁸⁰¹ Compare 4.4.3.4.3.2.

⁸⁰² See 2.4.2.4.

⁸⁰³ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 2; Finale Doshi-Velez et al, “Accountability of AI Under the Law: The Role of Explanation” (2019) arXiv:1711.01134 [cs, stat], online: <<http://arxiv.org/abs/1711.01134>> at 8–9 arXiv: 1711.01134.

that contains a textual explanation of the previously described steps. This explanation is seen as so important that failure to explain why a decision was reached can lead to that decision being invalidated in many jurisdictions and types of decisions, according to a survey by Finale Doshi-Velez *et al.* Only a few exceptions, such as the exclusion of testimony in the United States, and privacy-sensitive divorce decisions in France, do not require an explanation.⁸⁰⁴

This explanation is an important tool for a number of actors in the legal system:

For the *parties*, the decision is the explanation of why the judge reached a certain outcome. It provides the legal grounds, enabling them to understand the process. One of the key functions of the court system is giving individuals their day in court. The written document plays an important role in showing the parties that their arguments were heard and considered. The explanation and communication of the reasoning behind a decision gives a decision its legitimacy, as the parties can verify that the decision was reached in accordance with the law, and not based on arbitrary whims of the judge. The explanation is also a crucial instrument in deciding whether they want to appeal the decision or not – only if the parties know why a decision is reached can they decide whether the judge has erred in applying the law, and a higher court might therefore overturn the decision.⁸⁰⁵

For *other participants in the legal system*, the decision is a statement by the judge on how a specific legislation should be interpreted. It can tell lawyers and parties how they might reason in court, and what they can expect from taking a case to court. The decisions are also very important for judges. Many jurisdictions have some notion that courts of a lower instance have to follow rulings of courts of a higher instance. This is especially pronounced in the common law, where the concept of stare decisis specifies that courts must generally follow decisions by an equal or higher instance court.⁸⁰⁶ The explanation

⁸⁰⁴ Doshi-Velez et al, “Accountability of AI Under the Law”, *supra* note 803 at 8–9.

⁸⁰⁵ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 2.

⁸⁰⁶ Holland & Webb, *supra* note 552 at 166.

thus becomes an important element in shaping future decisions. In 4.9, I will examine how legal case law is used in the common versus the civil law.

For the *public*, the explanation is an important tool to ensure the rule of law in a certain jurisdiction. The legislator tends to give general rules, but does not specify how these should be applied in real situations. Only by reading the explanations provided by the courts can the public understand how a legislation is applied in reality. This is an important step for individuals in determining how the law might apply to them, and how to act accordingly. This predictability of the law is a crucial aspect of the rule of law in society.⁸⁰⁷ Further, the democratic process relies on the public being able to understand how laws are applied in practice, to be able to criticize or vote for alternatives if laws are applied in ways that are not desirable to individuals. Without the explanation provided by the judge, the public can only guess as to how laws are applied in practice and why, limiting their ability to participate in the democratic process.

What exactly, then, is an explanation? Atkinson *et al* suggest that we look to the social sciences to understand particularities of decision explanations. Miller arrives at four points that are specific to how explanations are made:⁸⁰⁸

- Explanations are often contrastive, i.e. they describe why a decision was taken as opposed to another decision.
- Explanations are selected, i.e. curated to include a few key reasons, rather than any possible reason. As Yablon notes, there could be many ways to explain why a fire started in an apartment, such as the presence of oxygen in the air. However, a better explanation might be that someone smoked in their bed.⁸⁰⁹
- Explanations are rarely based on probabilities.

⁸⁰⁷ See 3.4.2.

⁸⁰⁸ Tim Miller, “Explanation in artificial intelligence: Insights from the social sciences” (2019) 267 *Artificial Intelligence* 1–38 at 6; Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 2.

⁸⁰⁹ Charles M Yablon, “The Indeterminacy of the Law: Critical Legal Studies and the Problem of Legal Explanation Responsive Scholarship from outside the Movement” (1984) 6:4 *Cardozo L Rev* 917–946 at 926–927.

- Explanations are social, i.e. are a part of a transfer of knowledge in an interaction or conversation.

We can see many of these factors apply in legal decisions. Legal decisions often seem to gloss over undisputed or obvious conclusions to legal criteria, choosing instead to focus on the contentious points. They often use contrastive explanations, such as saying what would have to occur for another outcome to have been selected, while rarely relying on probabilities. Further, judges are clearly aware of the many functions of the judgment to transfer knowledge, often shaping their explanations to have the potential to teach the parties, participants in the legal system and the public why a decision was reached. Courts can be seen as being in conversation with other courts, by referencing and analyzing previous decisions.

4.7.2 Example

Let us examine how a judge may explain the decision taken in our example mentioned above. After deciding to award the tenant with a rent reduction, they prepare a written document to summarize their arguments. Here, they go through the different steps they performed to arrive at the conclusion, including the analysis of the evidence, the establishing of legal criteria, the application of the laws and the determination of an outcome. The explanation is likely to focus on key disputed issues, such as whether the proof is sufficient to show that the heating was not working, and whether this counts as affecting the peaceable enjoyment of a property by the tenant. Issues such as whether a lease exists are less likely to be featured prominently in the decision, as they were not disputed by the parties. Thus, the parties know why the judge assumed that the lease existed, and do not need an explanation for this.

4.7.3 Automating explaining the decision

Let us examine how explanation works in automated legal reasoning tools. As we have seen, explanation is a crucial feature of legal decision making. In developing software tools to predict legal outcomes, often an explanation is even more important than the

prediction itself.⁸¹⁰ A prediction without an explanation merely tells us what the system thinks of our case. As we have seen, predicting cases or performing legal reasoning in many instances relies on steps that are AI-complete, leading to potentially inaccurate outcomes. Without an explanation, the user cannot understand how the system works, and verify that the outcome is sensible.

With an explanation, the user can learn more about the law and their case by understanding the consequences of the factors present in their case, and how other factors may change the outcome. This can increase the trust of an individual in such a system. The explanation can also be used to generate arguments that the party can use in court. Further, explanations can lend legitimacy to the prediction by referencing relevant previous court cases and the applicable statutes and explaining how they apply to a case.

If a system is intended to assist a judge, the need for an explanation becomes even more pressing. It is unlikely that a system that, i.e. provides draft of judicial decisions would be taken seriously if it does not explain its decision, so that the judge can audit the functioning of the system.

Legislators have responded to this importance of explaining automatic decision making with regards to individuals, especially if such decisions have significant effects or are intended to be used by a judge to render a decision. For example, the General Data Protection Regulation (GDPR) in the European Union,⁸¹¹ arguably contains a right to an explanation. This right gives the individual affected by an automatic decision that produces legal or similar effects the right to an explanation about the logic involved in reaching the decision.⁸¹² Likewise, the European Union draft regulation on artificial intelligence includes a rule that AI systems “intended to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of

⁸¹⁰ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 1.

⁸¹¹ *General Data Protection Regulation*, *supra* note 278.

⁸¹² Selbst & Powles, *supra* note 278.

facts” should be considered as high-risk systems,⁸¹³ giving rise to a number of obligations including transparency and human oversight.⁸¹⁴

Due to this importance of explanations, the question of how predictions and reasoning can be explained has received significant attention in the field of AI & Law. What should be noted is that generating the explanation is typically not an additional step in the systems previously presented, but rather stems from the method used to perform the reasoning or prediction itself. In fact, generating the explanation for why a side might win was the purpose in creating many of the systems in the first place.⁸¹⁵

Atkinson *et al* wrote an overview over the different ways explanations have been generated in automated legal reasoning systems.⁸¹⁶ They explore the types of explanations by the different technologies used. Let us examine how explanations can be generated using some of the systems we have previously discussed.

4.7.3.1 Case-based reasoning systems

Case-based reasoning systems rely on explanations through example. Atkinson *et al* argue that explanations typically take the following form: ”the case should be decided in this way because it is like these cases, and unlike these other cases.”⁸¹⁷ Systems employing this methodology, such as HYPO, CATO and IBP (described in 4.4.3.4.2.2) generate explanations by highlighting similarities between cases favoring one outcome with the current fact pattern, but also providing a possible counter-argument that distinguishes those cases.⁸¹⁸

⁸¹³ *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS*, *supra* note 111 annex III, article 8(a).

⁸¹⁴ *Ibid* article 13, 14.

⁸¹⁵ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 2.

⁸¹⁶ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46.

⁸¹⁷ *Ibid* at 3.

⁸¹⁸ *Ibid* at 5.

4.7.3.2 *Rule-based reasoning systems*

Rule-based reasoning systems (such as the one described in 4.5.3), are also explainable. Here, the explanations take the form of showing which rules were traversed and which rules lead to which outcome. The rules can be referenced by the statute or case they were extracted from. These explanations can further be enhanced by adding the possibility to examine what happens if a condition would have been different (see e.g. the system in 4.5.3.2.7, which offers the possibility to see how unknown criteria affect the outcome) or by allowing the examination of multiple, contradictory rules (see 4.4.3.4.2.7).⁸¹⁹

4.7.3.3 *Machine learning systems*

Generating explanations for machine learning systems can be more difficult. Some argue that generating useful explanations about machine learning systems would require the development of a rule-based or case-based model in the domain, obviating the need for the machine learning system itself.⁸²⁰

Some machine learning models are explainable, such as decision trees and linear regression. In linear regression, for example, it is possible to examine the weight each factor is assigned, thereby showing how a decision was reached.⁸²¹ Of course, this may not lead to the type of explanations that we expect from the legal system.

Some algorithms, such as neural network, can contain billions of interacting parameters.⁸²² Generating explanations for these systems is very tricky, but is a crucial aspect of using these systems in the law. Explaining the predictions of these models is thus an active field of research.⁸²³ As we have seen in 4.4.3.2.3.2, textual models

⁸¹⁹ *Ibid* at 6–7.

⁸²⁰ *Ibid* at 14–15.

⁸²¹ Surden, “The ethics of artificial intelligence in law”, *supra* note 541 at 731.

⁸²² *Ibid*; Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 806–807.

⁸²³ Mouhamadou-Lamine Diop, “Explainable AI: The data scientists’ new challenge”, (14 June 2018), online: *Towards Data Science* <<https://towardsdatascience.com/explainable-ai-the-data-scientists-new-challenge-f7cac935a5b4>>; Alejandro Barredo Arrieta et al, “Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI” (2019) arXiv e-prints, online: <<http://adsabs.harvard.edu/abs/2019arXiv191010045B>>; David Gunning et al, “XAI—Explainable artificial intelligence” (2019) 4:37 *Science Robotics*.

sometimes face difficulties generating useful explanations for understanding why certain cases are successful and others are not. Atkinson *et al* explore hybrid systems, which add a layer of symbolic reasoning to neural models, as a possible solution,⁸²⁴ but argue that none of the current machine learning based systems are able to generate explanations with the same quality as explanations generated by traditional symbolic systems.⁸²⁵

We have now seen how judges and other legal decision makers tend to reach legal decisions. Often, such reasoning involves identifying a guiding legal rule, assessing the evidence to find the relevant facts, deciding whether these facts fulfill certain legal criteria, logically traversing a legal rule to arrive at an outcome, determining discretionary aspects of the outcome, and explaining the decision. As we have seen, some of these steps may be AI-complete, making it difficult to replace the decision maker with automated systems. Therefore, in this thesis, I focus on legal decision support tools, which can support the user in understanding their rights, but do not attempt to tell them what they should do.

In building this kind of system, it is important to understand the different legal sources that are available to us, that may influence the way legal decision makers apply the steps described above. Two such sources are likely to be the statutes (which stem from the legislative bodies) and legal decisions (which stem from the judiciary). Let us examine how the legislator and the judiciary interact to shape how rules apply in society.

4.8 The interplay between the legislator and judiciary

In automating judicial reasoning, it is important to understand the dynamics at play between the legislator and judiciary. In most jurisdictions, these are independent branches. However, they are nonetheless engaged in a kind of conversation as to how cases should

⁸²⁴ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 18.

⁸²⁵ *Ibid.*

be dealt with.⁸²⁶ Understanding this dynamic can help us understand which data sources are required to model legal reasoning in an area.

4.8.1 The legislator

The legislator deals with cases in the abstract. They are able to create and alter the structure of legal rules, for example by adding new prerequisites, exceptions or conclusions to legal rules, through the legislative process.⁸²⁷ In doing so, the legislature does not have a specific case in mind.⁸²⁸ Rather, it creates rules on how to deal with certain issues in the abstract.⁸²⁹ This is important, as it is impossible to think of all situations that may arise in advance. Further, leaving ambiguities in the text may facilitate political compromise.⁸³⁰ Ashley and Rissland compare this to the process of “eager learning” in computer science, where an algorithm is devised to solve new cases before a new query is observed.⁸³¹

In our *example*, the legislator has decided that the tenant has the right to a rent reduction if the landlord fails to perform their duties. They may also specify what these duties consist of, such as arguing that a landlord has to provide the tenant with peaceable enjoyment of a property. However, they do not specify exactly what is meant by “peaceable enjoyment” – the interpretation of this criteria is left to the courts. As described above, the rule implements an abstract norm, that has to be applied to particular facts by courts. This allows the rule to work in any conceivable situation that may consist of peaceable enjoyment of a property, without the legislator having to enumerate all of these situations.

⁸²⁶ Ashley & Rissland, *supra* note 602 at 21.

⁸²⁷ *Ibid* at 18.

⁸²⁸ *Ibid* at 19.

⁸²⁹ Jan Komárek, “Reasoning with Previous Decisions: Beyond the Doctrine of Precedent” (2013) 61:1 American Journal of Comparative Law 149–171 at 158.

⁸³⁰ Ashley, *supra* note 44 at 40.

⁸³¹ Ashley & Rissland, *supra* note 602 at 19.

4.8.2 The judiciary

The judiciary, on the other hand, deals with solving individual cases. Faced with a set of facts, the court needs to decide how the law should apply in a specific situation. While the decision only targets a single situation, the court may be mindful of how the reasoning and decision may apply in future cases.⁸³²

While the legislator is relatively free in deciding the content of an abstract norm, the judiciary is more constrained in reaching a decision. Firstly, they are constrained by the structure set out by the legislator. This is the case especially in civil law systems.⁸³³ In common law systems, on the other hand, legislation seems to play a lesser role in judiciary decision making – in fact, entire legal areas are frequently created by courts, such as the area of product liability.⁸³⁴

However, the judiciary is able to determine how the constituent terms of a legislation is interpreted, in light of the facts that arise in a certain case.⁸³⁵ Ashley and Rissland compare this sort of reasoning to “lazy learning”, a learning method in computer science where the algorithm to classify a new case is only devised once the computer is faced with the new case.⁸³⁶ Of course, even here, courts are not able to arbitrarily decide the legal criteria that a set of facts corresponds to. Often, courts are bound by previous decisions.⁸³⁷ Below in 4.9, I will describe how such reasoning differs between the common law and the civil law.

In our *example*, the court is faced with a situation of inadequate heating, and needs to determine whether this should be seen as the tenant lacking peaceable enjoyment of the property. The legislator has not given guidance about this specific situation – as such, the judiciary is responsible for filling out what “peaceable enjoyment” means in this context.

⁸³² *Ibid.*

⁸³³ Joseph Dainow, “The Civil Law and the Common Law: Some Points of Comparison” (1966) 15:3 *The American Journal of Comparative Law* 419–435 at 433–434.

⁸³⁴ Ashley & Rissland, *supra* note 602 at 22–25.

⁸³⁵ *Ibid* at 20.

⁸³⁶ *Ibid.*

⁸³⁷ *Ibid* at 19.

In reaching this decision, the court may be constrained by previous decisions from superior instances that have reached a decision regarding this interpretation (see 4.9).

4.8.3 Conclusion

As we can see, both the statutes enacted in a legal area, and the court decisions that interpret the statutes are important legal sources when aiming to build legal decision tools. They are complimentary systems that interact to give a clear picture over how cases in an area should be dealt with. Reading either just legal decisions or legal cases does not give a full picture over how the law applies. Therefore, I believe it is important to build legal decision tools that model both legal rules and legal cases. Such systems are referred to as hybrid systems, and include systems such as CABARET and IBP, described above. Popple similarly argues that case-based and rule-based approaches should be combined in order to build legal expert systems.⁸³⁸

As discussed, there may be different ways that courts can reason with previous decisions. Understanding these ways is crucial in deciding how to automate this kind of reasoning. Thus, let us explore different fashions of reasoning with case law.

4.9 Reasoning with Case Law

Now that we have determined that reasoning with case law is often an important component of building legal decision support tools, let us explore how such reasoning is done in the legal system. Generally, we expect cases that are similar to be treated similarly.⁸³⁹ What makes cases similar, and how previous case law is reasoned with, depends to a large extent on the legal system. Let us explore how cases are reasoned with in common law systems and civil law systems, and how courts align their decisions to harmonize some decisions for cases.

⁸³⁸ James David Popple, *Legal Expert Systems: The Inadequacy of a Rule-based Approach* (Oxford: Blackwell, 1987).

⁸³⁹ Ashley & Rissland, *supra* note 602 at 28.

Of course, it should be noted that common law and civil law are far from the only legal systems in the world. Peter Glenn, in his work “Legal Traditions of the World”, distinguishes between 7 legal traditions (chthonic, talmudic, civil, islamic, common, hindu and confucian), and states that there are even more traditions.⁸⁴⁰ Glenn argues for conceptualizing the world in terms of legal traditions, passed on through time, rather than legal systems, which allows the non-conflictual understanding of the different legal methods.⁸⁴¹ These points are well taken, but beyond the scope of this thesis. Here, I focus on common law and civil law, since these systems are prevalent in Quebec, and have been the focus of most AI & Law research. However, I believe that studying how AI & Law research may interact with the other legal traditions to be a fascinating area of future work.

4.9.1 Common Law

In common law systems, case law is the basis of the law.⁸⁴² Historically, in England, courts were the most important source of law. The law consisted of rules that could be generalized from previous court decisions.⁸⁴³

Stability of this system is guaranteed by the doctrine of precedent.⁸⁴⁴ In common law systems, courts are bound by a concept referred to as *stare decisis*. This means that after a court has made a decision in a certain case, any court of an equal or lower level has to follow the reasoning in this case. A new case that is similar to previously decided cases has to be decided in the same way.⁸⁴⁵ Of course, this raises the questions of which cases are similar enough that the reasoning has to apply, and which cases have differences that are substantial enough that the reasoning of the old case does not have to apply in the

⁸⁴⁰ Patrick Glenn, *Legal Traditions of the World: Sustainable Diversity in Law*, 5th ed (Oxford, New York: Oxford University Press, 2014) at 361.

⁸⁴¹ H Patrick Glenn, “Doin’ the Transsystemic: Legal Systems and Legal Traditions” (2005) 50 McGill LJ 863–898 at 897–898.

⁸⁴² Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 434.

⁸⁴³ *Ibid* at 424–425.

⁸⁴⁴ *Ibid* at 425.

⁸⁴⁵ Holland & Webb, *supra* note 552 at 166; Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 425.

new case.⁸⁴⁶ The judge has significant flexibility in making this decision, and can thus allow social factors to influence whether a previous decision should control a new situation or not.⁸⁴⁷

According to Levi, the reasoning process in the common law consists of the following steps:⁸⁴⁸

- Similarity is seen between cases
- The rule of law inherent in the first case is announced
- The rule of law is made applicable to the second case

Establishing the *similarity* between the cases is thus the first step. Of course, no two cases are exactly the same. Likewise, every case has something in common with every other case.⁸⁴⁹ The judge thus has to consider whether the similarity between two cases is *relevant*. A previous case is relevant to a current case if the reasoning from the previous case can be applied to the current case.⁸⁵⁰ In the common law, this means that the material facts (i.e. the facts that are relevant to the legal issue at hand)⁸⁵¹ are shared between the two cases. The judge may also choose to distinguish a previous case, by arguing that it differs in some important way from the current case and does therefore not have to be applied.⁸⁵² By distinguishing previous cases a new decision can be made consistent with all previous decisions in a certain area, even when those cases point in different directions.⁸⁵³

The judge of the current case has a margin of discretion in determining which facts of the previous case are material facts. They can reinterpret previous cases, to argue that certain facts that the initial judge thought were material are in fact not material, and vice-

⁸⁴⁶ Holland & Webb, *supra* note 552 at 167.

⁸⁴⁷ Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 425.

⁸⁴⁸ Levi, *supra* note 593 at 501–502.

⁸⁴⁹ Cass R Sunstein, “On Analogical Reasoning” (1993) 106:3 Harvard Law Review 741–791 at 774.

⁸⁵⁰ Holland & Webb, *supra* note 552 at 197.

⁸⁵¹ *Ibid* at 227.

⁸⁵² *Ibid* at 218–219.

⁸⁵³ Komárek, “Reasoning with Previous Decisions”, *supra* note 829 at 152.

versa.⁸⁵⁴ Further, they can choose which level of abstraction material facts should be understood at.⁸⁵⁵ For example, a previous case involving a car could be seen as covering other cases involving to a) cars of the same brand, b) any car, or c) any type of moving vehicle. This, of course, has an impact on which cases are similar, and which are different. Levi sees this as the rules being discovered as they are applied to new cases.⁸⁵⁶

Once a case has been found to be similar, the judge must determine the *ratio decidendi* of a previous case. There seems to be some debate about exactly what is meant by this, but in general it seems to refer to the binding part of a decision, i.e. the material facts and the reasoning applied to them.⁸⁵⁷ This *ratio decidendi* is the principle that has to be applied in the current case, if the previous case cannot be distinguished. What exactly the *ratio decidendi* of a previous case is can also be reinterpreted by judges, for example by altering which of the facts are considered material.⁸⁵⁸ Parts of a judgment that are not part of the ratio decidendi are referred to as *obiter dictum*, and are not binding.⁸⁵⁹ It should be noted that *the ratio decidendi* is generally seen as being how a case was decided, rather than the words used by a judge in a decision.⁸⁶⁰

Once two cases have found to be similar, and the ratio decidendi has been identified, the judge applies the previous principle to come to a decision on the current case.⁸⁶¹ In doing so, they introduce a new interpretation of the *ratio decidendi* of a previous case, which will be considered by judges who later face similar cases. In this way, the law develops and adapts to new situations.⁸⁶²

Legislation also plays a role in common law. However, historically, judges mistrusted the legislation, which was introduced by parliament, while the courts were tied to the king.

⁸⁵⁴ Ashley & Rissland, *supra* note 602 at 27.

⁸⁵⁵ Holland & Webb, *supra* note 552 at 228–229.

⁸⁵⁶ Levi, *supra* note 593 at 502.

⁸⁵⁷ Holland & Webb, *supra* note 552 at 201–206.

⁸⁵⁸ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1596.

⁸⁵⁹ Holland & Webb, *supra* note 552 at 210–212.

⁸⁶⁰ Komárek, “Reasoning with Previous Decisions”, *supra* note 829 at 151.

⁸⁶¹ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 432.

⁸⁶² Westermann, *supra* note 607 at 198.

Therefore, the courts would often interpret legislation very narrowly, in order to retain decision areas. Parliament responded by drafting very specific legislation to ensure the fulfillment of the legislative intent.⁸⁶³ Likewise, legal doctrine plays a much smaller role in the common law than in civil law countries, and usually focuses on analyzing decided cases.⁸⁶⁴

Many automated legal reasoning systems were created in common law systems.⁸⁶⁵ For example, HYPO, CATO, IBP and SHYSTER all stem from common law domains. These systems heavily focus on representing case law and comparing previous cases to a current fact pattern by the user. The systems are then able to highlight similarities between previous cases and the current case or distinguish previous cases that have undesirable conclusions. As we have seen, this type of reasoning is very typical for the common law.

4.9.2 Civil Law

In civil law systems, there is typically a much larger focus on codified law than in common law systems.⁸⁶⁶ Many such jurisdictions have civil codes, which are books that contain rules regarding the relationship between individuals, aiming to provide comprehensive principles for the entire subject matter of civil law, including persons, family, ownership, contracts, obligations and sales etc.⁸⁶⁷ This code, and other statutes in civil law jurisdictions, are seen as the basic source of the law.⁸⁶⁸

The legal reasoning procedure in civil law systems is much more focused on the logical deductive part described in 4.5, where the reasoning starts with a broad principle, then considers the facts of the case and finally logically applies the principles to the facts, in

⁸⁶³ Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 425–426.

⁸⁶⁴ *Ibid* at 428.

⁸⁶⁵ Kevin D Ashley, “Case-Based Models of Legal Reasoning in a Civil Law Context. Invited paper” (2004) International Congress of Comparative Cultures and Legal Systems of the Instituto de Investigaciones Juridicas, Universidad Nacional Autonoma de México.

⁸⁶⁶ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1594.

⁸⁶⁷ Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 424.

⁸⁶⁸ *Ibid*.

order to arrive at an outcome.⁸⁶⁹ There is an assumption that there is a *right* way to rule in every case, and that disagreement must be due to a mistake in the reasoning process.⁸⁷⁰ The law is seen as stemming from rationality, and is used as an instrument of reason, to construct the modern state.⁸⁷¹ The judge thus has the duty to apply the written law, as elaborated by the statute and academic commentators.⁸⁷² While the judge can interpret the rules, they are not seen to change them – this is left to the legislator.⁸⁷³

However, this does not mean that case law does not play a role in civil law systems. In general, there is no concept of binding precedent, where cases of the same or other courts have to be followed. Each new decision is expected to be grounded in legislation.

However, certain courts may be more authoritative in deciding how a law should be interpreted, meaning that lower courts are expected to follow their reasoning.⁸⁷⁴

Compared to the common law, the facts of a case play a lesser role in the civil law. Rather, the clarification of how a statute should be interpreted is the core component of cases that will be considered by succeeding courts. The European Court of Justice, for example, often concludes its decisions with a bold section, that states a textual explanation of the rule of a decision. This is interpreted almost as in legislation, where the text itself matters more than the context and facts of the case, comparable to interpreting statutory law.⁸⁷⁵ Komarek refers to this style of reasoning as the “legislative” reasoning with cases, while the common law approach is the “case-bound” way of reasoning with previous decisions.⁸⁷⁶ However, he also points out that civil law courts occasionally engage in case-bound reasoning, and common law courts also engage in

⁸⁶⁹ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1596; Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 428.

⁸⁷⁰ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1596.

⁸⁷¹ Glenn, *supra* note 840 at 151–153.

⁸⁷² Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1597; Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 428.

⁸⁷³ Stein, “Roman Law, Common Law, and Civil Law Symposium”, *supra* note 549 at 1600.

⁸⁷⁴ Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 426.

⁸⁷⁵ Komárek, “Reasoning with Previous Decisions”, *supra* note 829 at 156–157.

⁸⁷⁶ *Ibid* at 157–158.

legislative reasoning.⁸⁷⁷ Ashley seems to make a similar distinction, referring to style of reasoning relying on the tests formulated by a previous decision as “Abstract Precedent Scenario”, versus the reasoning relying on factual occurrences, which he refers to as “Fact-Based Precedent Scenario”.⁸⁷⁸

Legal decisions thus look very different in civil law compared to in the common law. They are typically much shorter. They do not focus on the facts of a case, but rather give a short overview, and then explain which principles and rules are applicable to a case. Finally, the judge concludes the results stemming from the application of the law to the case.⁸⁷⁹

There have been a number of AI & Law reasoning systems built in civil law countries. For example, PROLEG (see 4.5.3.2.5), is built in Japan, which is a civil law system.⁸⁸⁰ Likewise, the Loge-expert system (see 4.5.3.2.8) was built in a civil law context, modelled after the Code Civil du Quebec.⁸⁸¹ Compared to some of the systems based on the common law, we can see that these systems rely much more heavily on statutes rather than case law. Even in areas where case law is used (for example in Loge-expert to provide information)⁸⁸², the cases are presented as being oriented around rules stemming from legislation, rather than being recorded based on their factors and outcomes, as was the case in many of the systems based in common law systems.

4.9.3 Harmonization of decision making

Previously, we discussed ways of qualitatively reasoning using previous case law. The decision, which often stems from a higher court, is read, analyzed, and applied to the current case. This can occur by establishing the similarity of cases based on the facts of a

⁸⁷⁷ *Ibid* at 158–160.

⁸⁷⁸ Ashley, *supra* note 865 at 5.

⁸⁷⁹ Dainow, “The Civil Law and the Common Law”, *supra* note 833 at 432; Ashley, *supra* note 865 at 3.

⁸⁸⁰ Veronica Taylor et al, “Introduction: Nature of the Japanese Legal System” (2008) 1 Business law in Japan 3–8 at 1.

⁸⁸¹ Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563 at 254.

⁸⁸² *Ibid* at 257.

case, and then applying the reasoning to the current case, as is often the case in common-law systems. It can also occur by reading the previous decision as an application and clarification of the law, which can guide judges in how to apply the legislation to specific cases, as is often the case in the civil law.

However, there also seems to exist another form that judges may seek guidance from previous cases. In some cases, the judges need to make discretionary decisions. This could be, for example, the amount of damages to award to an individual based on a certain occurrence, or the length of a prison sentence for an individual.⁸⁸³

Such decisions can be seen to rest in the discretion of the judge, i.e. they can freely decide on the appropriate consequence.⁸⁸⁴ However, we also expect a certain level of consistency in the decisions. To maintain trust in the judicial system, it is important that decisions about cases that are similar also have similar outcomes. This has not always been the case – in Maryland, for example, prison sentences were estimated to differ by up to 20% based on the ethnicity of the defendant.⁸⁸⁵

Lawlor differentiates between three types of *stare decisis*:⁸⁸⁶

- **Traditional stare decisis** – This seems to be the kind of qualitative reasoning discussed above in 4.9.1 and 4.9.2, where similar decisions from the same or higher courts lead to the same decision.
- **Local stare decisis** – Judges at the same court try to exercise discretion in a consistent manner.
- **Personal stare decisis** – An individual judge is likely to try to exercise discretion in a consistent manner.

⁸⁸³ See 4.6.

⁸⁸⁴ Lord Justice Bingham, *supra* note 771 at 28.

⁸⁸⁵ Shawn D Bushway & Anne Morrison Piehl, “Judging Judicial Discretion: Legal Factors and Racial Discrimination in Sentencing” (2001) 35:4 Law & Society Review 733.

⁸⁸⁶ Reed C Lawlor, “What Computers Can Do: Analysis and Prediction of Judicial Decisions” (1963) 49:4 American Bar Association Journal 337–344 at 340; Stranieri & Zeleznikow, *supra* note 784 at 105–106.

Traditional *stare decisis* relies on landmark cases, that change the law and are likely to be cited as precedents for a certain decision. In local and personal *stare decisis*, landmark cases are less useable, since they do not allow the discovery of how discretion is *usually* exercised, but rather focus on definitional issues or ambiguity. Instead, commonplace cases have to be used, that show how rules are applied in regular court cases.⁸⁸⁷

It seems like the implementation of guidelines, as described in 4.6.1, can be seen as an attempt of implementing local *stare decisis*. The guidelines typically attempt to encode many commonplace cases into an easily applicable algorithm, that gives the judge a way to calculate an outcome that they can then diverge from in exceptional cases.

These kinds of guidelines may also play an important role in administrative decision making, such as the decision on whether to grant social aid, or whether a person should be allowed to immigrate.⁸⁸⁸ Here, the guidelines set out which factors the decision maker should consider, in order to guarantee a fair and consistent application of the rules. These factors could stem from previous decisions where the questions were tried in court, or internal policy decisions.

As we can see, this type of reasoning with previous decisions is different from the common and civil law reasoning described above. The individual cases matter less, and are not necessarily cited or considered in detail to reach a decision. Rather, the aggregate trends of the cases, encoded in internal guidelines, matter for the decision maker. This method of decision making is nonetheless very important to guarantee the consistency and fairness of legal decision making.

In previous work, it seems like cases using machine-learning models are the most similar to the local and personal *stare decisis* described by Lawlor. We have seen examples of this when aiming to model the classification of facts in terms of legal concepts, including

⁸⁸⁷ Stranieri & Zeleznikow, *supra* note 784 at 108–109.

⁸⁸⁸ “Operational instructions and guidelines — Immigration, Refugees and Citizenship Canada”, (5 December 2017), online: *Immigration, Refugees and Citizenship Canada (IRCC)* <<https://www.canada.ca/en/immigration-refugees-citizenship/corporate/publications-manuals/operational-bulletins-manuals.html>> Last Modified: 2022-10-07.

the projects by Alarie *et al*, Yin *et al* and Branting *et al* (see 4.4.3.4.2.5). There have also been examples of aiming to model judicial discretion, such as in the projects by Dahan *et al* and Stranieri *et al* (see 4.6.3.2).

Instead of using a few significant decisions to generate arguments, these projects model hundreds or thousands of “commonplace” decisions. In essence, they seek to discover how the legislation is factually applied in an area. The systems can then predict new fact patterns based upon the facts of a case. This kind of reasoning seems akin to the local and personal stare decisis discussed in this section. It also seems related to legal realism,⁸⁸⁹ as the systems aim to discover the real-world outcomes of cases in court.

In a system focusing on increasing access to justice, this approach could be very useful. Individuals may be interested in the actual outcome that they can expect should they go to court. By analyzing many cases and identifying trends in areas of low-intensity, high-volume decisions, this information could be provided to the user.

4.9.4 Conclusion

We have seen three different ways of reasoning with previous case law.

In common law systems, cases are typically reported in terms of their facts and the reasoning about these facts. Judges or lawyers that aim to reason with these cases tend to select cases that are similar in terms of their facts to a new case and argue that the same reasoning should therefore apply to the current case. Or, they might try to distinguish a case, by saying that the facts in a previous case have significant differences to a new case. Ashley refers to this kind of reasoning with cases as “Fact-Based Precedent Scenario”.⁸⁹⁰

In civil law systems, the reasoning relies much more heavily on legislation. Cases are usually seen as elaborating or explaining a statutory provision. Previous cases are typically not reported with an in-depth reasoning about the facts in a case, but rather focus on setting out new rules for how the statute should be interpreted. Cases are

⁸⁸⁹ See 1.3.2.1.

⁸⁹⁰ Ashley, *supra* note 865 at 5.

reasoned with almost like legislation. Ashley refers to this kind of reasoning as “Abstract Precedent Scenario”.⁸⁹¹

Finally, we have seen a less formalized style of reasoning with previous cases, which stems from our expectation that courts and judges will decide similar cases in similar ways. Lawlor refers to this as “local stare decisis” and “personal stare decisis”.⁸⁹² It seems like this kind of reasoning is often supported by guidelines that collect cases in an area, and thus give judges an indication on the normal outcome (such as damages or prison sentence) in certain cases. However, they are free to diverge from these guidelines if the circumstances warrant it.

All of these styles of reasoning have been explored in AI & Law research. In general, it seems like case-based reasoning systems correspond most closely to the common law style of reasoning, the rule-based systems correspond most closely to the civil law style of reasoning, and the machine-learning based systems correspond most closely to the harmonization of decision making. Of course, these distinctions have to be taken with a grain of salt – just like judges in common and civil law systems sometimes switch over to reason using the other approach of using case law,⁸⁹³ the different approaches of building legal reasoning systems can be successfully applied in jurisdictions beyond the ones attributed above.

4.10 Conclusion

In this chapter, I have explored the different steps of legal reasoning, the previous attempts to automate these steps and the challenges of doing so. This serves as an important basis for my work in building a methodology to create decisions support tools that can increase access to justice.

⁸⁹¹ *Ibid.*

⁸⁹² Lawlor, “What Computers Can Do”, *supra* note 886 at 340; Stranieri & Zeleznikow, *supra* note 784 at 105–106.

⁸⁹³ Komárek, “Reasoning with Previous Decisions”, *supra* note 829 at 158–160.

First, this chapter has given an overview of how judges reason around cases, and the different steps involved with going from a case to a decision. Building a tool that can give individuals an idea of how a judge might decide in their case likely requires the simulating of some of these reasoning steps. Understanding the reasoning process a judge performs in the different steps will be a crucial legal background in creating this system.

Further, while emulating all of these steps may be difficult in a single system, understanding which steps a computer system is able to perform, and which steps must be performed by the user, is an important clarification of the purpose of a project. For example, systems that are able to classify facts in terms of legal concepts⁸⁹⁴ can be very helpful for users but rely on the user to assess whether the system applies to their case at all, how the facts should be proven in a court, and what the consequences the legal concept being applicable is in their case. This does not decrease the usefulness of the system, since classifying facts in terms of legal criteria is an important step in legal reasoning. However, being clear about which steps are handled by the system and which are not is crucial in not misleading the users of the system.

Second, understanding the previous systems aiming to emulate the different steps of legal reasoning provides me with a base of knowledge in how to construct my own methodology to build legal decision support tools. Understanding the tradeoffs of prior work is a crucial step in deciding which of these technologies and approaches I will make use of for my own system.

Third, exploring the different steps of legal reasoning has given me an appreciation for which steps in legal reasoning may be beyond the current state of artificial intelligence and symbolic systems. Some of the steps described above may be AI-complete, requiring “general artificial intelligence”, which is currently beyond the state of the art. This seems to be the case for many areas where there is an “unbounded” input space, i.e. the judge can consider any factor they see fit in making a decision. Examples of such unbounded

⁸⁹⁴ See 4.4

spaces include reasoning with evidence, classifying facts as legal concepts when there is no set list of factors to consider, determining the correct interpretation of the structure of legal rules and deciding on a discretionary outcome when there is not list of factors to consider. Systems that aim to automatically accomplish these steps may be difficult to build. Instead, building systems that give the user the tools to make the decisions for themselves may be more practical.

In other areas, the problem space is bounded, i.e. there is a set number of possible incomes and outputs. Here, building systems that provide meaningful statistics or predictions seems possible. Such areas include the assessment of the legal concepts that apply based on facts where there is a set list of factors to consider, the emulating of the logical application of legal rules when the intended syntactic application is clear, and the discretionary decision of an outcome when there is a set list of factors to consider. In these cases, it should be possible to build powerful models of an area. However, even here, judges may diverge in individual cases, which may be difficult to model.

Fourth, this chapter has highlighted the importance of providing explanations in the legal field. This is a legal requirement in many cases, and a crucial feature in building systems that are able to provide information in a useful and legitimate way. Some of the systems, such as case-based reasoning systems, prioritize delivering a useful explanation over giving an accurate prediction. Machine learning systems face some shortcomings in this area, as they often rely on models that are too complex to understand. The explanations delivered by such systems may not be explanations in a legal sense.

Fifth, in 4.8 and 4.9, I explored some important characteristics of interactions between the legislator and the courts, and between courts. From this analysis, it seems like it is important to model both legislation and case law, in order to understand the reasoning performed in certain legal areas. There are different ways of reasoning with case law. Some ways to reason with decisions, such as the reasoning in common and civil law systems, rely on the qualitative understanding of cases, either as seen through the facts present in a case or through the rules clarified by a case. Another way of reasoning with cases relies on the quantitative alignment of reasoning and outcome in commonplace

cases. It is important to understand these distinct styles of reasoning with legal cases, as the way that is chosen will influence how cases are represented and reasoned with.

With these insights in mind, let us move on to explore how to build a methodology to create legal decision support tools to improve access to justice. In the next chapter, I will discuss some constraints and choices that need to be considered in building such a system. This will inform the concrete implementation of the methodology that I will create.

Chapter 5 Design criteria

Research Objective: Determining the relevant design criteria (1.2.2.4)

Research Topics:

- Which design criteria should guide the design of my methodology using AI to increase access to justice and access to legal information?

5.1 Introduction

We have now investigated artificial intelligence and access to justice and analyzed how judges tend to deal with legal cases. As we have seen, the reasoning process is complex and involves multiple important steps with different styles of reasoning. We have also explored how previous research has aimed to automate these steps, and the challenges in doing so.

In this chapter, I will explore some crucial design criteria that need to be decided upon before developing a legal decision support methodology. It is crucial to determine certain criteria, such as the desired target user and the purpose of a system, before beginning development, as such choices can have a significant effect on the concrete implementation of a legal decision support tool.⁸⁹⁵

The overarching goal for my methodology is to increase access to justice (see Chapter 3). In order to achieve this target, I will focus on designing a system that fulfills the following four design criteria:

- The system should target laypeople
- The system should be able to handle areas of high-volume, low-intensity⁸⁹⁶ legal problems
- The system should give specific and useful insights

⁸⁹⁵ Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563 at 254.

⁸⁹⁶ See 3.2.1.

- The system should be practical, i.e. ready to use in real-world scenarios

Each of these design criteria has a number of implications for the concrete development of my methodology to develop legal decision support tools. I will explore the reasons for choosing these criteria and their implications in-depth in 5.2 - 5.5, in light of the capabilities of AI presented in Chapter 2, the needs discussed in Chapter 3 and the previous approaches examined in Chapter 4.

Each of these implications hold a part of the answer of my research goal (see 1.2). In 5.6, I will summarize these implications and explore the concrete impact they have on the design of legal decision support tools (see 1.2.2.5), namely:

- The coding of legislation and court cases
- The input of user data
- The analysis of user data to obtain useful information
- The presentation of information to the user

The purpose of this chapter is not to concretely detail the implementation of my methodology to build decision support tools. This description will follow in Chapter 6 and Chapter 7. Rather, the purpose is to illuminate the design criteria, following from my goal and the analysis in Chapter 4, that shape the development of a methodology able to support access to justice.

5.2 The system should target laypeople

A first important consideration to take regarding building a system emulating legal reasoning is the target audience of the platform. Different target users of such systems differ in their levels of legal knowledge and understanding. This imposes different constraints on the design of legal decision support tools. In prior work described in Chapter 4, we saw a multitude of such purposes.

Paquin *et al* argue that specifying a target user of the system at the beginning of building the system is an important prerequisite, since different target users present different challenges in building the system.⁸⁹⁷

Let us explore some of the possible target users of a legal decision support system, and the possible use cases for such systems. I will then elaborate on my choice of target user (namely the layperson) and explain some particularities of building decision support tools for this user.

5.2.1 Different target users and purposes

- **The public** – Some research in emulating legal reasoning has targeted the public.⁸⁹⁸ In this research, the purpose was e.g. to examine trends in legal decision making, by analyzing the factors (including legal factors such as the merits of a case, and extra-legal factors such as the time of day or ideological leanings of the judge) that may affect outcomes in court.
- **The judiciary** – Legal reasoning systems could also target the judiciary. Some have argued that it could be used to partially automate certain tasks, such as the triage of cases, or by generating decision suggestions for judges. Researchers believe that systems like this could increase the efficiency of the judicial system,⁸⁹⁹ and potentially eliminate sources of bias.⁹⁰⁰
- **Legal professionals** – Legal reasoning systems could also target legal professionals. Here, the purpose would be to support the generation of stronger

⁸⁹⁷ Paquin, Blanchard & Thomasset, “Loge–expert”, *supra* note 563 at 254.

⁸⁹⁸ Daniel L Chen & Jess Eigel, “Can machine learning help predict the outcome of asylum adjudications?” (2017) Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law 237–240.

⁸⁹⁹ Daniel Becker & Isabela Ferrari, “VICTOR, the Brazilian Supreme Court’s Artificial Intelligence: a beauty or a beast?” (2020) SIFoCC Standing International Forum of Commercial Courts, online: <<https://sifocc.org/app/uploads/2020/06/Victor-Beauty-or-the-Beast.pdf>>; Eric Niiler, “Can AI Be a Fair Judge in Court? Estonia Thinks So”, (25 March 2019), online: *Wired* <<https://www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so/>>; Tara Vasdani, “Robot justice: China’s use of Internet courts”, (5 February 2020), online: *The Lawyer’s Daily* <<https://www.thelawyersdaily.ca/articles/17741/robot-justice-china-s-use-of-internet-courts>>.

⁹⁰⁰ Alarie, Niblett & Yoon, *supra* note 682 at 5; see however Angwin et al, *supra* note 276; Starr, *supra* note 537; Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 803–804.

arguments, such as by surfacing relevant case law and suggesting possible arguments that the lawyer might use, potentially even with predictions of the likelihood of success of such arguments.⁹⁰¹ Such systems could also be used to support the lawyer in accomplishing certain tasks more efficiently.⁹⁰²

- **Law students** – Legal reasoning systems could also target law students.⁹⁰³ Law students could examine such a system to learn patterns of legal reasoning with case law and rules.

5.2.2 Laypeople

While the previously discussed applications of artificial intelligence to support actors in the legal system are very interesting, in this work I have decided to focus primarily on systems targeting laypeople, i.e. individuals without legal training. As discussed in Chapter 3, individuals today face significant issues with access to justice. Potentially, building legal decision support tools could provide these individuals with important information that is able to increase their access to justice. Note that targeting the methodology at laypeople does not preclude other target users from benefiting from a legal decision support tool. I will discuss how my methodology could support other target users in 9.2.

Let us explore some attributes of laypeople, and the implications these have for my research.

5.2.3 Laypeople think of their situations in terms of facts, or goals they want to achieve

Targeting laypeople has a number of implications for the user of the system. One implication is that the user is often likely to think about their case in terms of factual occurrences, rather than in terms of legal issues. As we saw in 4.2, identifying the legal

⁹⁰¹ Alarie, Niblett & Yoon, *supra* note 682.

⁹⁰² Zeleznikow, *supra* note 67 at 25; Richard Susskind, “Pragmatism and purism in artificial intelligence and legal reasoning” (1989) 3:1 AI & Soc 28–38 at 35.

⁹⁰³ Aleven, “Using background knowledge in case-based legal reasoning”, *supra* note 47.

rules that may be applicable to a factual situation is a crucial step in legal reasoning.⁹⁰⁴ However, laypeople may not be able to understand the legal aspects of their situation. Instead, they may think of their situation purely in terms of the factual occurrences. A study of self-represented litigants in Canada in 2013 found that many self-represented litigants had difficulty determining which form they should fill out.⁹⁰⁵ Likewise, Branting found that one of the biggest challenges faced by laypeople is their inability to know which facts they need to establish, and organize their facts in the appropriate legal manner.⁹⁰⁶

Let us assume that a user is facing an issue with heating in their apartment. Asking the user whether they would want a rent reduction, for example, may not be helpful, since the user does not know that a rent reduction is a possible consequence to issues with heating. Rather, the user has to be given an entry point that allows them to explore possible legal avenues, based on a factual occurrence. In the world of expert systems, this style of reasoning is referred to as forward reasoning, since it starts with the input of the user and then reasons “forwards” towards the possible goals based on these facts.⁹⁰⁷

On the other hand, in some cases, users of systems may have a particular goal in mind and wish to explore the possible avenues of arriving at this goal. For example, a user may wish to leave their apartment before the end of their lease. In this case, backwards-chaining, where an expert system starts with a goal and then explores the possible avenues of reaching this goal, could be more useful.⁹⁰⁸

Ideally, the developed methodology would therefore support both the forward and backward styles of reasoning. In prior work, systems such as CABARET are able to

⁹⁰⁴ Compare Thompson, *supra* note 75 at 22–23.

⁹⁰⁵ Macfarlane, *supra* note 323 at 59–60.

⁹⁰⁶ Branting et al, *supra* note 348 at 218.

⁹⁰⁷ See 2.5.1.

⁹⁰⁸ See 2.5.1.

support both styles of reasoning.⁹⁰⁹ We also saw systems such as Spot able to identify possible legal issues in plaintext descriptions of factual situations.⁹¹⁰

5.2.4 Laypeople may not understand legal concepts

The second implication is that the user cannot be expected to understand legal concepts, such as open-textured terms in the legislation. In 4.4, we discussed how judges aim to determine applicable legal concepts from the facts of a case, and how case-based systems can be used to emulate this style of reasoning. In 4.5.3.3.3, we explored how designers of rule-based expert systems dealt with this issue of open-textured concepts. Solutions include asking the user for whether a specific legal concept is fulfilled or not, optionally with the inclusion of examples from case-law or other explanations. Other approaches include integrating case-based reasoning into a rule-based system in order to model the classification of facts into legal concepts.

In building legal decision support tools for laypeople, it is crucial to be aware that we cannot ask them directly for whether certain open-textured concepts are fulfilled. Instead, we need to give them the context and explanations necessary to perform this assessment (see the “communication layers” used in Loge-expert)⁹¹¹ or build a system that can automatically infer fulfilled legal concepts from the facts of a situation.

5.2.5 Laypeople have trouble evaluating the quality of provided information

The third implication is that the user may not understand some fundamental aspects of the legal system, and artificial intelligence. There are many particularities of the law that are not obvious to laypeople and may lead to confusion or even harm if they interact with legal decision support tools. For example, laypeople may not be aware that every case is unique, and that a judge may diverge from previous case law in certain situations.

⁹⁰⁹ See 4.5.3.2.6.

⁹¹⁰ Colarusso, *supra* note 560.

⁹¹¹ Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563 at 256.

Likewise, individuals have a propensity to overestimate the capability of artificial intelligence systems.⁹¹² In 2.6.3, I described the limits of artificial intelligence, including a lack of common-sense reasoning. As we have seen in Chapter 4, this means that most AI & Law predictive systems are not fully able to predict the complex reasoning process of the judge at the different steps. For example, some systems worked well for standard cases but had difficulties dealing with exceptional cases.⁹¹³

Conveying this nuance to laypeople is an important step in building legal decision support tools. For example, when provided with an estimation that their case is 70% likely to succeed, or conversely 80% likely to fail, laypeople might base their entire decision on this estimation, rather than understanding the flaws and shortcomings of the selected AI methodology, or the exceptional nature of their case.

Providing laypeople with a probability of success, or even a yes/no prediction of whether they may win their case, may therefore be less useful.⁹¹⁴ It is very important to find ways to convey the ability and potential shortcomings of a legal decision support tool to the user. Ideally, legal experts should be able to understand and vet the potential information a layperson could be given by the system, to verify that it is accurate and not misleading.⁹¹⁵

5.2.6 Laypeople may not understand legal language and complex texts

Among laypeople, there may be different levels of reading comprehension. While lawyers can be assumed to be able to read legal statutes and decisions, the same cannot be said for laypeople. The Self-Represented Litigants project conducted a survey in 2013 in Canada, finding that almost every self-represented litigant had trouble understanding

⁹¹² Surden, “The ethics of artificial intelligence in law”, *supra* note 541 at 723.

⁹¹³ See 4.5.3.3, 4.4.3.4.3, 4.3.3.3, 4.6.3

⁹¹⁴ Compare Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 3.

⁹¹⁵ Thompson, *supra* note 75 at 28.

the language used in court forms.⁹¹⁶ Using plain and clear language has been found to be a crucial aspect in increasing access to justice of the population.⁹¹⁷

In order to reach the greatest possible audience, a system targeting laypeople should therefore attempt to minimize the amount of text a user has to read in order to understand the information given by the system.⁹¹⁸ Further, the vocabulary should be as simple as possible. If possible, legal information should be condensed to the useful parts, and summarized in plaintext language.

At the same time, it is important not to let perfect be the enemy of good. As will be further explored in 5.5, the goal of my methodology is to be useful in a practical sense. For example, initially targeting individuals that have a certain level of reading comprehension could therefore be a viable first step for building such a methodology. The approach could be expanded with additional accessibility features in future work, including, for example, the evaluation of how well the system can be used by people with disabilities.⁹¹⁹

5.2.7 Laypeople may have constraints in terms of technological access

Another aspect to consider is access to technology by laypeople. According to Statista, as of January 2022, 96.5% of the Canadian population were internet users.⁹²⁰ Globally, almost 5 billion individuals have access to the internet.⁹²¹ This suggests that the internet is a viable medium to reach a large amount of people with legal decision support tools.⁹²²

⁹¹⁶ Macfarlane, *supra* note 323 at 60.

⁹¹⁷ Hon Bridget Mary McCormack, “Access to Justice Requires Plain Language” (2021) 100:2 Michigan Bar Journal 44–46; “Plain language – essential for real access to justice”, (18 July 2017), online: *Provincial Court of British Columbia* <<https://www.provincialcourt.bc.ca/enews/enews-18-07-2017>>; “New guide for plain language in law”, (31 January 2019), online: *The Access to Justice Lab* <<https://a2jlab.org/new-guide-for-plain-language-in-law/>>.

⁹¹⁸ Cabral et al, *supra* note 433 at 275.

⁹¹⁹ *Ibid* at 262–263.

⁹²⁰ “Number of internet users in Canada 2022”, online: *Statista* <<https://www.statista.com/statistics/243808/number-of-internet-users-in-canada/>>.

⁹²¹ Simon Kemp, “Digital 2022: Global Overview Report”, (26 January 2022), online: *DataReportal – Global Digital Insights* <<https://datareportal.com/reports/digital-2022-global-overview-report>>.

⁹²² Karim Benyekhlef et al, “ICT-Driven Strategies for Reforming Access to Justice Mechanisms in Developing Countries” (2015) 325 *The World Bank Legal Review* at 342–343.

Zelevnikow suggested using a web-based interface in 2002, but cautioned against technical challenges, such as difficulties identifying rules in a large knowledge base, and difficulties scaling the system.⁹²³ These challenges may be less pronounced today, with the incredible advances in computing power.

While computers are pervasive, phones may be even more widespread. In 2017, 89.5% of households owned mobile phones, while 84.1% of households had owned home computers in Canada. In the first income quintile, defined as earning less than 32,914 CAD per household, 73.1% of households owned mobile phones, while 63.4% owned home computers.⁹²⁴

This shows that not all households may have access to smartphones and an internet connection. Low-income individuals, that have the greatest need for improved access to justice, may also have the lowest level of access to technology. However, just like in the previous section, the adage of “perfect is the enemy of good” applies in this case. Even if not all lower-income households can be reached by a system targeting smartphone users, a significant portion of people can. Beyond this, deploying the developed legal decision support tools to legal aid clinics or libraries may be able to further increase the penetration of such a computer tool.⁹²⁵

I therefore target users with access to a computer, tablet or smartphone for my legal decision support tool. This requires the website to be optimized for both desktop and mobile use.⁹²⁶ Chavan suggests also optimizing websites for quick delivery, and touch-based interactions.⁹²⁷

⁹²³ Zelevnikow, *supra* note 67 at 19.

⁹²⁴ *Communications Monitoring Report 2019*, Reports, by Canadian Radio-television and Telecommunications Commission (CRTC) Government of Canada, Reports BC9-9E-PDF (The Canadian Radio-television and Telecommunications Commission (CRTC), 2020) Last Modified: 2020-01-21.

⁹²⁵ Cabral et al, *supra* note 433 at 247.

⁹²⁶ *Ibid* at 270–271.

⁹²⁷ *Ibid* at 277–278.

5.3 The system should be able to handle areas of high-volume, low-intensity legal problems

Now that we have discussed the type of target users, let us explore the type of legal areas that the methodology should primarily target. In order to be able to increase access to justice for laypeople, I have chosen to initially focus my research in areas of high-volume, low-intensity cases,⁹²⁸ such as consumer protection issues, debt problems and housing disputes. Such areas are frequently litigated in court. Further, the complexity of the legal issues and the amount of damages claimed tends to be relatively low.

The reasons for this choice of type of domain are two-fold. First, building legal decision support tools in such areas can have a significant impact on access to justice. Second, building legal decision support tools in such areas may be more feasible. I will explore these two implications below.

5.3.1 Impact on access to justice

Focusing on areas of high-volume, low-intensity cases is important in order to have the greatest possible impact on access to justice. Here, I will explore why this is the case.

5.3.1.1 Tools can help many individuals

The existence of a large number of cases is a strong indication that many people tend to be faced with such legal issues. This increases the potential impact an artificial intelligence system can have on the public. While building a system for cases that are only heard a few times per year could be useful for parties of those cases, it would not have the same impact on access to justice as cases that are heard thousands of times per year. Thus, building legal decision support tools in such areas holds the promise of having a broad impact on society, as many individuals may benefit from the information provided by the system.

⁹²⁸ Benyekhlef & Vermeys, *supra* note 305.

5.3.1.2 Cases more likely to affect laypeople

Further, high-volume, low-intensity legal areas are more likely to affect laypeople. As we saw above, a significant portion of individuals reported being affected by such legal issues, such as consumer, debt, employment and housing issues.⁹²⁹ Laypeople are less likely to be involved in complex legal cases regarding i.e. intellectual property, corporate law or constitutional law. As we saw above, laypeople often face issues in knowing how to deal with everyday legal issues.⁹³⁰ Building legal decision support tools in areas of high-volume, low-intensity legal disputes could be an important step to support such users.

5.3.1.3 More likely to be carried out by self-represented litigants

Further, low-intensity cases are more likely to be carried out by self-represented litigants. Since the monetary values are relatively low, hiring lawyers at hundreds of dollars per hour may eclipse the value of the dispute.⁹³¹ Further, many individuals that are affected by low-intensity legal disputes may not have the resources to hire a lawyer.⁹³² As we have seen, the court experience for self-represented litigants is not ideal.⁹³³

In Chapter 4, we saw that many components of legal reasoning are very complex and may be beyond the current scope of AI systems. Lawyers are able to go much further than AI systems, by recommending next steps to their clients, drafting documents and representing them in court.

Using a legal decision support tool may inform the user that they have a legal right, and thus make them aware that they might want to hire a lawyer in the first place, or seek another way of resolving their issue. Further, in situations where individuals are unable to hire a lawyer, e.g. for financial reasons, the legal decision support tools could be an

⁹²⁹ See 3.2.1.

⁹³⁰ See 3.2.2.

⁹³¹ Benyekhlef & Vermeys, *supra* note 305.

⁹³² See 3.2.4.2.

⁹³³ See 3.2.4.3.

important step in helping the user understand their rights, which could support them in preparing their case when going to court.

5.3.1.4 Decrease load on court system

Finally, building AI systems in areas where there is a significant number of cases would also potentially support the court system. If the court system is overloaded with cases, individuals may not be able to get their dispute resolved in a timely matter, limiting their access to justice. As we have seen, self-represented litigants frequently place an additional heavy burden on the court system.⁹³⁴

Legal decision support tools could improve this situation in multiple ways. First of all, individuals may recognize that their case is not viable, or does not correspond to their expectations in terms of monetary reward.⁹³⁵ These cases may thus not be introduced in court, decreasing the number of cases that a court has to deal with. Of course, there are also risks with this situation, as described above in 3.6.2.1 and 3.6.2.3.

Second, the information provided by a system may allow the user to settle their case on their own. They may talk with the other party, in context on the information provided and possible outcomes, and find an amicable solution.⁹³⁶ These cases could thus be dealt with outside of the court system, allowing the court system to focus on more complex cases.

Third, the individuals that do go through court may do so with more information, making the process go more smoothly. The tool may guide such people towards the right forms, and inform them of their rights and which facts they need to prove. Such information may decrease the load on the court system, as the court officers and justice are able to more smoothly deal with individuals that are better informed.⁹³⁷

⁹³⁴ See 3.2.4.4.

⁹³⁵ See 3.6.1.2.

⁹³⁶ See 3.6.1.3.

⁹³⁷ See 3.6.1.4.

Overall, building decision support tools in areas of high-volume, low-intensity cases could thus have an important impact on decreasing the caseload of the court system, and giving judges more time to focus on more complex cases.⁹³⁸

5.3.2 Feasibility of building legal decision support tools

As we have seen, focusing on areas of high-volume, low-intensity cases will likely have the highest impact on access to justice. However, this is not the only reason for choosing such domains. Areas of high-volume, low-intensity cases may further be important to increase the feasibility of building systems that can treat case law using artificial intelligence. Let us explore why this is the case.

5.3.2.1 Requirements for a lot of data in machine learning

The volume of cases is important for multiple reasons. First of all, as we have seen above, the use of machine learning requires large datasets.⁹³⁹ If there are only a few cases available, the artificial intelligence method may not be able to extract the relevant signal from the data, leading to poor performance. This is, of course, dependent on the specific methodology used and the way that cases are used and encoded. However, in general, approaches that use some kind of statistical learning of patterns likely benefit from the availability of a lot of case law, which exists in areas of high-volume, low-intensity cases.⁹⁴⁰

Further, the amount of *accessible* data is important for the training of an AI algorithm. In order to develop AI methodologies, cases have to be accessible in a bulk format. As we have seen, having public access to legal information is not a given in all jurisdictions.⁹⁴¹ Even in jurisdictions where case law is accessible via the internet, the downloading of bulk data, which is required to train AI models, is often not supported.

⁹³⁸ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 796.

⁹³⁹ See 2.6.3.1.

⁹⁴⁰ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 796.

⁹⁴¹ See 3.4.

5.3.2.2 Less complex legal reasoning

Selecting an area with low intensity of cases is further beneficial due to the legal reasoning being less intricate.

Cases that are complex and raise complicated legal questions are likely to be much more difficult for an AI system to analyze and present. In fact, some of the reasoning that a judge carries out when dealing with a complex constitutional law question, including reasoning related to policy arguments and purposive interpretations of the law, may require a so-called “strong” AI system, as described above. For commencing the development of a methodology to use AI to increase access to justice, such a domain may therefore be tricky to deal with.⁹⁴²

In cases dealing with simple legal issues, such as debt, employment, neighbor and debt problems, the focus of a court procedure is likely to be on the parties providing evidence, or the judge determining whether a certain situation reaches up to a certain standard set out in the law. The monetary amounts in question may be comparatively low, and the individual is likely to be self-represented. Different cases may share large similarities, and deal with similar topics but in different contexts. These properties make it more feasible and practical to create legal-decision support tools in these areas. Benyekhlef and Zhu agree that the use of AI in such areas of low intensity is likely much more feasible than in areas that require subtle and delicate legal reasoning, incorporating socio-economical contexts, such as the area of human rights.⁹⁴³

5.3.2.3 Legal pathways “explored” by courts

If there are a lot of cases in an area, it means that many of the legal pathways may already be explored and clarified by previous cases. In 4.5.3.3.4, I explored the difficulty of dealing with syntactic ambiguity in legal rules without AI-complete systems. If there are many cases in a legal area, it is more likely that a given legal pathway has already been tried by a court, clarifying the syntactic structure of a rule. Thus, the creator of the AI

⁹⁴² Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797.

⁹⁴³ *Ibid.*

system does not have to speculate about how the syntactic components of a rule may be interpreted – they can instead observe how judges practically interpret a rule and encode the rule into the computer system. Rather than encoding the legislation itself, we may here encode the way the legislation was concretely applied in court.

5.3.2.4 Cases more likely to form clusters

Having a large amount of court cases means that it is more likely that a case that is similar to that of the user has already arisen.⁹⁴⁴ If there are few cases, every case may be unique in some way. Computer systems struggle with handling new situations, such as cases involving previously unseen factors.⁹⁴⁵ If there are larger amounts of cases, clusters of similar cases may emerge, that can allow the computer to build a more accurate model of cases involving the same issues. The emergence of similar cases further allows the system to surface cases that are similar to that of the user, which can help them in understanding how their case might be treated.

Targeting areas of high-volume, low-intensity cases has a number of implications for how cases should be encoded, which will be explored below in section 5.4.5.

We have now seen which target user (laypeople) and which type of legal areas (areas with high-volume, low-intensity cases) the methodology should target. Next, let us explore the kind of information the system should aim to provide to the user.

5.4 The system should give specific and useful information

The ambition of my methodology is to be able to give specific and useful information to laypeople. I will briefly elaborate on what is meant by these terms, and then explore the implications of this constraint on the developed methodology.

⁹⁴⁴ note 485 at 129.

⁹⁴⁵ See 4.4.3.4.3.2 and 4.6.3.3

- **Specific** - There exist several legal information websites that aim to give access to legal information in simplified language.⁹⁴⁶ These tools can be very useful. However, for my methodology I would like the information to be specific, i.e. respond to the situation of the user to provide specific information. Providing the user with specific information also makes it likely that the information is shorter, which is an important criterion I discussed in 5.2.6.
- **Useful** - Further, I want the information provided to the user to be useful. The user should receive information that can be used to achieve beneficial outcomes and support them in making decisions related to how to deal with their case.

Next, I will discuss the implications of aiming to build a system that should provide specific and useful information regarding everyday legal issues to the layperson user.

5.4.1 The system should inform users about their rights and possible outcomes

Let us explore the type of information that should be provided to the user, that may fulfill the requirements of being specific and useful. I will focus on providing the user with information about their rights, and which outcomes cases such as theirs generally lead to in court.

5.4.1.1 Allowing the user to discover their rights

One way a legal reasoning tool could be used to support individuals is by helping them understand their legal rights.⁹⁴⁷ Since the law can be difficult to understand and vague, laypeople may be unaware of the fact that they have certain rights when faced with certain problems.⁹⁴⁸ As we saw in 3.2, not knowing that a situation has legal characteristics, or being uncertain about legal rights in a certain situation, is one of the main reasons preventing individuals from addressing their legal issues.⁹⁴⁹ Similarly, not knowing which forms to use and understanding such forms is one of the significant

⁹⁴⁶ Cabral et al, *supra* note 433 at 248; E.g. “Éducaloi - Your starting point for legal information”, online: *Éducaloi* <<https://educaloi.qc.ca/en/>>.

⁹⁴⁷ Cabral et al, *supra* note 433 at 247.

⁹⁴⁸ See 3.4.

⁹⁴⁹ Currie, *supra* note 297 at 56.

barriers faced by self-representing litigants.⁹⁵⁰ A system that is able to triage the problem of the user and describe the possible legal rights and avenues available to them could thus bring significant gains in access to justice.⁹⁵¹

The system should function in the following manner. The user would indicate a factual situation or a goal to the system. In response, the system would need to determine the key guiding legal issue relevant to the user, and elicit the relevant information from the user, without requiring the user to understand legal concepts (see 5.2.4). Then, the system would need to understand whether the answers provided by the user indicate that certain legal criteria may be fulfilled, and the consequences thereof.⁹⁵² The system could then provide them with information about which rights may apply to them, and the avenues available to them to enforce their rights.

This information would be *specific* to the case of the user, corresponding to their individual situation. It would also be *useful* in helping the user understand the legal rights, which could lead them to resolve their issues in ways they were not previously aware of. For example, they may use the information to settle their case with the other party. Further, the information provided by the system on possible next steps could inform the user of the avenues of obtaining additional help (such as hiring a lawyer), or how to enforce their rights.

Likewise, such a tool could be integrated into an ODR platform.⁹⁵³ Benyekhlef and Zhu suggest using artificial intelligence as a first step for ODR platforms, by informing the user of the validity of their case, the competency of the ODR platform and possible outcomes of the conflict.⁹⁵⁴ It could thus serve as an entryway, directing the user toward the correct type of claim in such a platform. The use of the methodology is not a main focus of this thesis, but will be explored below as future work, see 9.4.4.2.

⁹⁵⁰ See 3.2.4.3.

⁹⁵¹ Cabral et al, *supra* note 433 at 294–295.

⁹⁵² See 4.2, 4.4 and 4.5.

⁹⁵³ See 3.5.2.4.1.

⁹⁵⁴ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797.

5.4.1.2 Informing users of possible outcomes

Beyond informing the user of the existence of a right, the system may provide the user with an estimation or a range of possible outcomes for their case, should it go to court. For example, if a tenant is told that a certain specific issue they face generally results in damages of 200 CAD – 400 CAD being ordered by the judge, they can incorporate this information in deciding on whether they should pursue the case in court. They can weigh the value of the possible court decision versus costs of appearing in court, which may be higher than the value of the dispute itself.⁹⁵⁵ In terms of legal reasoning, providing this information is akin to modelling the outcome of a case, described in 4.6.

Providing this kind of information is *specific*, as it should relate to the specific situation of the user. This requires the analysis of why two cases are similar, and why certain cases are not, which is one of the key research questions of building the methodology.⁹⁵⁶

I further believe this information to be very *useful* to the user, by providing important context to their decision making. As described in 3.5.2.2, this information can be used as a BATNA, in order to align the expectations of the users, and encourage them to settle.

A concrete example will show why this kind of information could be useful for the parties. If one of the parties expects a case to lead to damages of 5,000 CAD, while the other party is only willing to pay 200 CAD, the case may be difficult to settle. However, if the parties can see that similar cases that go to court generally end up with one of the parties paying between 400 CAD and 600 CAD, they may be able to adjust their expectations accordingly and come to an agreement without going through the arduous court procedure. This could be highly useful information for the user, as it would allow them to benefit from the advantages of using alternative dispute resolution, such as lower costs, and quicker and more amicable solutions, and also reduce delays in the justice

⁹⁵⁵ Benyekhlef & Vermeys, *supra* note 305.

⁹⁵⁶ See 1.2.2.5.

system (see 3.6.1.7).⁹⁵⁷ A vast majority of individuals seem to prefer this kind of negotiated solutions over using the court system.⁹⁵⁸

Even if the user is not able to settle their case based on the overview of previous outcomes, showing previous outcomes may be useful by providing important context to the decision of the user on how to further pursue the case. For example, the user may decide that the effort required to enforce their case in court is not worth the likely compensation.⁹⁵⁹

The display of a BATNA could also be very effective when integrated into an ODR platform (see 3.5.2.3). This would allow parties to conduct their negotiations in light of the statistics provided about previous cases, potentially increasing the possibility of amicable settlements.⁹⁶⁰ The system could also continuously collect the data from settlements reached inside the system, and present these to users of future similar cases.⁹⁶¹ This kind of data would potentially be more accurate and relevant for the users of the platform, since it is collected in the same context.⁹⁶² The use of AI in ODR has been referred to as ODRAI,⁹⁶³ and will be explored as future work in 9.4.4.2.2.

There have been a number of systems that aim to contribute to settlement through tools such as supporting the communication of the user or using game theory approaches to get to a solution.⁹⁶⁴ These are beyond the scope of this thesis.

⁹⁵⁷ Benyekhlef et al, *supra* note 322 at 10.

⁹⁵⁸ Benyekhlef et al, *supra* note 922.

⁹⁵⁹ See 3.6.1.2.

⁹⁶⁰ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797; Zeleznikow, *supra* note 330 at 39; Carneiro et al, “Online dispute resolution”, *supra* note 385 at 21.

⁹⁶¹ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 796–797.

⁹⁶² Compare 3.6.2.6.

⁹⁶³ note 485 at 134.

⁹⁶⁴ Zeleznikow, *supra* note 67 at 41; Emilia Bellucci, Arno R Lodder & John Zeleznikow, “Integrating artificial intelligence, argumentation and game theory to develop an online dispute resolution environment” (2004) 16th IEEE International Conference on Tools with Artificial Intelligence 749–754; Lodder & Thiessen, *supra* note 501; Cabral et al, *supra* note 433 at 291.

5.4.1.3 Filling out forms and generating arguments?

Another pertinent use of legal decision support tools could be related to filling out forms and assembly of documents.⁹⁶⁵ There exist multiple software solutions that allow the creation of such systems, such as DocAssemble⁹⁶⁶ and A2J Author.⁹⁶⁷ Such an application will be explored in 9.4.4.1.3.

Further, I will not explore the generation of arguments for laypeople to use in court. While this could also be very important and helpful, it may raise issues with regards to the distinction between legal advice and legal information, see 5.5.6. However, I will explore the use of my methodology to generate legal arguments in future work, see 9.4.4.1.1.

5.4.2 Which steps of legal reasoning should be covered?

In Chapter 4, we went through the different steps of legal reasoning. As discussed in 4.10, understanding which of these steps of legal reasoning are covered is an important requisite to understanding how to build a legal decision support tool. Let us examine which steps will be covered by my created methodology:

- 1. Identification of the guiding legal rule (4.2)** – Identifying the guiding legal rule that an individual can rely on in obtaining a certain remedy is a crucial prerequisite to the individual understanding their rights. This step thus has to be covered by the legal decision support tool.
- 2. Finding of facts (4.3)** – The finding of facts relies on reasoning with evidence, which is an important step. However, as discussed in 4.3.3.3, emulating this kind of reasoning is difficult. In this work, I therefore do not explore the assessment of evidence. The results of the system are therefore conditional on the individual being able to prove their facts to the judge. It is very important to make this clear

⁹⁶⁵ Cabral et al, *supra* note 433 at 251–252.

⁹⁶⁶ note 456.

⁹⁶⁷ note 455.

to the user. I will describe possible applications of my methodology for structuring and collecting evidence in 9.4.4.1.1.

3. **Establish fulfilled legal criteria based on proved facts (4.4)** – Establishing the fulfilled legal criteria based on facts is an important step in legal reasoning and is typically opaque to laypeople.⁹⁶⁸ Emulating this process is thus an important part of providing legal decision support.
4. **Applying legal rules to legal criteria (4.5)** – Understanding the structure of the law can be complicated for laypeople. Supporting the user in understanding how the rules may apply to their situation is therefore an important step in providing an individual with information about their rights and avenues.⁹⁶⁹
5. **Determining the outcome of a case (4.6)** – Likewise, providing the user with information about the discretionary outcome of their case is important for providing the user with a BATNA.⁹⁷⁰
6. **Explaining the decision (4.7)** – Explaining the decision, whether through providing the rules that were followed or through showing similar cases, is another important step in providing useful legal information to the user.

It is important to note that these steps will not necessarily be “automated”. They are merely areas that are included in the methodology I suggest, whether through providing information related to a specific reasoning step to a user or analyzing data about a specific step using machine learning.

5.4.3 The system will rely on legal rules and court decisions

Now that we have explored what kind of information we would like to supply to the user, let us examine which legal source material could be analyzed to generate the information.

⁹⁶⁸ See 5.2.4.

⁹⁶⁹ See 5.4.1.1.

⁹⁷⁰ See 5.4.1.2.

This determination relates to the purpose of the system. However, since one of the goals of the methodology is to be practical, the data should also be available for the creation of such systems.

Possible documents that could be analyzed in order to build legal reasoning systems include:

- The legislation – The legal rules governing a certain area.
- Case law – a written summary of a previous case, how the judge reasoned about this case and their conclusion.
- Party submissions – the submissions of the parties to the court.
- Court dockets – These contain the procedural history of a case.
- Evidence – the evidence used by a party to argue a case.
- Legal doctrine – academic research in an area, explaining the functioning and details of a legal area.

Not all of this information is generally available. For example, party submissions, while perhaps corresponding most closely to the perspectives of laypeople, are generally difficult to access.

In my research, I focus on the legal sources of legislation and case law. Both of these are usually public, although there may be privacy issues limiting jurisdictions from providing access to decisions in bulk.⁹⁷¹ Further, as described in 4.8, both legislation and case law are crucial in understanding how the law in an area is practically applied. The legislation in a legal area is needed to set out the structure of the rules applicable to a situation, while case law is needed to understand how the open-textured terms of the legal rule are applied in concrete cases, and the discretionary decisions judges take about the outcomes of cases.

⁹⁷¹ Compare Nicolas Vermeys, “Privacy v. Transparency: How Remote Access to Court Records Forces Us to Re-examine Our Fundamental Values” in Karim Benyekhlef et al, eds, *eAccess to Justice* (University of Ottawa Press, 2016) 123.

In order to generate explanations of legal terms and outcomes, my methodology will also be able to incorporate plaintext summaries and explanations of legal areas, if available.

5.4.4 Legal rules will be represented in a rule-based system

To allow the user to explore their rights, my methodology thus needs to be able to encode legal rules. This system needs to be able to emulate the reasoning by a judge performed in the step described in 4.5, i.e. applying legal rules to legal criteria. As such, it needs a way to represent the logical connection between different legal concepts and outcomes described in the law.

In 4.5.3, I explored the automation of systems reasoning about legal rules. The systems generally rely on logical representations and connections (such as AND/OR) to represent legislation. Some of the systems additionally contain mechanisms to deal with evidence. Loge-expert was a system targeted at laypeople. It used strategies such as adding a communications layer to translate legal language to layperson language and providing court case examples to allow the layperson to understand legal concepts.⁹⁷²

For my methodology, I will build a system to represent rules in a rule-based encoding, so that the system is able to reason about legal rules in an area and give the user an idea over the rules that are applicable to them.

5.4.5 Everyday decisions will be represented in term of the merits

The system will likewise require previous legal decisions for part of the reasoning process. These decisions will be needed to give the user information about the legal criteria that may be fulfilled by the facts of their situation, see 5.4.1.1. Likewise, case law will be used to give the individual information about the potential outcome of the case, i.e. how much damages they may be able to expect should their case go to court, see 5.4.1.2. Let us explore some aspects of choosing the type of decisions, how they should be represented and how they should be reasoned with.

⁹⁷² Paquin, Blanchard & Thomasset, “Loge-expert”, *supra* note 563.

5.4.5.1 The system will use commonplace cases

In prior work, we saw systems that rely on different types of cases. Systems such as TAXMAN II, HYPO, CATO, IBP and GREBE seem to focus on few, but important cases that have complex representations, often including domain models. These cases seem to correspond to what Stranieri and Zeleznikow refer to as landmark cases, that are “interesting”, often cited as precedent, and reported by the courts. They often introduce new categories of facts, new principles or have unexpected outcomes.⁹⁷³

Other systems rely on a higher number of cases, with simpler representations. These include the systems presented in 4.4.3.4.2.5 and 4.4.3.4.2.6. Here, the representations are usually less complex (such as a list of factors that appear in a case), however the number of cases is higher, see e.g. Yin *et al* which have encoded 900 cases.⁹⁷⁴ It seems like these kind of representations correspond more closely to what Stranieri and Zeleznikow refer to as commonplace cases, that are “uninteresting” on a legal level, but can give an indication of how judges and courts tend to decide in certain areas (see 4.9.3).⁹⁷⁵

In 5.3, I established that the system should target areas of high-volume, low-intensity cases. However, even in such areas, there may be cases that are landmark cases. For example, if a case reveals challenging factual situations, it may be appealed and decided in a higher instance. This case could then act as a landmark case, by influencing the judges in the first instance. Should the methodology focus on these landmark cases, or on commonplace cases in the first instance?

In order to provide the user with the information regarding the legal rights that apply to them and possible outcomes of their case,⁹⁷⁶ I believe commonplace cases are the most important data source. Commonplace cases can give us a window into how judges in reality reason about facts relating to legal concepts, and the outcomes they tend to ordain

⁹⁷³ Stranieri & Zeleznikow, *supra* note 784 at 108.

⁹⁷⁴ Yin, Zulkernine & Dahan, *supra* note 685.

⁹⁷⁵ Stranieri & Zeleznikow, *supra* note 784 at 108.

⁹⁷⁶ See 5.4.1.

in certain cases.⁹⁷⁷ Since, as discussed above, the legal issues raised by high-volume, low-intensity cases may not be the most complex, analyzing commonplace cases to inform the user of how regular cases are typically treated would be a useful pursuit. Therefore, I will focus on such commonplace cases.

Of course, landmark cases may also play a role in areas of everyday legal issues. For example, a landmark case may change the way a certain criterion is interpreted by the courts. Such decisions may be reflected in the commonplace cases that are decided after this landmark case, as judges adhere to the new precedential decisions. Focusing exclusively on landmark cases may not give an accurate window of what the user could expect if they go to court. Landmark cases are, by definition, extraordinary, and thus may not correspond to the situations that most laypeople would encounter.

5.4.5.2 Cases will be represented in a way that allows the comparison between previous cases and user cases

In prior work, we saw a number of ways of representing cases. Some systems rely on the text of a case, or the metadata of a case.⁹⁷⁸ Other systems rely on the merits of a case.⁹⁷⁹ In the latter category, there are a number of possible ways to represent cases, including using factors and deformations,⁹⁸⁰ dimensions and factors,⁹⁸¹ and semantic networks.⁹⁸²

For building my methodology, a key requirement is that the representation should allow the comparison between cases of the user and cases stored in the database. This means that the features of a case must correspond to values that the user can provide about their own case.

This requirement is trickier than it might seem. As described in 5.4.1, we want to provide the user with information regarding the legal rights that may be available to them, and an

⁹⁷⁷ Compare 1.3.2.1.

⁹⁷⁸ See 4.4.3.2 and 4.4.3.3.

⁹⁷⁹ See 4.4.3.4.

⁹⁸⁰ See 4.4.3.4.2.1.

⁹⁸¹ See 4.4.3.4.2.2.

⁹⁸² See 4.4.3.4.2.3.

idea of the possible consequences of taking their case to court. Providing this information is the most helpful right after a situation has occurred, when a user does not know about their legal rights, and what to do about their situation.

All features that we use to encode case law must thus be information that the user knows about their case *before* it has gone to court. This excludes information such as which judge will preside over their case, or which law firm represents the opposing party – before a case has gone to court, it is impossible to know these features.⁹⁸³ Likewise, we cannot ask the user to provide text of how a judge would describe their case to compare it to previous judgments. Therefore, we cannot directly use the textual content of legal decisions to represent them in our system.⁹⁸⁴

However, the user does know what has happened (i.e. the factual situation of their hypothetical case) even before their case has gone to court. Therefore, if we can extract the factual occurrences of previous cases to represent them, we can ask the user to enter their factual situation, and then compare it to the situation of previous cases, to provide them with information about the legal rights applicable in their case, and outcomes of cases that were similar. Of course, as we will see, classifying the facts of a case can be a very difficult problem.

5.4.5.3 Cases should be analyzed to provide useful information

Finally, the cases must be analyzed in a way that allows the providing of the information outlined above in 5.4.1.

As we discovered in 4.9, there are a number of different ways that judge's reason with caselaw. In general, common law jurisdictions seem to understand cases in terms of the factual situation that has occurred in a case, and compare new cases to previous cases by selecting similar cases, applying their reasoning, and distinguishing similar cases that

⁹⁸³ See 4.4.3.3.3.

⁹⁸⁴ See 4.4.3.2.3.1.

have different outcomes. In previous work, there have been a number of systems that emulate this style of reasoning with cases, such as TAXMAN II, HYPO and GREBE.⁹⁸⁵

Civil law cases, on the other hand, seem to be seen more in terms of the legal clarification they provide to certain statutory rules. Systems emulating this style of reasoning include PROLEG and Loge-expert.⁹⁸⁶

However, courts also reason with cases in another way, in order to ensure the consistency of judicial decision making. Here, cases are used more like quantitative datapoints to discover trends in judicial decisions, such as which facts generally correlate with certain legal requirements being fulfilled, or which facts generally lead to which kind of discretionary outcomes. These systems have been emulated using machine learning methods.⁹⁸⁷

In my case, the purpose of the system is to provide information about possible legal avenues and outcomes to the user of the system. I believe using commonplace cases to be the fitting data source for this kind of reasoning. Likewise, I believe the final style of reasoning, i.e. exploring how discretionary decisions are made, to be the appropriate method to analyze and reason with these cases. On the level of everyday cases, the statutory language is likely to be relatively clear, and to have been applied and clarified often in previous cases. The uncertainty faced by a judge in these cases frequently lies in understanding whether certain legal criteria apply to certain situations, and what outcomes to award. I believe it to be possible to model these factors using commonplace cases.

The methodology will thus focus on analyzing commonplace cases in a way that allows the user to understand which legal criteria may apply to their situation, and the outcomes seen in previous cases similar to theirs.

⁹⁸⁵ See 4.9.1.

⁹⁸⁶ See 4.9.2.

⁹⁸⁷ See 4.9.3.

5.5 The system should be practical

Another consideration to make is whether the developed legal reasoning framework should be focused on practical applications or on capturing legal reasoning steps for research purposes.

Susskind has written a paper detailing the difference between “pragmatism” and “purism” in AI & Law.⁹⁸⁸ He argues that *pragmatists* aim to develop working computer systems that “assist in the solving of legal problems”. Pragmatists are not always educated in law, and prefer the use of rule-based shells, that run on personal computers.⁹⁸⁹ *Purists*, on the other hand, are more interested in clarifying central concepts, rather than developing working products. Susskind splits the purists into AI purists, that focus on the computational issues, and jurisprudential purists, that focus on frame-based and extended logic.⁹⁹⁰ Thompson likewise differentiates between sophisticated systems that can autonomously engage in legal decision making, versus more modest and practical approaches.⁹⁹¹

In order to have the greatest real-world impact, I aim to create a methodology that can be used to create legal decision support tools that can be deployed to the real world, i.e. on the practical side. I believe this to be possible without compromising on the legal validity of the developed system.

In order to make a practically useful tool, I intend to not just devise a theoretical approach, but also develop everything that is needed to operationalize the research. This goal could be seen as achieved, since the methodology was used to develop a decision support tool that was deployed to the public in July 2021, see Chapter 8. Let us explore

⁹⁸⁸ Susskind, *supra* note 902.

⁹⁸⁹ *Ibid* at 29.

⁹⁹⁰ *Ibid*.

⁹⁹¹ Thompson, *supra* note 75 at 12.

the implications that stem from aiming to develop a system ready for immediate real-world deployment.

5.5.1 Focus on building a generalizable methodology

The first important implication of building a practical approach is that I want to build a methodology, instead of a single decision support tool. Building a tool that is only applicable in a single legal area can be interesting and a significant contribution. However, in order to achieve the greatest impact on access to justice, the developed methodology should allow the creation of legal decision support tools in multiple areas. This can also prove that such a methodology is able to capture the general idea of legal reasoning, rather than being limited to emulating the style of legal reasoning in a specific area. Susskind believes that the lack of a methodology to create legal expert systems can be one of the reasons for the lack of such systems.⁹⁹²

5.5.2 Focus on frequent types of cases

It is likely impossible to cover every single eventuality that may arise in a case. Even in high-volume, low-intensity legal areas, certain case types may be quite rare, while others are very frequent. Building a practical methodology means accepting that not every single eventuality will be covered by the system. Instead, in order to ensure that the legal decision support tool can have the greatest possible impact with the least amount of work required, the focus of the legal decision support tools will likely be on cases that arise in an area most frequently.

For example, in the area of landlord-tenant disputes in Quebec, the most frequent case is a landlord wanting to expulse a tenant for failure to pay rent.⁹⁹³ Starting with this area would thus lead to a system that is initially able to handle the greatest proportion of cases. The system can then be expanded to cover more types of cases.

⁹⁹² Susskind, “Expert systems in law”, *supra* note 598 at 4.

⁹⁹³ *Rapport annuel de gestion 2020-2021*, by Patrick Simard (Tribunal administratif du logement, 2021) at 45.

At the same time, it is important for the system to be able to adequately deal with cases that are not yet covered by the system. There should be some kind of escape hatch, that makes users aware that they fall outside of the scope of such a system.

5.5.3 Focus on practical approaches to encoding legal information

Building legal decision support tools can be a significant time sink, both when dealing with legal rules⁹⁹⁴ and with legal cases.⁹⁹⁵ Susskind believes the time and effort required to be one of the reasons that there are few working legal expert systems.⁹⁹⁶ In building my methodology, I will therefore focus on a practicable approach to encoding legal information. Such an approach should require as little work as possible to create a feasible legal information tool. The approach should be as simple as possible on the level of representation, while still allowing useful information to be provided. Further, I will prefer approaches that rely on encoding fewer cases, rather than approaches that rely on encoding many cases.⁹⁹⁷

5.5.4 Focus on building intuitive interfaces for the creation of the legal decision support tools

Another important focus of my thesis is the development of an interface that allows the creation of legal decision support tools, by encoding the necessary legal information into the system. Thompson argues for the need of such a system to facilitate the entry of expert system into a computer.⁹⁹⁸ Likewise, Al-Abdulkarim *et al* argue that tools that support a developed methodology are important in making a methodology more teachable and quicker to implement.⁹⁹⁹

The tool should be simple and intuitive enough that individuals without a computer science education are able to gain access to the tool, follow a brief training or watch a video, and start developing their own legal decision support tools. Designing this system

⁹⁹⁴ See 4.5.3.3.1.

⁹⁹⁵ See 4.4.3.4.3.3.

⁹⁹⁶ Susskind, “Expert systems in law”, *supra* note 598 at 4.

⁹⁹⁷ See 4.4.3.4.3.4.

⁹⁹⁸ Thompson, *supra* note 75 at 42.

⁹⁹⁹ Al-Abdulkarim et al, “Factors, issues and values”, *supra* note 702 at 1.

will involve challenges discussed by the human-computer interaction community, regarding how to make interfaces that are easy to learn and use.¹⁰⁰⁰ I am targeting the web as a deployment platform for this creation interface, in order to reach the largest target audience possible and not require individuals to install a program to start using the platform.

I will also explore uses of machine learning in this interface to support the person entering the information into the system.¹⁰⁰¹ As we have seen in 4.4.3.4.2.6, using automatic systems to help with the encoding of information is a very promising approach.

While ambitious, I believe the creation of such an interface is important in order to allow the greatest possible impact on access to justice of my methodology. The ideal group to implement a decision tool is a person or team with legal expertise in a certain area. Typically, these users have to work with engineers to create legal decision support tools. With my interface, this may no longer be necessary, giving the legal experts the tools and methodology, they need to start implementing legal decision support tools in their areas of expertise. Many such tools could rapidly be deployed to increase access to justice across many different areas.

Likewise, such an interface allows the creator of the system to update it to reflect changes in the legislation and jurisprudence. This is crucial in the legal area, since the law changes and adapts.¹⁰⁰² Even beyond the law, expert systems have long suffered from difficulties of maintenance.¹⁰⁰³

5.5.5 Focus on building intuitive interfaces for the end-user

Further, my research will include the development of an interface that allows the use of the system by the public. I will aim to develop a system that makes it easy and intuitive

¹⁰⁰⁰ See 1.3.5.1.

¹⁰⁰¹ See 5.5.3.

¹⁰⁰² Compare Amanda Deitz, “Artificial Intelligence | Towards A Law Assisted By Algorithms”, (8 May 2022), online: *AI Magazine* <<https://ai-magazine.com/problems/artificial-intelligence-towards-a-law-assisted-by-algorithms/>>; Latifa Al-Abdulkarim, Katie Atkinson & Trevor Bench-Capon, “Accommodating change” (2016) 24:4 *Artif Intell Law* 409–427.

¹⁰⁰³ See 2.5.2.

for users to enter their situation into the system and receive information regarding their case. This system should be able to intelligently ask the user for information regarding their specific case, with questions adaptively chosen based on previous answers provided by the user.¹⁰⁰⁴

In building this interface, it is important to make sure that the interface is as clear and easy to use as possible. Researchers in the field of Human-Computer Interaction have dealt with the question of how to make systems that allow the user to explore and navigate enormous troves of information cooperatively with the computer. These ideas are very relevant and important in designing interfaces that allow users to explore their legal rights, previous outcomes, and legal avenues.¹⁰⁰⁵

As discussed in 5.2.7, I believe the web to be a viable platform to reach as many users as possible. Developing an interface for the web that allows the user to interact with the legal decisions support tool is a crucial step in making sure that the system is practical and can be implemented in the real world. As far as I can tell, this is quite unique in the world of AI & Law, as most research does not include the development of interfaces that are ready to be used by end-users.

5.5.6 Focus on giving legal information

I intend my methodology to focus on giving legal information, rather than giving legal advice. I believe this to be an important prerequisite to deploy legal decision support tools, both from a legal perspective and a functional perspective.

From a legal perspective, focusing on providing legal information ensures that the system can be deployed without breaking rules targeting the provision of legal advice by non-lawyers, that exist in many jurisdictions.

¹⁰⁰⁴ Compare Thompson, *supra* note 75 at 16–17.

¹⁰⁰⁵ See 1.3.5.

In Quebec, lawyers have the exclusive right to give an opinion on something related to the law, or drafts documents that can be used in court.¹⁰⁰⁶ For the Barreau du Québec, it seems giving an opinion on an area where there can be multiple different opinions counts as giving legal advice.¹⁰⁰⁷ Similar rules exist in many jurisdictions. In order to not run afoul of these rules, I focus on creating a system that allows the user to explore their own rights, rather than giving an opinion or suggesting what they should do.

Focusing on giving legal information rather than giving legal advice is also important from a functional perspective, in my opinion. Using algorithms to predict the outcome of cases is tricky (just as it is for human lawyers), and are unlikely to achieve perfect accuracies. Even if these algorithms have high accuracies, using them to provide advice risks providing wrong or misleading advice to users in certain cases.¹⁰⁰⁸ Therefore, instead, I aim to conceptualize the system so that it provides the user with information rather than advice and give the user the tools to verify and understand this information. They can then choose whether to let it influence their decision regarding how to proceed with their situation.

Rather than using the system in the sense of “artificial intelligence”, it instead becomes an instance of “augmented intelligence”, giving the user the tools and information to augment their own intelligence, by providing relevant tools and context to allow them to explore their own situation.¹⁰⁰⁹ Therefore, I adopt the terminology of “legal decision support tool”.¹⁰¹⁰

¹⁰⁰⁶ *Act respecting the Barreau du Québec*, *supra* note 510 Article 128.

¹⁰⁰⁷ note 511.

¹⁰⁰⁸ See 3.6.2.1.

¹⁰⁰⁹ Compare Deitz, *supra* note 1002.

¹⁰¹⁰ For discussion, see 1.3.5.3

5.6 Application in building legal decision support tools

In the previous section, we explored the implications of a number of important goals for my methodology. In this section, I will summarize these implications in terms of the different steps of creating legal decision support tools, see 1.2.2.5.

5.6.1 Data representation

The first step in concretely implementing a legal decisions support tool is deciding how the data should be encoded in this system. This approach should be generalizable, to allow its application in many different legal domains (5.5.1). With my methodology, I initially target areas of high-volume, low-intensity cases (5.3).

I aim to build a methodology that allows the encoding of both rules and case law (5.4.3). The legal rules will be represented in terms of logical connections, i.e. using a rule-based approach (5.4.4). Cases will be represented in terms of their merits, in a way that allows the comparison between hypothetical situations of users and previous cases (5.4.5). The encoding will be practical, i.e. focus on simple but useful representations (5.5.3).

I will build a web-based interface to allow the input of rules and cases into the system in an intuitive way, supported by machine learning methods (5.5.4).

5.6.2 Capturing the data of the user

The next step in a legal decision support tool is to capture the information of the user of the tool.

In my methodology, this will allow users to enter their situation in terms of the facts (i.e. what has occurred) or the goal the user wants to achieve (5.2.3). It will not require the user to understand legal concepts (5.2.4).

In order to target the greatest possible amount of people, the system will be intuitive (5.5.5) and web-based (5.2.7), and rely on plain and clear language as far as possible (5.2.6).

5.6.3 Analyzing the data of the user

At this stage, the system will have the prospective case by the user, and a database of stored information, such as case law or legal rules. The next step is to analyze the data of the user, in order to generate useful results.

In my methodology, the system will analyze the case of the user, in order to generate information able to support the user in discovering the rights that they may have in specific situations (5.4.1.1) and the outcomes of previous similar cases (5.4.1.2).

5.6.4 Presenting the results to the user

Finally, once the case of the user has been analyzed in light of the previously stored information, the system has to present the information to the user in a way that is useful to them.

In my methodology, the system will provide the user with information regarding their case, not advice (5.5.6). The information should be as simple as possible (5.2.6) and not require the user to evaluate the quality of the information (5.2.5).

5.7 Conclusion

In this chapter, I have explored a number of implications of the goals I have set for the methodology. These goals include that the methodology should be targeted at laypeople, target areas of high-volume, low-intensity legal issues, be able to provide specific and useful information, and be practical.

I analyzed these goals in light of the previous work in Chapter 4. The result is a number of criteria, summarized in 5.6, on how my legal decision support methodology will be conceptualized.

In the next chapters, I will describe the concrete implementation of a methodology adhering to these criteria. In Chapter 6, I explore the first such methodology (FactorBot), which consisted of encoding cases in a factor-based representation and using this to detect trends in the legislation. While I ultimately focused on another approach, this

methodology was an important step in exploring what is possible and feasible, and laid the groundwork for the final methodology.

In Chapter 7, I present the JusticeBot methodology, which is the main contribution of this thesis. After learning the lessons from FactorBot, this methodology relies on a hybrid case-based/rule-based approach, which works in a collaborative fashion between the user and the computer, i.e., using augmented intelligence.

Chapter 6 FactorBot – factor-based modelling of decisions

Research Objective: Designing a methodology for creating legal decision support tools
(1.2.2.5)

Research Topics:

- How should legislation, court cases and legal information be encoded in order to arrive at a useful result to increase access to justice?
- How can an interface accurately capture the features of a user’s potential dispute?
- How can the information provided by the user be analyzed in order to identify relevant information and relevant previous cases?
- How can information be shown to a user of this system in a way that supports them and encourages the amicable settlement of their dispute?
- How can the accuracy of the system be evaluated, and potential sources of bias be eliminated?

6.1 Introduction

Now that we have determined the criteria that the developed methodology should fulfill, let us explore the concrete implementation of such a methodology. In this chapter, I will present the FactorBot approach, which I developed between the fall of 2018 and the summer of 2019. This approach represents case law based on the binary presence or absence of factual occurrences. The user is then asked to provide their factual situation through an interface, which is compared to previous cases in order to predict applicable legal concepts and provide statistics about the outcome.

After the summer of 2019, the approach used evolved towards a more practical methodology, namely JusticeBot, which will be presented below in Chapter 7. However, FactorBot was an important milestone in my research, and the discoveries made during this time informed many of the choices made in JusticeBot. The research also resulted in

a paper published at ICAIL 2019.¹⁰¹¹ (The ICAIL paper refers to JusticeBot rather than FactorBot. Here I have chosen to refer to this approach as “FactorBot” in order to avoid confusion with the approach described in Chapter 7.)

Therefore, I will present the FactorBot approach in this chapter. In 6.2, I present the hypothetical user interface that could be used to interact with the legal decision support tool. In 6.3, I present the steps used to build this methodology. In 6.4, I discuss the methodology and analyze the promises and shortcomings of the approach, that eventually lead to the shift towards the JusticeBot approach. I then assess to which extent the design criteria set out in Chapter 5 are fulfilled in the developed methodology in 6.5. In 6.6, I summarize the main findings of the section.

6.2 A User Journey through FactorBot

First, let us examine how the user might interact with the FactorBot. It should be noted that the presented interface is a mockup, as the interface for FactorBot was not developed fully, unlike the interface for JusticeBot.

6.2.1 User input

In Figure 5, you can see the planned interface for users to provide their information to the system.¹⁰¹²

¹⁰¹¹ Westermann et al, *supra* note 536.

¹⁰¹² *Ibid* at 133.

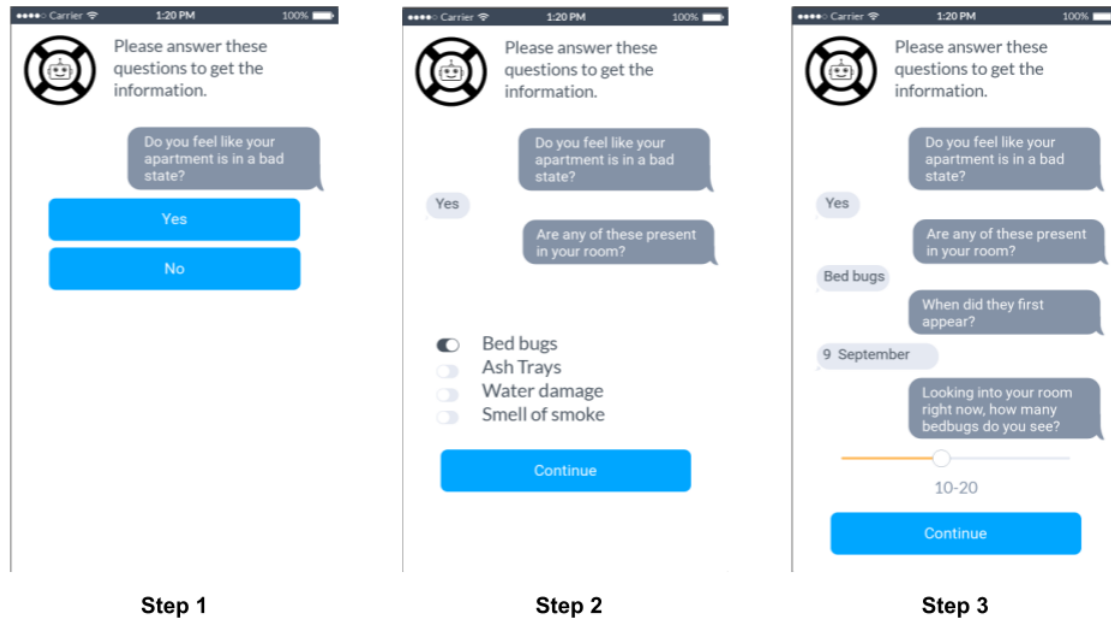


Figure 5 - User Input in FactorBot

As we can see, the system is conceptualized as a mobile app or website.¹⁰¹³ Further, the interface of the system mimics a chat interface, where the user responds to a number of questions in a back-and-forth with the system. This is a common interface paradigm, meaning it would be easy to use for individuals. The texts used are both brief and use simple language.¹⁰¹⁴

The steps the user would follow to provide the details of their situation are shown in Figure 5. In step 1, the system aims to determine the potential legal issue faced by the individual. To determine whether an individual may want to pursue remedies relating to the condition of their apartment, the system asks the user whether they *feel* like their apartment is in a bad state. Phrasing the question in this way allows the triage of the case of the user, without making any assumptions about their case, and without relying on legal concepts.¹⁰¹⁵

¹⁰¹³ See 5.2.7.

¹⁰¹⁴ See 5.2.6.

¹⁰¹⁵ See 5.2.3.

Since the system has determined that the user wishes to explore their options based on the condition of the apartment, in step 2 the system asks the user to provide a list of issues with their apartment. Again, the system does not ask for legal concepts, but rather inquires about concrete factual occurrences, such as bedbugs being present in the apartment of the user.¹⁰¹⁶

In step 3, based on the selection of the user, the system asks more in-depth questions, such as in this case the length of the infestation of bedbugs, and the number of bedbugs the user can see. These details allow the system to capture the situation of the user in more detail.

6.2.2 Providing information

Next, the system shows information to the user. The interface for this is presented below in Figure 6.

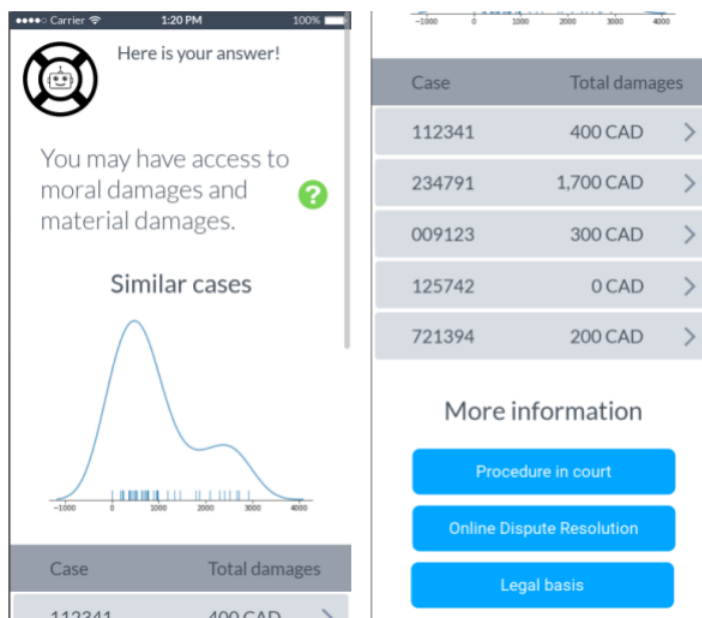


Figure 6 - Information output in FactorBot

¹⁰¹⁶ See 5.2.4.

As we can see, the user is given information about previous similar cases – a list of such cases in order for the user to read and understand the cases, and a curve giving a statistical view over the damages awarded in previous cases.¹⁰¹⁷ Further, the system provides the user with an indication of the remedies they may have access to, based on their situation, and information about next steps (“More information”).¹⁰¹⁸ However, the user is not told explicitly what to do, they are merely given information that may be relevant to them.¹⁰¹⁹

Now that we have seen the planned interface for the FactorBot system, let us examine the process of building this system.

6.3 Building FactorBot

Let us explore the process of building FactorBot, from the perspective of the steps introduced in 1.2.2.5.

6.3.1 Encoding legal data

The first step in building a legal decision support tool is finding a suitable way to encode legal information. Let us explore how this was done in the FactorBot.

6.3.1.1 Representing the Legal Rules and Issues

In order to represent the rules applicable to a certain situation, the FactorBot used the Default Logic Framework.¹⁰²⁰ It represents rules in a rule tree, with the conclusion at the top and the propositions required for the conclusion to be true nested below this conclusion.¹⁰²¹ For FactorBot, we represented the rules relating to the termination of a lease in Quebec in this framework. An extract of this representation can be seen below in

¹⁰¹⁷ See 5.4.1.2.

¹⁰¹⁸ See 5.4.1.1.

¹⁰¹⁹ See 5.5.6.

¹⁰²⁰ Westermann et al, *supra* note 536 at 134; Walker, *supra* note 579.

¹⁰²¹ See 4.5.3.2.4.

Figure 7. Ivan Galindo da Fonseca and Vern Walker worked on producing this representation.

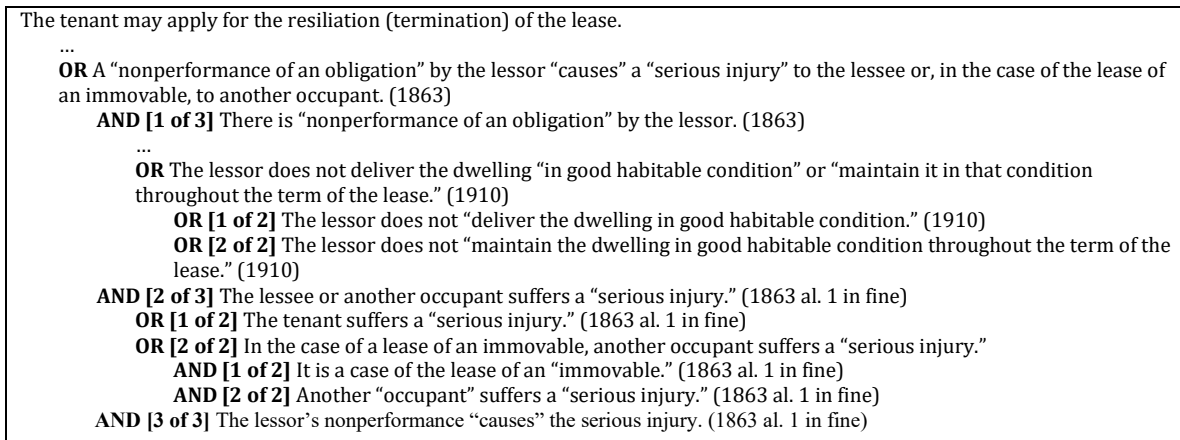


Figure 7 - Partial rule tree for termination of a lease due to the dwelling not being in a “good habitable condition”¹⁰²²

The rule tree represents the necessary conditions that can lead to the termination of a lease. We can see that three requirements need to be fulfilled in order for a lease to be terminated due to the nonperformance of an obligation by the lessor, represented by the AND clauses in the figure. Each AND clause has children that define nested requirements for that conclusion to be fulfilled, which may themselves have children.

This representation thus allows us to logically reason about how rules apply to certain legal criteria, as described in 4.5. For example, if we want to terminate the lease, the system could backward-chain through this rule-tree, to ask whether the criteria that could lead to the lease being terminated are fulfilled.

However, once rules start referring to open-textured legal concepts, it is no longer possible to create further logical rules. For example, one of the conditions in Figure 7 is: “The lessor does not “deliver the dwelling in good habitable condition.” (1910). In order

¹⁰²² Westermann et al, *supra* note 536 at 135.

to support the user in understanding whether their case fulfilled this condition, we decided to analyze case law, to see how judges previously applied this concept.¹⁰²³

6.3.1.2 Representing case law

In order to understand the meaning of the legal concept “deliver the dwelling in good habitable condition”, we used cases from the dataset of decisions by the Tribunal Administratif du Logement, described in 8.3.2.

While the entire dataset comprises around 1 million decisions, we focused on decisions from 2017, which was the most recent year available to us. There were 38,286 cases from 2017. We further narrowed this down by choosing cases mentioning article 1910, 1854 and/or 1864, which are relevant to the discussion of good habitable condition. This yielded 594 cases, of which we analyzed 202 randomly chosen cases. Some of these turned out to be irrelevant, as determined by the annotators, while some cases were annotated by multiple annotators to determine the reliability of the annotations. This left 149 cases that were ready for analysis.¹⁰²⁴ The annotations were performed by 10-20 law student volunteers that participated in weekly seminars. I am very grateful for their contribution.

One problem that we immediately ran into was the fact that the cases were widely divergent in terms of the facts that could lead to an apartment being considered not in a good habitable condition. Unlike some areas, where there is a bounded “shopping list”¹⁰²⁵ that judges may consider in determining whether a legal criterion is fulfilled, there is an immense number of situations that can lead to an apartment not being in a good habitable state. To overcome this, we used a method known as the Grounded Theory Method.

The Grounded Theory Method, described in 1.3.3.2, does not depend on having a theory about the data before analyzing the data. It consists of reading multiple documents,

¹⁰²³ *Ibid.*

¹⁰²⁴ *Ibid* at 136.

¹⁰²⁵ Stranieri & Zeleznikow, *supra* note 784 at 104.

highlighting “memos” of recurring themes, and refining these memos until they can capture the relevant information in the documents.¹⁰²⁶

In our case, this method was used to discover and annotate facts that are present in decisions relating to the good habitable condition of an apartment. We pursued the method in three partially overlapping phases.

6.3.1.2.1 Phase 1 – Factor identification

In the first phase, we aimed to discover the possible factors that contribute to a judge deciding that an apartment is not in a good habitable state. The idea of factors here is similar to that described in 4.4.3.4.2.2 and 4.4.3.4.2.5. We focused on factors that were objective (i.e. do not depend on the subjective assessment of the judge), so that we could ask the user for their situation and compare this to previous cases, annotated in terms of the factors.¹⁰²⁷

In order to identify the factors, we read a number of decisions, and then discussed facts that appeared in the cases. We determined whether the fact represented the type of recurring situation that would be useful to capture for the analysis of cases. If this was the case, we assigned a hashtag to the fact, that could be used to demarcate the existence of such a factor in other cases. Figure 8 shows an excerpt of the hashtags. In total, we identified 44 such factors, forming a taxonomy that could be used to annotate previous cases.¹⁰²⁸

<p>Were any of the following factors present in the case you read?</p> <p>Other deficiencies</p> <p>#Dirty - The apartment is dirty.</p> <p>#Moisture - The apartment has issues with moisture.</p> <p>#Mold - There is mold in the apartment.</p> <p>Infestation</p> <p> #Bedbugs - The apartment is infested with Bedbugs.</p> <p> #Rats - The apartment is infested with rats.</p>
--

Figure 8 - Extract of taxonomy used for classifying case factors¹⁰²⁹

¹⁰²⁶ Webley, *supra* note 42 at 943–945.

¹⁰²⁷ Westermann et al, *supra* note 536 at 136.

¹⁰²⁸ *Ibid.*

¹⁰²⁹ *Ibid* at 137.

6.3.1.2.2 Phase 2 – Factor refinement

The previously identified factors were able to capture the overall facts that were at issue in a case. However, there could be substantial differences between different cases, even when they share factors. For example, a judge would likely consider spots of mold in the bathroom differently from mold that covers the entire apartment and emits heavy smells. Directly comparing these two situations may not be useful. In phase 2, we therefore created in-depth, factor-specific taxonomies that aim to capture more details about the factors. In our initial experiments, we chose to focus on 39 cases that dealt with bedbugs. For these cases, we created a special taxonomy, that you can see in Figure 9.

<p>Which intensity of bedbugs were present?</p> <ul style="list-style-type: none"> • Low (Few bugs, few bites) • Medium (if not specified) • High (Intense infestation, entered furniture) <p>How long were bedbugs present? (in weeks)</p> <p>Is it possible to tell who caused the issue?</p> <ul style="list-style-type: none"> • Tenant • Landlord • Not discussed/ attributed to a third party <p>How helpful was the landlord in solving the problem?</p> <ul style="list-style-type: none"> • Helpful (ordered help when requested) • Not helpful (responded slowly, did not order professional help) • Not discussed/applicable <p>How cooperative was the tenant in helping the extermination?</p> <ul style="list-style-type: none"> • Helpful (Prepared apartment) • Not helpful (Did not prepare apartment, enable access) • Not discussed/applicable
--

Figure 9 - Taxonomy of in-depth factors regarding bedbugs¹⁰³⁰

6.3.1.2.3 Phase 3 – Case annotation

Finally, we applied the created taxonomies to the cases, to annotate the factors that can be found in the cases. In annotating the factors, we focused on factors that the judge saw as proved or that were not contested, since those are the only factors that will influence the decision. We also annotated the outcome of the case, such as whether the rent was terminated and the amount of damages and/or rent reduction awarded.¹⁰³¹

¹⁰³⁰ *Ibid.*

¹⁰³¹ *Ibid.*

In order to annotate the decisions, we initially worked with the Gloss software, developed by Jaromir Savelka.¹⁰³² This software allows the assignment of cases to annotators, and the annotation of cases on a per-sentence basis. For our use case, we annotated decisions on an overall case basis, assigning the hashtags for the presence of factors to the entire case. Later, we switched to using Google Forms, where the annotators filled out a form for each case, selecting whether certain factors were present. This ensured consistency in the spelling of the factor names, while also allowing me to update the list of factors in real-time as new factors were discovered, as per the Grounded Theory Method.

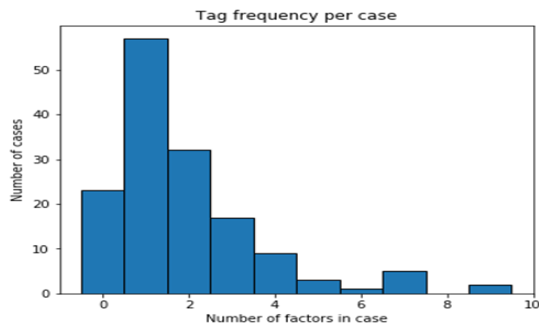


Figure 10 - Tag frequency per case¹⁰³³

Tag	Frequency	Relative frequency
#bedbugs	41	27.50%
#repairsnotconducted	27	18.10%
#mold	16	10.70%
#rainleakage	16	10.70%
#waterleakage	13	8.70%
#heating	11	7.40%
#otherinfestation	11	7.40%
#accesstopremises	10	6.70%
#constantrepairs	10	6.70%
#cockroaches	10	6.70%
#noise	10	6.70%
#intruderprotection	9	6.00%
#landlordunresponsive	9	6.00%
#otheraccessories	8	5.40%
#wallrepair	8	5.40%
#bathroomutilities	8	5.40%
#moisture	7	4.70%

¹⁰³² Jaromir Šavelka & Kevin D Ashley, “Segmenting U.S. Court Decisions into Functional and Issue Specific Parts” (2018) 313 *Legal Knowledge and Information Systems (Frontiers in Artificial Intelligence and Applications)* 111.

¹⁰³³ Westermann et al, *supra* note 536 at 137.

#isolation	7	4.70%
#danger	6	4.00%
#parkingaccess	5	3.40%

Figure 11 - Most frequent factors¹⁰³⁴

In total, we identified 202 factors across the 149 cases. Some cases had no factors at all, likely due to the annotation of only proved factors. Many cases had several factors, showing that multiple issues with apartments were often present simultaneously. Figure 10 shows the number of factors per case. Figure 11 shows the most frequent factors. Bedbugs being present in a case was the most common factor, but we also deliberately oversampled cases with bedbugs in order to experiment with the in-depth taxonomy developed for these cases. In total, we applied the in-depth bedbug-focused taxonomy to 39 cases.

6.3.1.2.3.1 *Inter-annotator agreement*

In order to investigate the reliability with which the annotations were applied, 14 of the cases were annotated by multiple annotators. The results reveal the difficulty of reliably identifying factors in decisions. In only 6 of the 14 cases (43%) did the annotators agree exactly on which factors were present. Difficulties in determining whether the judge saw a fact as proved or not, as well as annotator error, may have contributed to this relatively low agreement.¹⁰³⁵

For the cases targeting bedbugs, the results were similar. The annotators agreed on how long the bedbugs were present in only 45% of the cases. Even this comparatively simple metric proved to be much more difficult to apply than expected. For example, sometimes bedbugs were present for a period, treated for another few weeks, and then returned, although the parties disagreed on how long they returned for. Often, it was not exactly clear how long the judge considered the bedbugs to be present in a specific case. This shows the difficulty of fitting reality into neatly defined categories.

¹⁰³⁴ *Ibid* at 138.

¹⁰³⁵ *Ibid*.

This section has showed the process to identify the factors that are present in a case, and annotating cases in terms of these factors. As we can see, the process is far from simple, something that we will get back to in the discussion section. First, let us see how the situation of the user might be captured, and how the data could be analyzed and displayed to the user.

6.3.2 Obtaining user input

The next step in the legal decision support tool pipeline is to capture the situation of the user. As I mentioned, the FactorBot frontend was not fully developed, since I moved on to the JusticeBot methodology instead. However, let us examine how an interface could be used to capture the information of the user, corresponding to the encoded legal data described in 6.3.1. A mockup of such an interface can be seen in 6.2.1.

6.3.2.1 Forward-chaining

The basic unit of encoding case law in the FactorBot is the factors that are captured from cases. Thus, in order to compare a new case to the coded cases, we need to capture the new case using the same representation. One possible way of doing this could be through *forward-chaining* – the user would simply be presented with a long list of possible facts, and the system would calculate the possible legal conclusions that could be fulfilled from the provided facts. This could work but would present the user with a very long list of factors. In our example above, we annotated 44 factors for a single criterion, plus the potential in-depth factors relating to specific criteria. A real-world system may need to assess many such criteria to determine, e.g., whether the lease of a tenant could be terminated.

6.3.2.2 Backward chaining

A more intelligent approach would be to tie the questioning of the present facts to the possible goals that the user would like to assess. Once we have determined the general legal issue faced by the user of the system, we could ask them to specify the possible goals that they may want to achieve, and *backwards-chain* from these goals to ask only for the relevant facts. The default logic framework allows us to computationally determine the possible legal paths to achieve a certain goal, which would allow the

system to ask the user about only the facts that are relevant to criteria relating to the goal that the user would like to achieve.

Let us assume that a tenant is interested in terminating their lease. As we can see in Figure 7, in order to be able to terminate a lease, three cumulative criteria need to be fulfilled. One of these is the nonperformance of an obligation by the lessor, such as not presenting the apartment in a good habitable condition. If the user wants to terminate the lease, the first step is thus to ask the users for the facts that may relate to the condition of the good habitable state of the apartment. These could include mold, bedbugs, heating issues etc. Once these factors are captured, the system could predict whether this particular criterion is fulfilled (see below), and if so, ask for factors relating to the other required criteria, such as the nonperformance of the obligation causing a serious injury to the tenant. With this method, obtaining the input of the user is thus dependent upon, and intermingled, with the system analyzing the situation of the user.

No matter which strategy is chosen to capture the situation of the user, the next step is for the system to analyze the responses to provide the user with information regarding the legal criteria that could be fulfilled in their situation, or the outcomes of previous similar cases. In the next section, I will look at how this analysis can be undertaken.

6.3.3 Analysis of the case of the user

In the FactorBot, there are two uses for annotated case law. The first use is to predict the applicability of a certain legal criteria for the case of the user. This use is described in 6.3.3.1. The second use is to provide the user with statistics about previous cases that are similar to theirs. This use is described below in 6.3.3.2.

6.3.3.1 Prediction of applicable legal criteria

Once the user has provided their facts relating to a legal issue, the system needs to analyze these facts to determine whether a certain legal criterion is applicable or not. Legal criteria may have immediate consequences, or they may be a pre-condition for another legal criterion as can be seen, for example, in Figure 7. By assessing the

individual criteria, the legal rule tree can be traversed in order to arrive at the final conclusion, such as whether the lease can be terminated or not.

We performed a number of experiments to see whether the 149 annotated cases were sufficient to determine whether the criteria of not providing “good habitable condition” was fulfilled. We then performed similar experiments for the 39 cases annotated with the bedbug taxonomy.

It was not always clear whether a judge found an apartment to be in “good habitable condition”, making annotating this criterion difficult. Instead, we used the award of rent reduction as a proxy for the fulfilment of this criterion, since awarding rent reduction is a direct consequence of an apartment not being considered in “good habitable condition”. Our simplified representation thus considered that the tenant had “won” (i.e. the apartment was not in a good habitable state) when a rent reduction was awarded, and “lost” if a rent reduction was not awarded.

In order to predict whether a judge would award rent reduction or not in a case, we trained a random forest classifier¹⁰³⁶ on the annotated data. The input to the classifier was the presence or absence of the 44 factors in a case. The output was whether the judge would award rent reduction or not. Such a system could then be used as described above, to input the factors of a user case, and predict whether the legal criterion of good habitable condition is fulfilled or not.

We evaluated the system using cross validation. This means that $7/8^{\text{th}}$ of the cases are used for training, while the remaining $1/8^{\text{th}}$ of the cases is given to the system without the label of whether the case won or lost, i.e. for testing.¹⁰³⁷ The predicted values are then compared to the real values to generate a score for the model. This entire process is performed 8 times, rotating which $1/8^{\text{th}}$ of the cases are used for testing. This way, a realistic result can be reported even with the comparatively low amount of 149 samples.

¹⁰³⁶ See 2.6.1.3.

¹⁰³⁷ Compare 2.6.1.2.3.

Overall, the results were rather sobering. The prediction of whether the case would win or lose achieved a precision of 66.5%. Always predicting that there will be rent reduction would achieve a precision of 57.7%. Our prediction was thus slightly stronger. However, using only cases where at least a single factor was present removed this advantage.¹⁰³⁸

We repeated the experiments for the data annotated with the in-depth bedbug taxonomy. Here, the results were similar – the trained model was not much stronger than the baseline of always predicting that the rent would be reduced.¹⁰³⁹

It thus seems like predicting whether a legal criterion is fulfilled based on factors is difficult. We will get back to why this may be the case in 6.4. First, let us discuss another use of the cases – as datapoints that can be shown to the user as outcome statistics.

6.3.3.2 Using cases for outcome statistics?

Let us explore the use of the annotated cases to provide statistics to the user. For each case, we annotated how much rent reduction and moral damages were awarded to the tenant.

In this use-case, the user could provide their factors, and be given a list of previous cases that shared a factor, along with information on the damages awarded in previous such cases. The cases are thus not used to model the classification of facts into legal concepts, but rather to provide statistics on discretionary outcomes of cases to users.¹⁰⁴⁰ As discussed in 5.4.1.2, this kind of information could be very interesting for a user, in allowing them to determine whether to pursue their case, and which amounts they may want to settle on.

¹⁰³⁸ Westermann et al, *supra* note 536 at 139.

¹⁰³⁹ *Ibid* at 140.

¹⁰⁴⁰ See 4.6.




Figure 12 - Rent reduction in months versus number of factors¹⁰⁴¹

Figure 12 shows one instance of treating the cases in this way. The graph shows the correlation between the number of tags that were annotated for a case and the rent reduction awarded, in months. This measurement comes from the fact that rent reduction is often awarded based on the base rent that the tenant pays. For example, a judge might award a rent reduction of 10% of the rent, for four months. In this graph, this would correspond to a total rent reduction of 0.4 months of rent. As we can see, there seems to be a correlation between the number of factors that were present in a case, and the amount of rent reduction awarded.

Tag	Total	Win	Loss	Win %	Avg rent reduction	Avg Moral damage
Who was at fault for the infestation of bedbugs?						
Landlord	1	0	1	0.0%	\$0	\$2,000
Not discussed	36	25	11	69.4%	\$481	\$270
Tenant	2	2	0	100.0%	\$460	\$0
How intense was the infestation of bedbugs?						
High	9	9	0	100.0%	\$889	\$300
Medium	23	14	9	60.9%	\$379	\$392
Low	7	4	3	57.1%	\$216	\$0
Was the landlord helpful in exterminating the bedbugs?						
Helpful	23	17	6	73.9%	\$489	\$183
Not discussed	9	6	3	66.7%	\$442	\$189
Not helpful	7	4	3	57.1%	\$429	\$830
Was the tenant cooperative with the exterminators?						
Helpful	11	9	2	81.8%	\$790	\$437
Not discussed	17	9	8	52.9%	\$383	\$335
Not helpful	11	9	2	81.8%	\$275	\$109

Figure 13 - Statistics in data based on bedbug taxonomy¹⁰⁴²

Figure 13 shows information about damages awarded in cases where the bedbug taxonomy was used. For each factor, the table shows how often it was annotated in total,

¹⁰⁴¹ Westermann et al, *supra* note 536 at 139.

¹⁰⁴² *Ibid* at 140.

how often the cases annotated with it “won” (i.e. rent reduction was awarded) and lost, the percentage of wins, and the average rent reduction and moral damages awarded.

While the data is not always conclusive, there are a few interesting trends that emerge. For example, we can see that cases with a higher “intensity” of bedbugs typically lead to a higher amount of rent reduction. Likewise, cases where the tenant cooperates with the landlord typically lead to higher amounts of rent reduction. The landlord not being helpful typically leads to increased moral damages.

These results have to be taken with a grain of salt, as they are not based on a huge dataset. However, they might still be interesting for users, who could enter their own data, and see how judges previously treated similar cases, and the monetary amounts awarded. Of course, it is important to be cognizant that the individual results may heavily depend on the individual situation and the claims of the party.

Now that we have seen how the case of the user can be analyzed to generate information, let us see how this information can be presented to the user.

6.3.4 Providing results to the user

Once the user has entered their information, and the information has been analyzed, the user should be provided with information about their case. In 6.2.2, I demonstrated how such information could be provided.

One possible output could be the possible legal avenues available to the user. The system could output the predicted fulfilled legal criteria (see 6.3.3.1) and the remedies connected to these criteria, as determined by the legal rule tree. This can be seen as a sort of explanation with rules, since we could show the legal path traversed to arrive at a result.

Further, the user could be provided with information about previous similar cases, and a curve of the outcomes of the cases. Here, the system does not need to predict the applicability of the legal concepts, rather cases are surfaced only based on the shared facts of the user case and previous cases. The cases could also be presented in a list, allowing the user to browse the cases and read them to understand the reasoning typically

applied by a judge. The cases could also be presented in an abstracted way, such as in a curve showing the frequency of different outcomes arriving.¹⁰⁴³

This concludes the overview of the FactorBot system. Let us analyze the particularities of this system, in the light of prior work.

6.4 Discussion

In this chapter, I have presented the FactorBot approach. The system relied on the encoding of legal cases as vectors of facts, that could then be linked to specific legal criteria, represented in a rule tree. We built a part of this system by annotating 149 cases relating to the requirement of an apartment being in a good habitable state.

Let us examine the lessons learned from this project, including the comparison of the system to prior work, and the advantages and disadvantages of the system.

6.4.1 Comparison to prior work

First, let us compare how FactorBot compares to previous work.

6.4.1.1 Encoding cases

FactorBot relies on encoding cases in terms of the factors that appear in that case.

As we saw in 4.4.3.4.2, a number of such encodings have been used in prior work. HYPO used dimensions to represent factors (stereotypical factual occurrences) in trade secret cases. CATO and IBP represented the cases in terms of the presence or absence of binary factors. Similar representations were used in systems relying on nearest neighbor approaches (such as Mackaay *et al* and SHYSTER) and machine-learning based approaches, such as the systems suggested by Alarie *et al* and Yin *et al*. These systems also used binary representations of whether facts were present or absent in prior cases, which allowed them to identify similar cases and predict the outcome of cases. Just like

¹⁰⁴³ Compare Davide Carneiro *et al*, “Using case-based reasoning and principled negotiation to provide decision support for dispute resolution” (2013) 36:3 Knowledge and Information Systems 789–826 at 809.

FactorBot, these systems can be seen to use factors to predict whether a certain legal criterion is likely fulfilled by the appearance of a number of factors.

As I see it, there are two main differences between the approach taken in FactorBot and prior work in encoding cases.

First, a lot of the prior work is concerned with “bounded” domains, where there are a number of predetermined factors that are analyzed in court in order to come to a decision on a legal criterion. These factors can arise from legislation, prior court cases or doctrine. They provide a good structure for annotating cases, as it is likely that the factors capture most of the important information that led to the decision in a case.¹⁰⁴⁴

In our case, there is no such list of factors that tend to be relevant. A vast number of factual situations could influence whether an apartment is in “good habitable condition”. In order to overcome this challenge, we used the Grounded Theory Method to discover the cases while simultaneously annotating cases, by exploring an initial number of cases to identify some factors, and then refining this list as we noticed new factors in cases. This allowed us to come up with a comprehensive list of 44 factors corresponding to possible facts related to the habitable condition of an apartment. I think this is an interesting approach to be able to explore and annotate the factors prevalent in a legal area. However, the approach also carries some challenges, as we will see below.

Second, while a lot of prior work is focused on supporting lawyers or law students in generating arguments or estimating the chances of success of a case, FactorBot is targeted at laypeople. In prior work, factors can sometimes be partly legal in nature, since lawyers are able to reason about whether a legal criterion is applicable and enter this into the system. In FactorBot, this was not the case, as laypeople are unable to do such assessments – the only things they know are the facts of their situation.¹⁰⁴⁵

¹⁰⁴⁴ See 4.4.3.4.3.2.

¹⁰⁴⁵ See 5.2.3.

The factors we annotated were thus things like bedbugs being present in an apartment, the landlord not conducting repairs, mold being present in an apartment and issues with heating. In theory, this is the kind of information that a user of the system can know about their situation, even before it has gone to court. It is thus a “neutral” (or objective) representation, that allows the comparison of new situations to previously decided cases on the same axis, based on objective assessments such as the presence of facts.¹⁰⁴⁶ In 6.4.2.7, I discuss possible difficulties with this approach.

6.4.1.2 Reasoning with cases

FactorBot uses cases for two purposes: Predicting the classification of facts into legal concepts and providing statistics of discretionary outcomes to the user.

6.4.1.2.1 Predicting the classification of facts into legal concepts

In FactorBot, I built a machine learning model that aimed to predict the applicability of legal concepts based on the presence or absence of factors.

Other researchers have used case law for similar purposes.¹⁰⁴⁷ Projects such as HYPO, CATO and IBP, emulate a common law style of reasoning, in order to generate arguments and/or predict the applicability of legal concepts. HYPO retrieved cases similar to a user case based on shared factors, and used these to generate arguments, by highlighting cases that are similar based on some measure, and distinguishing cases that are different. CATO and IBP used similar methodologies to predict the outcome of cases and generate arguments.

This style of reasoning is somewhat different from what we used in FactorBot. The cases in FactorBot are not “interesting” cases, that are used as precedents by courts and that can be cited in order to generate arguments. Rather, they are “commonplace” cases, that illustrate how certain legal concepts were applied in previous cases.¹⁰⁴⁸ Rather than keeping the cases in the database, in order to be able to generate arguments citing the

¹⁰⁴⁶ See 5.4.5.2.

¹⁰⁴⁷ See 4.4.3.4.2.

¹⁰⁴⁸ See 5.4.5.1.

cases, we instead tried to build a machine learning model that is able model the pattern underlying those cases and apply this pattern to new cases.

This approach is more similar to the approaches used by Alarie *et al* and Yin *et al*.¹⁰⁴⁹

These systems encode hundreds of cases and build a machine learning model that is able to predict future cases with high accuracy.

6.4.1.2.2 Output of statistics of similar cases

As described in 6.3.3.2, FactorBot could also use the encoded cases in order to provide statistics on outcomes previously awarded in similar cases, such as the amount of rent reduction. Here, instead of using the cases to assess whether a situation fulfills certain legal criteria, cases are used as examples of the quantitative outcomes that the judge might award based on the facts of a case. For example, a user might enter the fact that their apartment has issues with mold and receive statistics of the typical damages awarded in such cases.¹⁰⁵⁰

There have been a few projects that use case law for the same purpose in prior work. Dahan *et al* used encodings of cases related to employee termination to predict the length of notice period. Stranieri and Zeleznikow used encodings of cases to predict the division of assets in divorces.¹⁰⁵¹

Once more, many such previous systems are developed in areas that seem “bounded”, i.e. have predetermined “shopping lists” of factors that a judge might examine. FactorBot performs similar reasoning in an area where there are no such clear factors. In 6.4.1.1, I describe how we overcame this issue.

6.4.1.3 Connection to legal rule tree

In the FactorBot, the plan was to connect the output of the machine learning algorithm to the rules pertaining to a certain legal area, encoded in the default logic framework. We

¹⁰⁴⁹ See 4.4.3.4.2.5.

¹⁰⁵⁰ See 6.4.1.2.2.

¹⁰⁵¹ See 4.6.3.2.

moved to the JusticeBot approach before completing this step but let us nonetheless compare this approach to prior work.

The rules were encoded in the default logic framework, meaning that the FactorBot shares a lot of context with this work by Walker.¹⁰⁵² We planned to build a machine learning system for each end-node of the rule tree, in order to predict the applicability of each open-textured legal concept based on the facts present in a case.

The rule-tree part of FactorBot further shares similarities with the other work presented in 4.5.3.2, which used rule-based reasoning to solve new cases. Many of these systems faced challenges dealing with open-textured legal concepts.¹⁰⁵³ Here, we aimed to overcome this challenge by introducing machine learning models specific to the legal concepts.

There are a number of systems in prior work that combine rule-based reasoning with case-based reasoning in a similar way. Both CATO and IBP contained a domain-model of the area of trade secrets. This model could be used to arrange legal arguments based around legal issues that contribute to the decision of whether a case contains a trade-secret misappropriation. Each issue has its own set of factors, that can be reasoned about and used to retrieve relevant case law.¹⁰⁵⁴ Similarly, CABARET connected case-based reasoning with a rule-based representation of an area, and used several mechanisms to selectively switch between these two styles of reasoning.¹⁰⁵⁵ Even though the style of reasoning of these systems is different, this manner of structuring and reasoning with cases around legal issues shares a lot of similarity with the FactorBot approach.

The approach in FactorBot can further be compared to Split-up, which used per-step neural networks connected to a structured representation of relevant factors to predict the distribution of assets between divorcees. They built over 20 networks, one for each

¹⁰⁵² See 4.5.3.2.4.

¹⁰⁵³ See 4.5.3.3.3.

¹⁰⁵⁴ See 4.4.3.4.2.2.

¹⁰⁵⁵ See 4.5.3.2.6.

argument, and then applied them to the information of the user in turn.¹⁰⁵⁶ In split-up, the purpose was to predict the percentage split rather than the applicability of legal concepts, however.

6.4.2 Discussion

Now that we have compared the FactorBot to previous work, let us examine some of the positive and negative aspects of this approach.

6.4.2.1 High level of automation

The FactorBot has a high level of automation. It models multiple legal reasoning steps, including the identification of a guiding legal rule (4.2), the assessment of the applicable legal concepts (4.4), the application of rules to the established legal concepts (4.5) and the assessment of the possible legal outcomes of a case (4.6). This allows the system to in theory inquire into the base-level facts of the user situation and perform the steps necessary to arrive at the possible legal consequences and information about the outcome of similar cases.

Providing this kind of information could therefore support laypeople gain a better understanding of their case, without needing any legal knowledge. As far as I can tell, it is one of few approaches that aims to carry out this level of support for the user. The information is further well suited for the user to understand the possible consequences of taking their case to court (by better understanding the possible monetary award) or settling their case (by understanding their BATNA, to allow them to align expectations).¹⁰⁵⁷

Of course, the high level of automation is a double-edged sword when it comes to the prohibition of giving legal advice. An argument could be made that the analysis of the specific facts of a user and prediction of how a judge would classify this could have

¹⁰⁵⁶ Stranieri et al, *supra* note 775 at 171–173.

¹⁰⁵⁷ See 5.4.1.

elements of giving legal advice.¹⁰⁵⁸ This could make deployment of the system challenging from a legal perspective.

6.4.2.2 Effort and difficulty of annotation

Annotating cases using the FactorBot methodology is a time-consuming and difficult endeavor. This is a problem that was also present in prior work.¹⁰⁵⁹ We worked with a number of law student volunteers, who read the cases and assigned factors to them. In a number of weekly sessions, we annotated 149 cases with 44 factors and the in-depth bedbug factors. This may seem like a lot of cases, but it should be noted that these cases were relating to a single legal concept, namely the good habitable condition of an apartment. In order to produce a system dealing with more legal issues, many more cases would have to be annotated, relating to many other legal concepts. Further, from the perspective of machine learning systems, 149 cases are relatively few. This may explain our difficulties in predicting case outcomes.

Beyond the time required for annotation, the process was also surprisingly tricky. We used the Grounded Theory Method to construct a flexible annotation schema able to as fully as possible capture the particularities of each case. However, this also meant that already annotated cases may suddenly no longer be correctly annotated, as new categories were added.

Many factors were very difficult to annotate reliably. Sometimes, judges may not clearly state which fact was proven and which was not, making it difficult to know whether a specific factor should be included. Judges also tend to offer a selective explanation of their decision, focusing on certain disputed legal issues.¹⁰⁶⁰ Annotating which factors were relevant with regards to which legal issue, or even which legal concept was fulfilled or not, thus became a difficult task.

¹⁰⁵⁸ Compare 5.5.6.

¹⁰⁵⁹ See 4.4.3.4.3.3.

¹⁰⁶⁰ See 4.7.1.

Sometimes, the categories were simply difficult to apply. Frequently, putting messy, real-world situations into neat boxes is very tricky. Even factors that may seem simple to annotate, such as the length of bedbugs being present in an apartment, can thus be difficult to capture between cases.

All of these difficulties can be seen in the inter-annotator agreement of the annotations, which was relatively low. This makes building models of the data tricky, as the models may be confused if factors are not applied consistently.

6.4.2.3 Data loss of annotation

Even if annotation does work perfectly, the process of annotating the cases for use in our analysis carries with it a significant data loss compared to the data a judge has access to in making a decision.¹⁰⁶¹

When making a decision, a judge has access to an enormous amount of data. They will listen to the testimony of the plaintiff and defendant, and any witnesses. Subtle nuances of the demeanor of the individuals, such as the facial expression, pauses in responses and the voice and body language may have an influence on how the judge decides. Likewise, the judge may assess documentary evidence, pictures or videos. In writing a legal decision, the document we have access to, only a tiny amount of this information is likely to be included.

Next, we attempt to extract legal factors from this written summary of the case, in order to compare the situations across multiple situations. Here, we reduce a potentially infinitely complex situation, with many particularities and events, to a set of 44 discrete factors. Once more, a lot of the data that the judge might base their decision on is lost in the process.

Finally, in order to make a fact-based representation that we can compare to a new user situation, we have to exclude some relevant context. For example, the user of a new case

¹⁰⁶¹ Westermann et al, *supra* note 536 at 141.

cannot know in advance which counterarguments the other party will raise. Therefore, we have to exclude them from our representation. However, these counterarguments likely play an important role in affecting whether the judge sees certain criteria as fulfilled or not.

The representation we use to capture the cases is thus highly lossy. Even using the Grounded Theory Method, only a tiny fraction of the potentially relevant factors in a real case can be captured. Not having access to the same datapoints as the judge likely makes it difficult to accurately predict the decision of a judge from factors available before a case goes to court.

6.4.2.4 Small dataset

Another issue related to the amount of data that we have access to.¹⁰⁶² One reason is the feasibility of annotating large amounts of case law, described above. However, even beyond this, obtaining enough cases to reliably model the connection between facts and fulfilled legal criteria may be difficult. For example, we were very lucky to work with a dataset of 1 million decisions pertaining to landlord-tenant decisions. Out of these, however, we only managed to identify 39 cases featuring bedbugs, relating to the articles we selected in 2017. With the in-depth factors that we developed, there ended up being very few cases in a lot of categories (such as “the landlord was responsible for the infestation of bedbugs”, see Figure 9). This issue is likely to be exacerbated in areas where there are fewer cases. Generating reliable models from this little data is likely to be complicated.¹⁰⁶³

6.4.2.5 Is it possible to predict these cases using factors?

In the FactorBot, we were not able to create models that accurately predict the relationship between facts and legal concepts. Partially, this is likely due to the practical issues we discovered during the project. The inaccurate and difficult application of the taxonomy means that faulty data may have caused issues for the model. Further, the

¹⁰⁶² *Ibid.*

¹⁰⁶³ *Ibid.*

number of annotated cases may make it difficult to create models and overcome the curse of dimensionality.¹⁰⁶⁴

As discussed, the biggest problem is likely to be the data lost when going from all of the information present in a courtroom to the factor representation we train the model on. In order to be able to translate between the case of a user and the stored cases, we need to use a general representation that only focuses on facts. In doing so, many possible factors that may influence the outcome have to be dropped. Building a model that can accurately predict individual cases based on this reduced amount of information may not be possible.

Even in a hypothetical situation where all possible factors can be captured, I am not sure that legal prediction using artificial systems would be possible. Judges are able to deal with any possible situation, understand how it interacts with the world, and decide whether a legal concept should apply or not. They are able to decide even completely new situations, based on common sense and elements such as societal need and legal intent, and their own discretion. This kind of reasoning is currently likely beyond the scope of artificial intelligence, which often requires thousands of examples to learn relatively simple patterns in data.¹⁰⁶⁵ Further, even lawyers may not always be able to predict what a judge will do – the task of perfectly predicting an outcome of a case may thus be an “impossible” task.¹⁰⁶⁶

In previous work, we saw a number of systems that are able to predict whether legal concepts apply, based on a number of factors, with high accuracy. These were typically in bounded domains, where there are a limited number of factors. Sometimes, cases would be incorrectly predicted. Dahan *et al*, for example, found that certain cases, where exceptional situations arrived, could not be captured by their model.¹⁰⁶⁷ Stranieri &

¹⁰⁶⁴ See 4.4.3.4.3.4.

¹⁰⁶⁵ See 2.6.3.1.

¹⁰⁶⁶ See 2.4.2.5.

¹⁰⁶⁷ Dahan et al, “Predicting Employment Notice Period with Machine Learning”, *supra* note 460 at 25–29.

Zeleznikow came to similar results.¹⁰⁶⁸ The capability of adapting to new situations is highly desirable in the legal context. It seems like AI often has trouble using this kind of reasoning, which means that the result of a prediction will not always match what a judge would decide.

Under these circumstances, it may not be wise to authoritatively tell a layperson user a prediction of how a judge will evaluate their case. They might overvalue the accuracy of the information given by the AI system, and thus make harmful choices. An individual that is told that their case is 90% likely to fail is unlikely to pursue their case, even if in reality, a new factor means that they will win their case in court. Likewise, a case that looks hopeful to the AI system may have uncaptured circumstances that make it unlikely to succeed in court.¹⁰⁶⁹

Based on these insights, it seems like building systems that aim to give a lay user information by autonomously carrying out the steps performed by the judge may not be feasible or desirable. As we have seen, artificial intelligence may not be able to understand the sophisticated world that we live in and predict how judges' reason in each individual case. In prior work, we saw that a lot of systems focus on explainable prediction, where the explanation and arguments advanced by the AI are more important than the prediction itself. According to Dumouchel, the inability of artificial systems to predict whether rules apply in individual cases is an inherent shortcoming on systems that are not embodied in the real world, as they cannot consider and weight multiple viewpoints, and weigh them to come to a fair decision.¹⁰⁷⁰

This led me to change the approach to determine the applicability of legal criteria in JusticeBot. Instead of *telling* the user whether their case is likely to fulfill certain legal criteria, JusticeBot gives them summaries of relevant case law, in order to *support* them

¹⁰⁶⁸ Stranieri et al, *supra* note 775 at 174–175.

¹⁰⁶⁹ See 3.6.2.1.

¹⁰⁷⁰ Dumouchel, *supra* note 98 at 256–258.

in making the decision for themselves. Instead of automatic prediction, the system thus aims to enhance the intelligence of the user by providing relevant context.

6.4.2.6 Is it useful to give statistics to the user?

While predicting whether legal concepts apply or not thus turned out to be problematic, the annotated cases could still be useful to give statistics about possible outcomes. The system may not be able to tell the user if their apartment will be seen to not be in a good habitable condition. However, they may still be interested in the fact that cases that are similar to theirs can lead to awards of 200 – 400 CAD, *given that* they live up to the standard of not being in a good habitable condition. For example, based on this information, an individual may decide that they do not want to pursue their case in court, as the possible results even if they win are not worth the effort.

In practice, there are some challenges with this approach as well. The difficulties of annotating factors in cases also apply to giving statistics. Likewise, a small dataset could lead to skewed statistics.

A system purely indexing decisions based on the factors would also benefit from a way to determine the interaction between multiple factors. For example, is it useful to provide a user that has issues with mold with statistics of cases that have issues with mold and water leakages? Determining how much each individual factor contributed to the outcome would be very challenging.

Finally, giving the user a curve of outcomes could also hide some context specific to each case. For example, perhaps the parties in many cases ask for only 200 CAD regarding a specific issue. In this case, the judge is unlikely to give them more, even if they could have theoretically obtained more damages. However, this will not be visible in the statistics.

6.4.2.7 Difficult to map factors to users

Finally, I want to highlight some potential difficulties of mapping the case in a decision to the case of a user. FactorBot captures the case of the user by asking whether certain facts are present. There are several potential difficulties with this process.

First of all, since the evidence is not captured in FactorBot, being able to compare the case of the user to the decided case means that the user needs to be able to prove the facts. It is important that the user is made aware of this, as they might otherwise believe that they can win, only to lose due to the rules of burden of proof.¹⁰⁷¹

Second, it can be questioned whether the factors selected are at the right “level of abstraction”. As discussed in 6.3.1.2.2, there could be enormous variance in the severity of factors, even when the same factor applies in a case. We tried creating an in-depth ontology to capture more specific factors, initially in cases with bedbugs. However, it is not sure if this approach is viable, as it exacerbates the difficulty of discovering factors and annotating the cases, and further increases the issues of small datasets.

Third, even though we specifically attempted to produce neutral factors, it seems like many of the factors may still involve some subjectivity of the judge. For example, the factors “repairs not conducted by landlord” required the judge to determine which repairs the landlord *should* have conducted. Likewise, the bedbug factor “How helpful was the landlord in solving the problem?” may depend on the judge deciding whether the landlord reacted quickly enough to inquiries etc. This could be overcome with even more factors relating to the meaning of “helpful”, but this, again, would increase the number of factors and thus make the process even more difficult.

6.5 Does the system fulfill the design criteria?

We have now seen the FactorBot system and analyzed some advantages and disadvantages of the system. Let us assess how well this system corresponds to the design criteria set out in Chapter 5.

6.5.1 The system should target laypeople

In theory, FactorBot corresponds well to the criteria set out in 5.2. It allows the user to start using the system without having a specific goal in mind (5.2.3) and does not require them to understand legal concepts (5.2.4). The text used is short and simple, as the user

¹⁰⁷¹ See 4.4.3.4.3.1.

only has to answer simple questions about their situation (5.2.6). Further, the system is well suited to be used on a web platform, such as on smartphones (5.2.7).

However, as discussed in 6.4.2.5, the system does rely on prediction to see whether certain criteria are fulfilled. This prediction did not perform very well and may not be suited to give information directly to laypeople (5.2.5).

6.5.2 The system should be able to handle high-volume, low-intensity legal issues

Another design criterion was that the system should be able to handle high-volume, low-intensity legal disputes, for reasons of feasibility and for a bigger impact on access to justice.

The FactorBot approach is appropriate to deal with such areas. One of the key requirements of the approach is that there are many cases that can be analyzed in order to build models and patterns of the correlation between factors and the outcomes of the cases. Thus, it is likely that the system would work best in areas of high-volume, low-intensity legal issues. However, as we discovered, even such areas may not have sufficient case law to accurately build a model.

6.5.3 The system should give specific and useful information

The system targets areas of high-volume, low-intensity disputes (5.3) in our initial implementation for landlord-tenant disputes. In theory, it allows the user to discover the possible remedies that they have access to (5.4.1.1), by using case-based reasoning to assess whether certain legal criteria are applicable based on facts (5.4.5), and rule-based reasoning through the default-logic framework to decide the consequences of this (5.4.4).

This kind of information is both specific and useful. However, as we saw, this usefulness depends on the prediction working well, which was not the case based on our initial experiments. The cases did however seem useful to provide information about the outcomes of previous similar cases (5.4.1.2), with some possible caveats.¹⁰⁷²

¹⁰⁷² See 6.4.2.6.

6.5.4 The system should be practical

The problem of the FactorBot project, which made me eventually move on to an evolved approach, is that it is not very practical. As we have seen, creating the representation and annotating cases takes significant work (5.5.3), which may impede the possibility of expanding the system to new areas (5.5.1).

Further, even with the annotated cases, the system built to predict the applicability of the legal criteria did not work very well. This limits the practical use of the system. Further, it can be questioned whether the autonomous classification of facts into legal concepts could be considered as giving legal advice – after all, lawyers could have different opinions on whether certain facts fulfil legal concepts or not (5.5.6). This could hamper the practical deployment of the system.

6.6 Conclusion

In this chapter, I presented the FactorBot project. It relies on capturing the facts of the case of a user and comparing their case to other previous cases in order to determine whether certain legal criteria were fulfilled or not, and possible outcomes the user could expect.

While very interesting, this approach faced some difficulties in the practical implementation. Annotating a sufficient number of cases turned out to be a time-intensive and difficult endeavor. Even so, the performance in predicting whether a legal concept applies or not was not very good, which would make it difficult to introduce such a system in practice. I discussed the implications for legal prediction. The results seem to indicate that the prediction of individual cases from pure facts is very difficult and may not be the best approach in building legal decision support tools.

For these reasons and with the lessons from FactorBot in mind, I moved on to the JusticeBot approach. Here, cases are not captured in discrete legal factors. Instead, we summarize the reasoning of the judge with regards to individual legal issues. These summaries can then be given to the user, in order for them to make an informed decision of whether they expect the criterion to apply in their case. The JusticeBot connects this to

a graph-based encoding of the legal rules in an area, in order to guide the user through the relevant questions and provide information and previous outcomes at the end.

Let us discuss how the JusticeBot works, and how it overcomes the issues with FactorBot, in the next chapter.

Part III

The JusticeBot Methodology

Chapter 7 JusticeBot – a hybrid rule-based/case-based augmented intelligence system

Research Objective: Designing a methodology for creating legal decision support tools (1.2.2.5)

Research Topics:

- How should legislation, court cases and legal information be encoded in order to arrive at a useful result to increase access to justice?
- How can an interface accurately capture the features of a user’s potential dispute?
- How can the information provided by the user be analyzed in order to identify relevant information and relevant previous cases?
- How can information be shown to a user of this system in a way that supports them and encourages the amicable settlement of their dispute?
- How can the accuracy of the system be evaluated, and potential sources of bias be eliminated? (see also Chapter 10)

Research Objective: Implementing the resulting methodology (1.2.2.6)

Research Topics:

- How can the methodology be implemented in a concrete, production-ready software stack allowing for the creation of legal decision support tools?
- How can the use of the methodology to create decision support tools be made as effective as possible, and require as little technical knowledge as possible?
- How can a front-end interface be designed for the methodology, that allows a layperson user to interact with the system in a useful way?

7.1 Introduction

After the insights gained from the FactorBot, the methodology evolved into the JusticeBot approach, which I will present in this chapter. The JusticeBot methodology is the most important research contribution described in this thesis.

In FactorBot, we tried to analyze a factual situation, by trying to find the correlation between factors in cases and the applicability of legal criteria. As we saw, this is difficult, as understanding a factual situation is currently beyond the scope of artificial intelligence. Further, it may cause issues with the prohibition of giving legal advice present in many jurisdictions.

The JusticeBot instead represents cases in terms of how a judge ruled on certain legal criterion, and the reasons for doing so. Instead of trying to predict the applicability of legal criteria in new cases, the system presents the user with summaries of how judges previously reasoned with regards the facts of a case, and the outcomes they tied to deciding legal criteria in certain ways. The user can then for themselves decide how judges might reason in their cases and explore the effects of this.

This approach has several advantages. Since cases are used as examples of reasoning rather than datapoints used to build a model, fewer cases are needed. Likewise, the system works together with the user to allow them to explore the outcome of their cases, by providing them with the necessary context and case-law. In this way, the potentially AI-complete task of determining whether a legal criterion applies to a diverse set of facts is “delegated” to the user of the system, with the support of the provided information. The implications of this approach are discussed in-depth below in 7.8.

The JusticeBot methodology goes beyond a theoretical system. Rather, I have created a toolchain that can be used to implement decision support tools using the methodology. This includes the JusticeCreator (described below in 7.3.1.5), which is used by legal experts to create legal decision support pathways, encode case law supported by machine learning, and preview the resulting tool. It also includes the JusticeBot frontend (described in 7.4.1), which is an application that allows the user to interact with the system, to answer questions about their own case and receive information about the potential legal pathways available to them, as well as the outcomes of previous similar cases. This thesis contains a potential step-by-step guide to build legal decision support tools using the JusticeBot methodology (see 7.3.3).

The methodology has been used to implement a concrete decision support tool, targeting landlord-tenant disputes in Quebec, available online at <https://justicebot.ca>. This deployed tool will be described in-depth in Chapter 8. Throughout this chapter, I will use screenshots from these tools to illustrate the functioning of the methodology. Currently, additional JusticeBot tools focusing on consumer disputes, data protection and family law are under development using these tools.

In this chapter, I will describe the different elements of the JusticeBot methodology. First, in 7.2, I demonstrate how the user journey through the JusticeBot looks. Then, I will describe the steps required to implement a decision support tool using the JusticeBot methodology. In short, the steps are as follows:

- **Legal information is encoded (7.3).** A legal expert encodes the legal rules judges apply to deal with certain legal issues into the system in a logical way (7.3.1). Cases are encoded into the system to discover the reasoning schema, to illustrate how legal criteria are applied, and the outcomes that judges award in certain situations (7.3.2).
- **The user enters their information (7.4).** The user navigates the interface to enter a hypothesis for how they believe judges will assess the individual criteria in their situations. To support them, they are shown summaries from the encoded cases, detailing the facts that judges typically consider in deciding on the legal criteria.
- **The user case is analyzed (7.5).** The system analyzes the information provided by the user, by retrieving relevant information encoded in the pathway, and the outcomes of cases that are similar to the case of the user.
- **The resulting information is shown to the user (7.6).** The user is presented with the information collected in the previous step and given a list of possible next steps to undertake.

This chapter further describes the mechanisms in the JusticeBot methodology to collect feedback from the user (7.7). I discuss the approach in 7.8 and wrap up this chapter in 7.9.

The succeeding chapters will explore other aspects of the methodology. Chapter 8 details a case study of the JusticeBot system, where we implanted and deployed a JusticeBot-based tool focusing on landlord-tenant disputes. Chapter 9 details some other areas where JusticeBot tools could be built, and future research work in expanding the platform. Chapter 10 wraps up the thesis.

7.2 A user journey through JusticeBot

Let us examine how a user would interact with a decision support tool created using the JusticeBot methodology. As mentioned, this methodology was fully implemented. The screenshots below are thus taken from a publicly deployed JusticeBot version, focusing on landlord-tenant disputes, available online at <https://justicebot.ca>. This version of the JusticeBot, along with feedback provided by the user, will be described more in-depth in Chapter 8. Here, I focus on the conceptual underpinnings of the system. Due to Quebec being a mostly French jurisdiction, the screenshots are in French, but I will provide relevant translations and explanations.

7.2.1 Interface overview

The user accesses the JusticeBot decision support tool by opening a URL (such as <https://justicebot.ca>). This can be done with a desktop computer or a smartphone. In general, the tool is structured around a number of screens that ask questions to the user in order to understand their situation. All of these screens share certain elements, displayed in Figure 14.

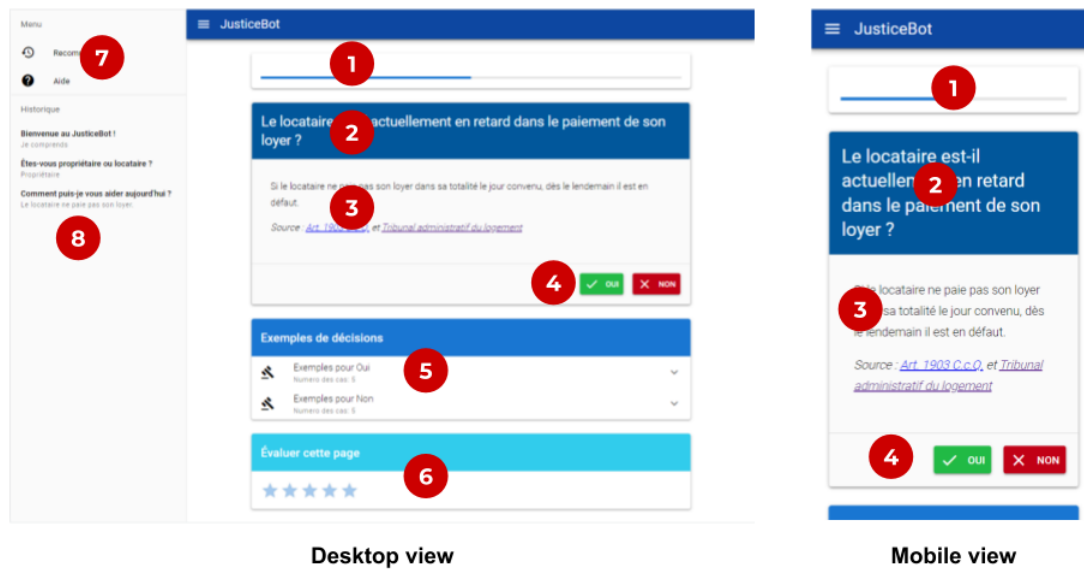


Figure 14 - An explanation of the elements present on a screen of a JusticeBot

The elements have the following functions:

1. **The progress bar** – the progress bar indicates how far the user has progressed through the JusticeBot. The user gets a visual indication of how close they are to the end of a pathway, which can prevent the number of questions from seeming overwhelming.
2. **The title of the question** – The current question that is being asked of the user.
3. **Description** – A plaintext explanation of the particularities of that question.
4. **Answers** – The possible answers a user can give to the question (e.g. Yes/No).
5. **Summaries of case law** – Summaries of how judges reasoned about the current question, sorted by relevance and the answer a judge found in previous cases. These will be explored more in-depth below.
6. **Rating** – Users can rate every single question, in order to provide feedback on confusing or unclear questions.
7. **Help and restart** – The user can select these options to either obtain help about how the JusticeBot works or to restart the pathway from the beginning.

8. **Answer history** – The user can see their previous answers. Further, they can click any question to go back to that question and choose a different answer, in order to explore how the outcome changes with differences in the circumstances.

The mobile view (to the right) demonstrates how the interface looks on phones, in this case the iPhone SE. As we can see, the interface is responsive, and adapts to be legible on the device the user uses. The options present on the left side of the screen on desktop are accessible under the hamburger menu on the iPhone.

7.2.2 User input

Let us examine the path a user would take through a JusticeBot-based decision support tool, here exemplified through the JusticeBot focused on landlord-tenant disputes.

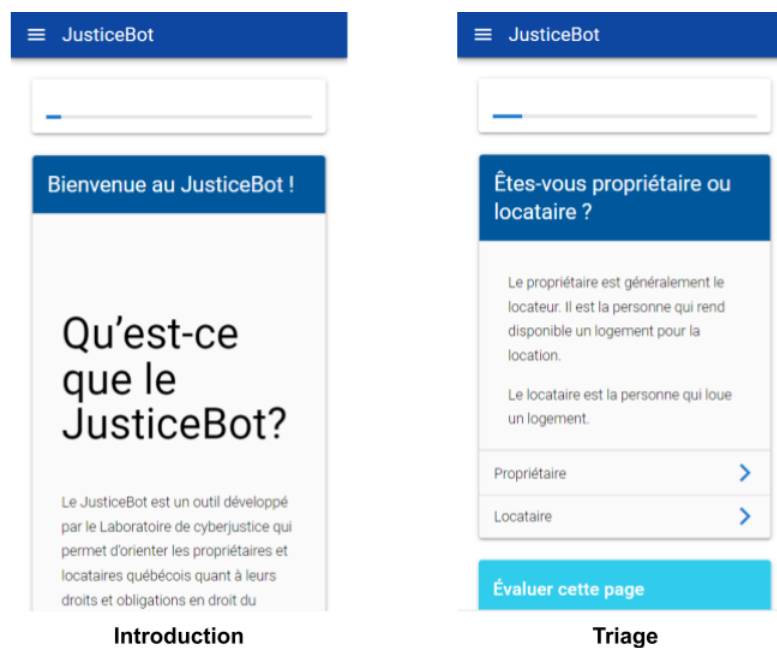


Figure 15 - Initial screens on JusticeBot

Figure 15 shows the initial screens the user sees when entering the JusticeBot. The first screen shows an introduction to the system, including information about the scope of the tool, how it works, what the user can expect and how they should understand the information provided. The page also includes a disclaimer, saying that the information provided information should not be seen as legal advice, but rather as legal information, and that the user needs to decide for themselves how to proceed.

Once the user has read and agreed to this information, they are taken to a triage screen (see Figure 15), that asks them whether they are a landlord (“Propriétaire”) or tenant (“Locataire”). They are further given information to help them in assessing this question.

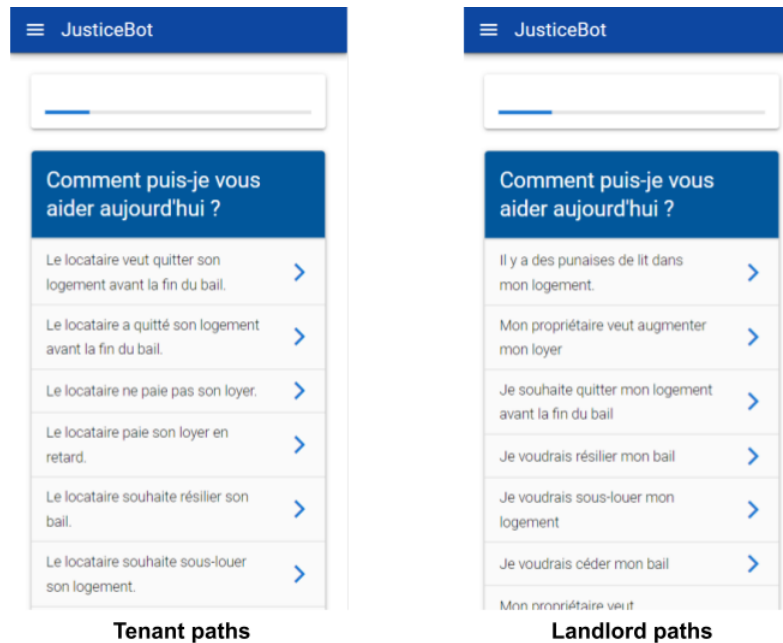


Figure 16 – Possible paths for landlord and tenants

Depending on whether the user chooses the tenant or the landlord option, they are presented with a portal screen, showing a number of paths that they can pursue, as seen in Figure 16. These paths can be oriented either towards achieving a certain goal (such as breaking the lease) or exploring the possible consequences of a factual situation (such as the tenant not paying the rent, or bedbugs being present in an apartment).

Once the user chooses a path, they are asked a number of detailed questions that aim to understand the particular situation of the user. Figure 17 demonstrates how this looks if a landlord wishes to explore their options when a tenant no longer pays their rent.

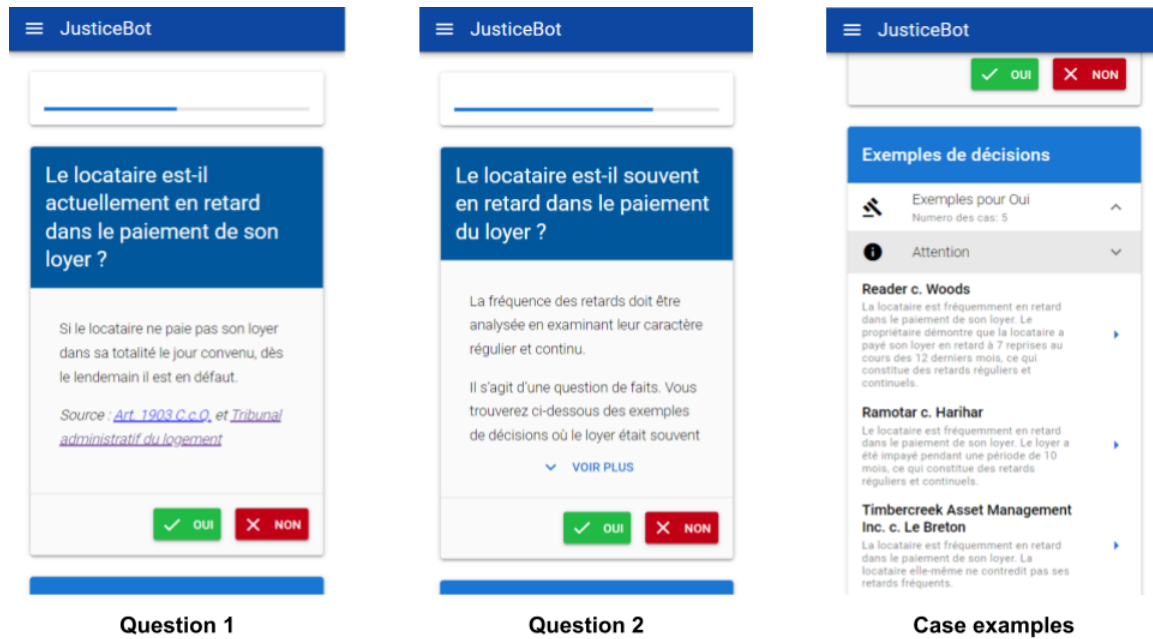


Figure 17 - Path for exploring options after tenant is late with paying rent

The first question displayed asks whether the tenant is currently late with paying their rent. As described above, the question has a title containing the main question, and a description containing an explanation of the question in plain language.

Depending on the answer of the first question, the interface adapts. If the user selects the “Yes” option, they would be asked if the tenant is more than three weeks late with paying their rent. This is a legal threshold that enables the termination of the lease.

If, however, the user selects the “No” option, indicating that the tenant is not *currently* late with paying their rent, the threshold is not relevant. Instead, the second question shown in Figure 17 is displayed, asking whether the tenant is *frequently* late with paying their rent. The rent being paid late frequently is also a legal criterion, that gives rise to certain rights. However, it is an open-textured term – what can be understood under “frequently late” is not detailed in the law, but rather determined by the judge regarding individual cases. Therefore, the user is presented with a number of summaries of how judges previously reasoned about the criterion of frequent lateness. The user can see summaries of five cases where the judge did consider the circumstances to fulfill the criterion of frequent lateness, and five cases where the judge did not find this to be the

case. Each case is summarized with regards to the facts that made the judge come to a certain conclusion, which allows the user to compare their own case to these previous cases to assess whether the judge is likely to find a certain criterion to be fulfilled or not. For example, in Figure 17, we can see that the judge previously found that the tenant paying their rent late in 7 out of the past 12 months could be considered to fulfill the criterion “frequent lateness” in a specific case.

If the user, in this case, selects the answer “Yes” and also answers “Yes” to the next question (which asks whether the frequent lateness of the tenant causes them a serious prejudice), they will arrive at the end of the pathway, and be presented the results of their inquiry.

7.2.3 Providing information

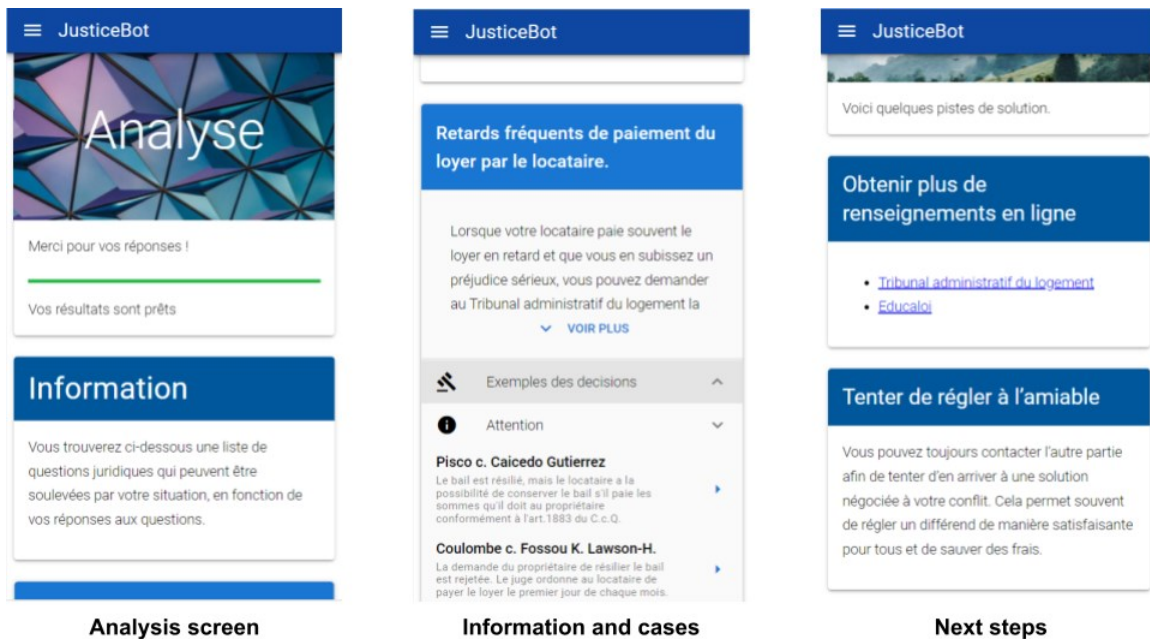


Figure 18 - Information provided after completed pathway

Once the user has arrived at the end of a given pathway, they arrive at the “Analysis” screen, as can be seen in Figure 18. This screen contains legal information that is relevant to the situation of the user. In this case, we can see that the user is informed that they may be able to terminate the lease, due to the frequent lateness of the tenant, and this frequent lateness causing serious prejudice to the landlord. Of course, as is made clear by the text

and a disclaimer, every case is unique, and the assessment depends on the judge agreeing with the user's assessment about the different criteria of their case.

The user is further given access to a list of previous cases. Here, the outcome of the case is presented, rather than the reasoning of the judge. Thus, the user can explore the real-world outcomes of cases that are similar to theirs. In this case, we can see that the judge does not always terminate the lease, even if the legal criteria are fulfilled. This information could be very relevant to help the user understand their situation and the possible outcomes.¹⁰⁷³

After this, the user is presented with a screen that gives them a list of possible next steps (see Figure 18). In this case, the screen refers the user to more information, and lists options including contacting a lawyer or trying to settle their dispute.

At the very end, the user is given the option to fill out a survey about their experience. They can restart the JusticeBot to explore a different issue or change any of their answers to explore how different aspects could lead to a different outcome. The answers given by the users with regard to the Landlord/Tenant JusticeBot will be described below in section 8.4.

This concludes a possible user path through the JusticeBot based tool. In essence, the user is guided through a guided pathway that asks them to specify how they believe a judge would see their situation. Based on this hypothesis, they are given information about the possible outcomes a judge may order, and what next steps they can undertake to address their situation.

Next, let us explore the process of building a JusticeBot legal decision support tool. The first step of the methodology is encoding the required legal information (including legislation and case law) that is required into a computer system.

¹⁰⁷³ Compare 5.4.1.2.

7.3 Encoding legal information

The first important step in creating a JusticeBot legal decision support tool is the encoding of legal rules and legal cases. Let us take a look at how this encoding is done. I will also present the JusticeCreator, which is the interface I created to allow legal experts to build legal decision support tools in a simple and intuitive manner.

7.3.1 Encoding legal rules

7.3.1.1 Introduction

JusticeBot uses the legal rules covering an area as the core for the system. These rules correspond to the reasoning of the judge described above in 4.5, namely starting with fulfilled legal criteria and then logically traversing legal rules to arrive at a legal conclusion.

The JusticeBot encodes this kind of reasoning in a rule-based system. For every legal area, we create a schema that captures the logical connections between the legal criteria a judge typically applies. As we will see, this schema has two purposes:

- It serves as an index to represent case law. Each case is represented in terms of how judges decided on the individual criteria, why they did so, and the outcome of the cases. These representations can then be compared to the case of the user, to surface relevant case law. This purpose of the schema is described below in 7.3.2.3.
- It serves as a pathway for the user when interacting with the JusticeBot, as a means to describe their case in terms of whether they believe that a number of legal criteria apply in their situation. This purpose of the schema is described below in 7.4.

The rules that are encoded into the JusticeBot will often stem from legislation. However, one of the key ideas of the JusticeBot methodology is that we want to capture the rules *as they are applied by judges*. Often, the reasoning steps applied by judges correspond closely to the legislation. However, occasionally, additional criteria targeting specific situations may have developed through precedents. To capture these, the encoding of the

rules is done in conjunction with reading case law. The use of case law in the methodology is discussed in 7.3.2.

In order to illustrate the encoding of legal rules into a standardized format, I will examine the following paragraph from the Code Civil du Quebec, governing whether a lease can be terminated as a result of the lateness of rent payment by a tenant:

1971. The lessor may obtain the rescission of the lease if the lessee is over three weeks late in paying the rent or, if he suffers serious injury as a result, where the lessee is frequently late in paying it.

The paragraph introduces two parallel ways that a lease can be terminated based on the non-payment of a tenant. The first way stipulates that a lease may be terminated if a tenant is currently over three weeks late in paying their rent. Alternatively, the landlord (lessor) may obtain a termination of the lease if the tenant is frequently late with paying their rent, *and* the lessor suffers serious injury as a result. Here, there is no requirement for the tenant to be currently late.

The article thus contains three important elements that tell the judge how to reason in a case that falls within the purview of the article: Legal criteria, a legal conclusion, and the logical connection between the criteria and the conclusions.

7.3.1.1.1 Legal criteria

The article contains a number of legal criteria that must be assessed by the judge. These legal criteria include:

- *The lessee is over three weeks late in paying their rent*
- *The lessee is frequently late in paying their rent*
- *As a result of the frequent lateness of rent payment, the lessor suffers a serious prejudice.*

The reasoning on whether such criteria apply to a case or not is described above in 4.4. As I noted, the determination often requires the interpretation and application of open-textured legal terms, such as above “frequently late” or “serious prejudice”.

7.3.1.1.2 A legal conclusion

The article assesses whether the judge should order the termination of a lease, based on a lessee not paying their rent. This is a legal conclusion, that comes as a logical consequence to the legal criteria being fulfilled.

Of course, the article only covers a specific situation. There may be other possible legal options to terminate the lease, described in other articles. Further, there may be other legal conclusions that apply based on other articles. For example, even if the tenant is only late by one week, they can be ordered to pay the rent, irrespective of whether they are frequently late with their rent. This criterion comes from article 1855 of the Code Civil Quebec, and general contract law.

Likewise, a judge may be given a margin of discretion in how they apply the conclusion. In this case, the law gives the judge the right to supplant the termination of the lease with an order to the tenant to pay on time in the future, see article 1973 of the Code Civil Quebec. Differences in outcomes that can be awarded in a case is discussed above in 4.6. In the JusticeBot methodology, multiple cases with different outcomes can be used to illustrate this discretion.¹⁰⁷⁴

7.3.1.1.3 The logical connection between the criteria and the conclusion

This paragraph above also specifies how different legal criteria may be connected to each other, using logical connection words such as “and” and “or”. If a landlord wants to terminate a lease based on this article, they are given two parallel paths of argumentation. First, they can prove that the tenant is currently more than three weeks late. *Or*, they can prove that the tenant is frequently late *and* that they suffer serious injury as a result. These paths are thus *alternative* – either can be fulfilled in order for the consequence (lease termination) to be decided upon.

If the landlord chooses to argue for frequent lateness of rent, they will also have to prove that they suffer a serious injury as a result of this frequent lateness. Both of these criteria

¹⁰⁷⁴ See 7.3.2.3.

have to be fulfilled in order for lease termination to be decided upon by the judge. They are thus *cumulative* requirements.

A logical display of the article may look as follows:

The lease can be terminated if:

The lessee is over three weeks late in paying their rent

OR

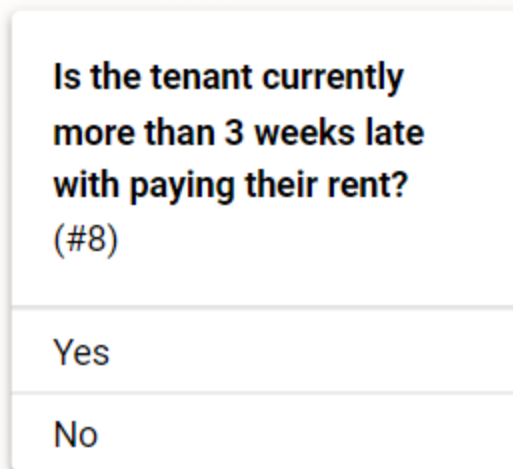
The lessee is frequently late with paying their rent AND as a result of this, the lessor suffers a serious prejudice.

In a way, the judge can be seen to traverse a graph of the different rules, by deciding on the individual legal criteria, and then moving on to the next criterion, until they arrive at a legal conclusion mandated by the law. For example, in order to decide upon lease termination based on frequent lateness, they would first determine whether the proved facts mean that the criteria of “frequent lateness” is fulfilled. If this is the case, they will move on to determine whether the landlord can be seen to suffer serious injury. If this is also the case, the judge will decide that the consequence should be a lease termination. If either of these criteria do not apply, the lease cannot be terminated due to this article of the law. This kind of reasoning is described above in 4.5. The JusticeBot will have to capture this kind of logical traversal in order to give specific information to the user.

Let us explore how these elements, namely legal criteria, legal conclusions, and the logical connections are represented in the JusticeBot framework.

7.3.1.2 Legal criteria

Legal criteria can be seen as questions that the judge needs to answer in order to arrive at certain legal conclusions. The criteria could stem from legislation or case law. In the JusticeBot framework, the legal criterion is referred to as a “question block”.



Is the tenant currently more than 3 weeks late with paying their rent? (#8)
Yes
No

Figure 19 – A question block, representing a legal criterion

Figure 19 demonstrates how a question block is visualized in the JusticeCreator. Each block has, attached to it, the following information:

A **title**, posing the legal criterion a judge will apply in a certain situation as a question. This title should correspond to the criterion a judge will apply in legal cases regarding a certain area. However, since this question will also be used as a method to allow end-users to explore their legal situation, the language used should ideally be simple enough not to present any comprehension difficulties for the target user. As we saw in 7.3.1.1, one legal criterion that the judge needs to assess is whether the tenant in a situation is more than three weeks late with paying their rent. The corresponding question can be seen above in Figure 19.

A **description** (not seen in the image) that explains and clarifies the specific criterion. In this case, this could be related to when the calculation for three weeks begins, whether a partial payment counts as not being paid etc. This description should be written in as simple language as possible, as its purpose is to support the user in deciding whether the criterion may be fulfilled in their case.

A number of **possible answers** to the legal criterion. In the figure above, the answers are "Yes" and "No". Each of the answers can be logically connected to a succeeding question or information block. This works well to build a logical flow of legal criteria that may apply or not, as we will see. However, it is also possible to change the text of the answer, or to add more than two answers.

Summaries of case law attached to each answer (not shown in the image). The JusticeBot framework uses the schema as a way to index case law. Each question block can thus contain a reference to case law, with summaries regarding how a judge has applied a certain criterion in previous cases. This supports the user in understanding whether a judge would find the criterion to apply in their case. The addition of case law to the schema is described more in-depth below at 7.3.2.2.

7.3.1.3 Legal conclusions

As described above, based on the applicable legal criteria, the judge finds certain legal conclusions. Some of these may have consequences attached to them immediately ("the lease should be terminated"). Others may be intermediary conclusions, that serve as input for other legal rules. In the JusticeBot, legal conclusions are referred to as "information blocks".

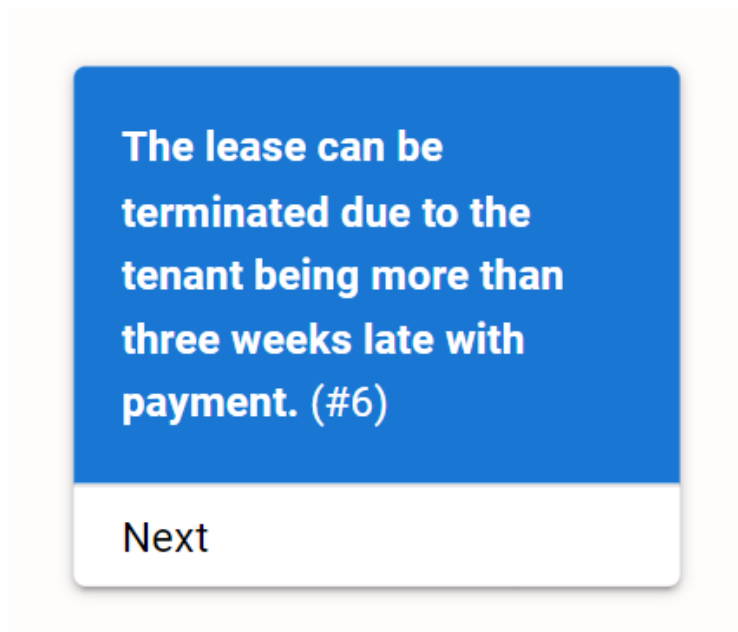


Figure 20 – An information block, representing a legal conclusion

The information block contains the following information:

A **title**, naming the legal consequences that a rule introduces.

A **description** (not seen in the image) that explains and clarifies the possible legal consequence. The description should contain plain language in order to be understood by layperson users of the system.

Summaries of case law attached to the legal conclusion (not shown in the image). Just like the question blocks, information blocks can contain a list of case summaries. These aim to summarize the consequences that a judge ordered in previous cases, so that the user can get an idea of how similar cases were dealt with. This way of encoding case law is described below in 7.3.2.3.

Each information block further has a single “next” connector that allows us to connect it to a question or information block. This allows information blocks to encode intermediary legal conclusions, as explored below under 7.3.1.4.3.

7.3.1.4 The logical connections

Question blocks and information blocks are the key fundamental components of the JusticeBot framework. In order to build these components into a useful rule-based representation of legal reasoning, they need to be connected together in a pattern that follows the reasoning of the judge.

7.3.1.4.1 Directed Acyclic Graphs

In the JusticeBot, I decided to encode the legal rules governing an area into a graph. More precisely, the system uses directed acyclic graphs (DAGs) to model the way a judge reasons about a legal issue. Such graphs consist of nodes (the circles in Figure 21) and edges connecting these nodes (the arrows in Figure 21). The specialty of the DAG is that each edge has a direction,¹⁰⁷⁵ and that the graph cannot form a closed loop (i.e. it is impossible to go backwards).¹⁰⁷⁶ DAGs have been used in AI & Law previously to model subtasks of legal prediction.¹⁰⁷⁷ Figure 21 shows an example of a DAG. The letters represent nodes, while the arrows are edges connecting these nodes.




Figure 21 - An example of a directed acyclic graph¹⁰⁷⁸

In the JusticeBot, I have chosen the DAG format to represent the rule traversal a judge is likely to engage in. In my opinion, the DAG structure closely matches how judges reason about a case, with the nodes representing legal criteria or legal conclusions, and the edges

¹⁰⁷⁵ K Thulasiraman & M N S Swamy, *Graphs: theory and algorithms* (New York: Wiley, 1992) at 97.

¹⁰⁷⁶ *Ibid* at 118.

¹⁰⁷⁷ Haoxi Zhong et al, “Legal Judgment Prediction via Topological Learning” (2018) Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing 3540–3549.

¹⁰⁷⁸ A3nm, *Tred-Gprime.svg* (2013) Wikimedia commons.

representing the connections between the two. Just like these graphs, legal reasoning is directed – the judge traverses the legal criteria in a certain direction, in order to arrive at legal conclusions.¹⁰⁷⁹ Just like the DAGs, judicial reasoning is also acyclic – if a criterion has been assessed with regards to a certain case, the judge does not have to return to assess the criterion again.

The DAG structure can be used to link the question and information blocks together in a way that approximates the legal reasoning of a judge. Below I will give a few examples of how this can be done.

7.3.1.4.2 Single criterion

In the simplest case, a single criterion may need to be assessed to arrive at a legal conclusion. For example, let us imagine a system that exclusively deals with the termination of the lease due to lateness of rent payment of more than three weeks. Figure 22 shows how this could look when represented in a DAG.

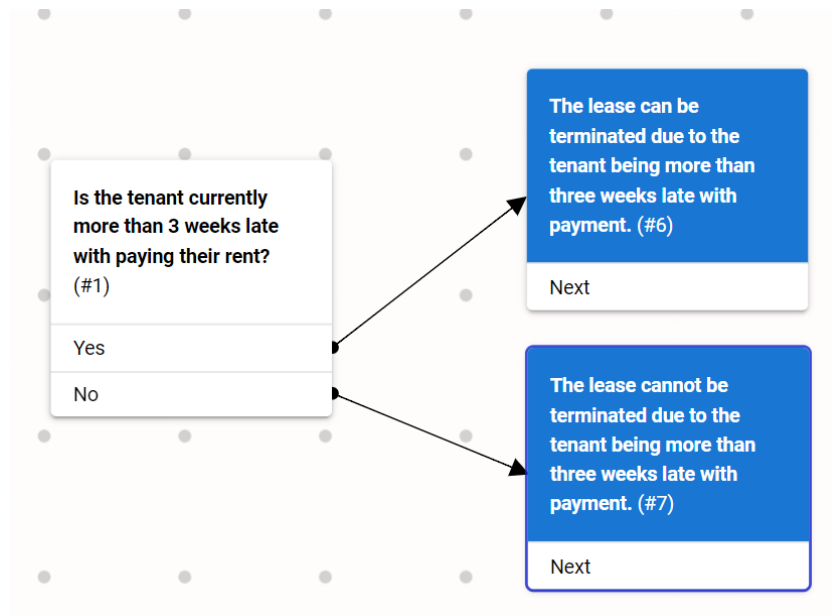


Figure 22 - DAG representing a legal criterion and two possible legal conclusions

¹⁰⁷⁹ Compare 4.5.

If the judge believes that the tenant is more than three weeks late with paying their rent, they will arrive at the legal conclusion that the lease can be terminated. Otherwise, they will likely say that the lease cannot be terminated. These connections can be read from the legislation quoted above.

7.3.1.4.3 Intermediary conclusions

I previously mentioned that we may want to introduce intermediary conclusions into the system. This could include, for example, the judge ordering the tenant to pay their rent if the tenant is late with paying. This conclusion can be encoded in the system in the following way:

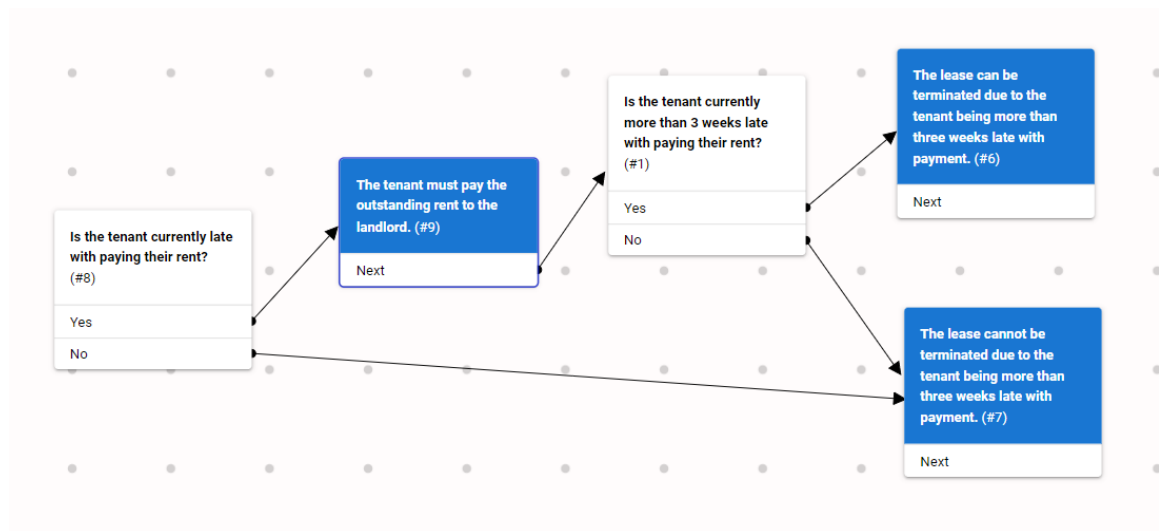


Figure 23 - A schema with intermediary conclusions

As we can see, if the tenant is currently late with paying their rent, the judge could arrive at the intermediary conclusion that the tenant must pay the outstanding rent to the landlord. However, the analysis of the case is not completed. The judge must also assess whether the tenant is more than 3 weeks late with paying their rent, in which case the lease shall be terminated. Of course, in case the judge finds that a tenant is, in fact, not late with paying their rent, the question of whether the tenant is more than 3 weeks late becomes obsolete. Instead, the judge could just skip to the end, concluding that this ground for terminating the lease does not apply.

7.3.1.4.4 Cumulative criteria

In order to terminate the lease for frequent lateness, two cumulative criteria must be fulfilled. Let us see how these could be encoded in the graph form.

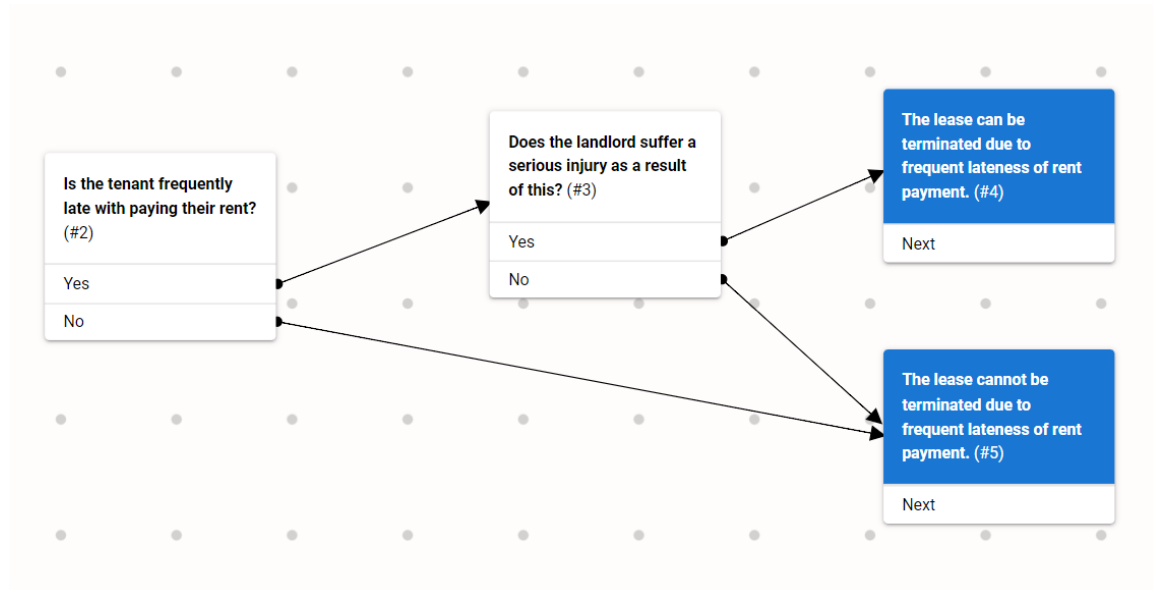


Figure 24 - Cumulative criteria in a DAG

This graph models the criteria described above for a judge to terminate the rent based upon the frequent lateness of a tenant. This requires two separate criteria to be fulfilled: The frequent lateness of rent payment and the landlord suffering a serious prejudice based on this frequent lateness.

The graph above captures the cumulative nature of these criteria. As we can see, if either of the criteria is answered in the negative, the resulting legal conclusion is that the lease cannot be terminated. However, if both are answered positively, the legal conclusion will be that the lease can be terminated.

7.3.1.4.5 Alternative criteria

In some cases, there may be two different legal criteria that are sufficient to obtain a legal conclusion. As noted above in 7.3.1.1.3, article 1971 of the CCQ presents two possible ways to obtain the termination of lease: Lateness of rent payment of three weeks, or frequent lateness of rent payment. These requirements are thus alternative. Figure 25

shows how such a logical connection could be represented. In this case, for the sake of clarity, the requirements for frequent lateness of payment and a serious prejudice have been merged into the same question block.

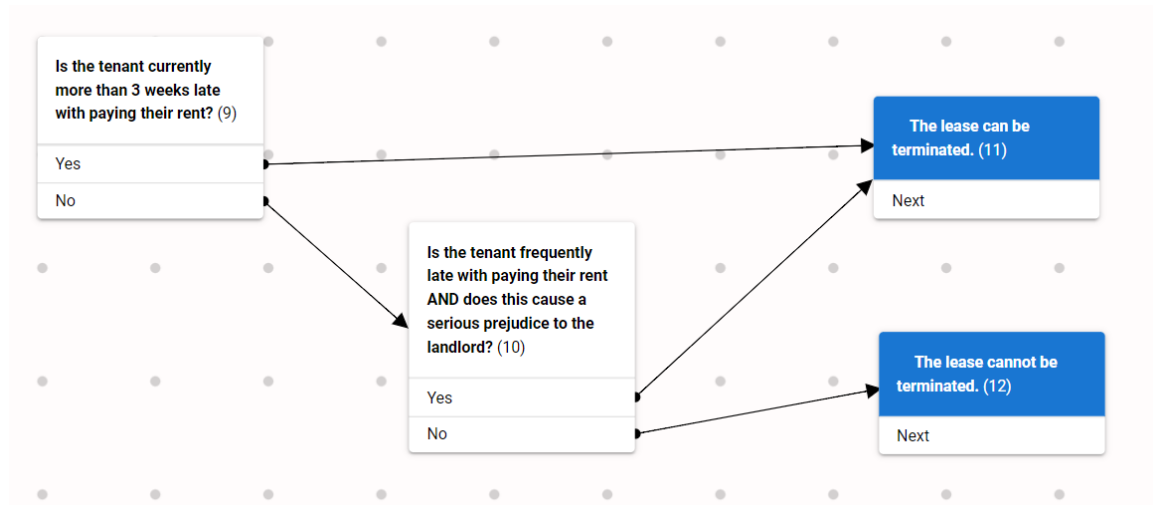


Figure 25 - Alternative requirements

As we can see in Figure 25, if either of the criteria are answered with “Yes”, the result will be that the lease can be terminated. Only if both criteria are found to not apply will the answer be that the lease cannot be terminated.

7.3.1.5 Encoding legal rules in the JusticeCreator

Now that we have seen the logical ways of representing different styles of reasoning in the JusticeBot methodology, I will describe the JusticeCreator, which is a tool to allow the intuitive creation of such schemas by legal experts without programming knowledge.

I designed and programmed the JusticeCreator in Spring 2021. It is now being further developed and enhanced by me and others at the Cyberjustice Laboratory. The JusticeCreator was used to build the JusticeBot TAL, which contains over 300 elements and has been accessed by over 17k users, as of February 2022. JusticeBot TAL will be described more in-depth below in Chapter 8.

The JusticeCreator is built using the Quasar framework,¹⁰⁸⁰ which leverages the reactive framework Vue.js¹⁰⁸¹ to produce web-based interfaces. It can thus be used on any computer with an internet connection. It exports the data in the form of a JSON-file, which can then be imported into the JusticeBot frontend and be made accessible to the public.

Next, I will describe how the JusticeCreator can be used to encode rules to create a new legal decision support tool.

7.3.1.5.1 The interface

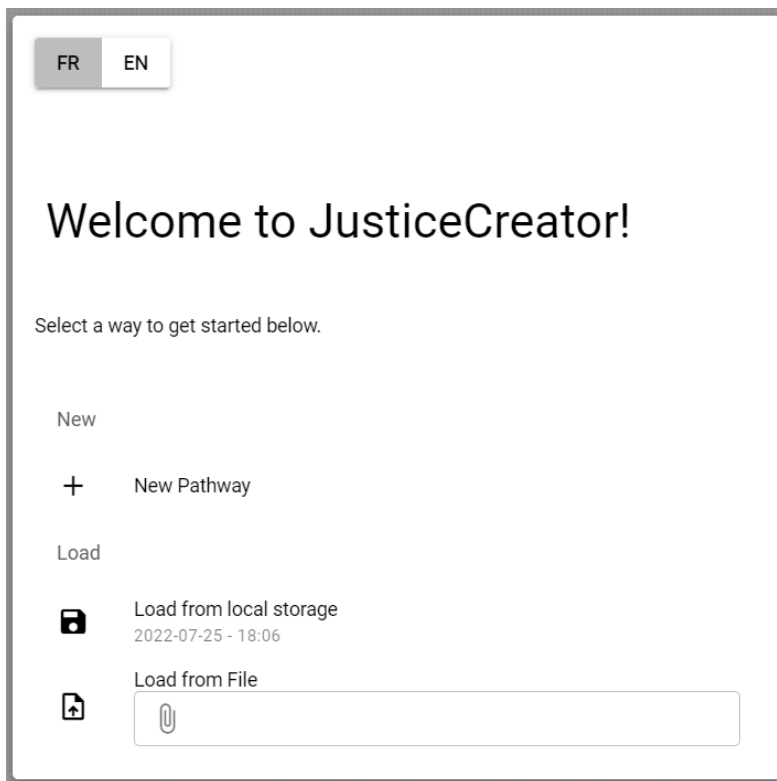


Figure 26 - Welcome screen of the JusticeCreator

¹⁰⁸⁰ “Quasar Framework - Build high-performance VueJS user interfaces in record time”, online: *Quasar Framework* <<https://quasar.dev/>>.

¹⁰⁸¹ “Vue.js - The Progressive JavaScript Framework”, online: *Vue.js* <<https://vuejs.org/>>.

Upon entering the JusticeCreator, the user is faced with the welcome screen shown in Figure 26, allowing them to either start a new project or load a previous project. Once they have made this choice, they are taken to the interface displayed below in Figure 27.

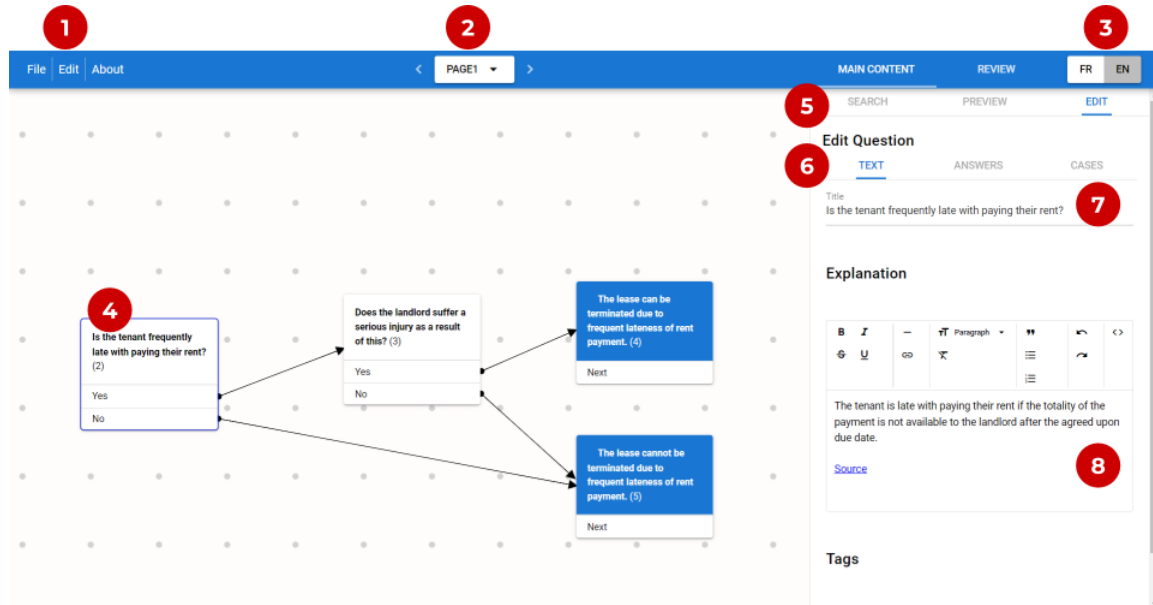


Figure 27 - Overview over the JusticeCreator interface

The design elements are as follows:

1. A menu allowing the user to start a new schema, save and load the current schema, and edit the languages available in the schema.
2. A page selector, allowing the splitting of pathways into multiple pages. This can greatly enhance the ease of editing complex schemas, as each individual pathway can be assigned its own page.
3. A language switcher, allowing the switching between different language versions of the schema and allowing the creation of multi-lingual schemas that share a logical flow.
4. The main schema view. It contains a graphical representation of the elements of the schema, and arrows displaying the logical connections between the different elements. The view can be zoomed and panned around, to focus on specific elements or get an overview over the entire schema. By right clicking an empty

- spot on the canvas, the user gets the option to add a new question block or information block.
5. The right pane menu. Here, the user can switch between three views:
 - a. The search view – This allows the user to search the entire schema to identify certain elements. This can be useful to locate an element in large schemas with hundreds of elements.
 - b. The preview view – This allows the previewing of how the currently selected element will look to a user of the JusticeBot. The preview uses the same styling as the front-end and can be interactively navigated by clicking the individual answers.
 - c. The Edit view – This is the most important view, that allows the editing of questions and information blocks. This is the view you can see in Figure 27.
 6. The “Edit Question” menu. Here, you can see the different sub-menus to edit the currently selected question block. The “text” tab (which is opened in the screenshot) allows the editing of the question block title and description. The “question” tab allows the editing, adding, removing and reordering of the possible answers to the question. The “cases” tab allows the attachment of cases to the question block (see below under 7.3.2).
 7. The title – this textbox allows the editing of the title of the current question block. Any changes made in this textbox are immediately reflected in the schema on the left.
 8. The description – this textbox allows the editing of the description field of the question block. It uses a WYSIWYG (What You See Is What You Get) editor that allows the addition of formatting (such as bolded and italics text), bulleted and numbered lists, and links to external resources. No coding is required to create clear and easily legible descriptions explaining the different legal criteria in simple terms.

Figure 27 shows the view of the JusticeCreator if a question block is selected. If an information block is selected, a similar editing view is shown, however the answers are not editable.

7.3.1.5.2 Creating a new schema

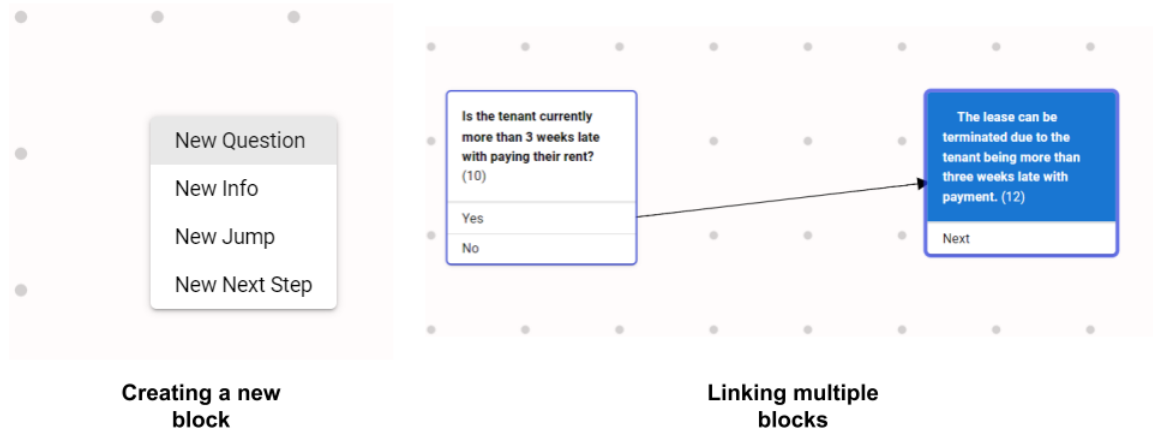


Figure 28 - Creating a new block and linking blocks in the JusticeCreator

The JusticeCreator contains the necessary tools to rapidly create new schemas, or edit existing schemas. When starting a new schema, the user is presented with an empty canvas. Right-clicking anywhere on this canvas brings up a context menu allowing the creation of a new question block or information block (see Figure 28). Once created, this block can be edited (to correspond to the legal criterion the user wants to encode) and dragged around (to be arranged in a logically coherent manner). Further, new blocks can be added. In order to link the two blocks together, the user simply drags an arrow from the appropriate answer in the first block to the second block (see Figure 28). This logically connects the blocks, allowing the encoding of the logical structure as described above.

The steps of creating new blocks and linking the blocks together may seem simple and can be learned in minutes by anyone with basic computer skills. However, it allows the creation of complex and explainable legal reasoning pathways. For reference, the published version of the JusticeBot targeting landlord-tenant disputes contains multiple legal pathways across multiple questions and is made of over 300 question and information blocks. The overall structure of the pathway can be seen in Figure 29.




Figure 29 - The legal pathway of the JusticeBot TAL

The visual representation of the pathway gives an intuitive overview over the structure of the legal reasoning pathway, making it easy for the creator to understand and reason about the path they are currently working on.

7.3.1.6 Conclusion

We have now seen how the legal reasoning structure can be represented in a DAG schema, and how this can be done in an intuitive way in the JusticeCreator software. Next, I will explore how cases are represented in the JusticeBot framework. After that, I will

concretely discuss how these tools can be used to create a legal decision support tool in a new legal area.

7.3.2 Representing court cases

We have discussed how the rules governing an area can be encoded in a graph structure. The graph encodes how the judge will deductively arrive at certain outcomes, based on the application of the legal criteria. This is equivalent to the logical reasoning with rules described above in 4.5 described above. However, the graph does not encode information relating to reasoning described in 4.4 (i.e. whether a certain factual situation will be seen to fulfill a certain legal criteria) or 4.6 (i.e. the concrete outcome a judge will decide upon).

Both of these properties are important for the user of the JusticeBot. Users are unlikely to reason in terms of legal criteria, but rather think of their situation in terms of facts.

Therefore, the JusticeBot needs a way to support the user in understanding whether their factual situation may trigger certain legal criteria. Likewise, a user is likely to want to know the concrete, real-world outcomes that a judge may award if their case goes to court. This can help provide context for whether they wish to introduce their case to the judicial system, or what a reasonable amount to ask in a settlement may be.

However, both of these reasoning processes are, as discussed above, difficult to carry out using artificial intelligence methods. Beyond the issues of data collection, it is very possible that new facts, policy decisions or issues with the annotation schema can cause the prediction to be wrong for individual cases. Further, providing the user with predictions may run afoul of the rules against non-lawyers giving legal advice prevalent in many jurisdictions.

In the JusticeBot, therefore, I chose to refrain from predicting these properties. Rather, the system indexes case decisions based on their relevance for the user and provides the user with the information so that they can make the determination of how a judge would

assess their facts and which outcome they can expect for their own case, using the examples of case law provided.

Each case has three separate uses:

- Cases are used to create the schema, and empirically validate its accuracy, in a way that resembles the Grounded Theory Method. The schema aims to encode the way judges tend to factually deal with cases of a certain nature. As cases are encoded, each case serves as a tool to verify and refine the correspondence between the schema and the legal reasoning steps employed by judges. This use of cases is discussed below in 7.3.2.1.
- Cases are used to illustrate how a judge applied a legal criterion in previous cases. This corresponds to the reasoning step where the judge assesses whether a certain set of facts correspond to an open-textured legal criterion or not, described above in 4.4. In JusticeBot, cases are summarized in order to illustrate how these criteria are applied, so that they can be used to support the user. I will describe this way of encoding case law below in 7.3.2.2.
- Each case can also be used as an example of the outcome a judge awarded in previous cases. This kind of reasoning corresponds to the style of reasoning described above in 4.6. The outcome of previous similar cases can be shown to a user in order for them to understand the possible consequences they could expect, should their case go to court. This way of representing court cases is described below in 7.3.2.3.

7.3.2.1 Creating and validating the schema

In 7.3.1, I described how the JusticeBot methodology encodes rules, in order to mimic the legal reasoning performed by judges. Case law is a crucial component in defining this schema, to be able to accurately capture the practical reasoning steps that judges tend to perform, in order to arrive at a decision.

To create the schema, the annotator might thus start to read the relevant legislation concerning a certain situation, and a few relevant cases. As the annotator sees which

criteria were applied by the judge, they can add these reasoning steps to the schema in the JusticeCreator. For example, an annotator might find that judges tend to first ask whether a tenant is frequently late with paying their rent, and if that is the case examine whether this lateness causes frequent lateness to the landlord, and if so terminate the lease. This shows that these criteria are applied in this order in practice, and should thus be encoded into the JusticeCreator.

While in some areas the reasoning might very closely mimic the criteria set out in the legislation, in other areas the reasoning may be performed differently. For example, when building the JusticeBot focused on landlord-tenant disputes, we analyzed the area of bedbug infestations. The legislation sets out general requirements for liability but does not specify particulars. In order to handle the situation of bedbug infestations, judges introduced a presumption of non-liability for the tenant¹⁰⁸² and assessed whether the landlord had been diligent in addressing the situation, as well as whether the tenant had cooperated with their efforts. None of these requirements stem from the legislation, yet they are crucial to help a layperson understand the consequences of going to court.

In such areas, where the court cases expand upon the legislation, the annotator should use something like the Grounded Theory Method to explore the area. By reading and encoding cases, they will see patterns emerge, that show the criteria that judges apply, and the outcomes that they order. Each case thus represents a puzzle piece to capture the overall reasoning schema that judges apply in solving a certain type of case.

As more and more cases are encoded into the system, these cases serve to *validate* the encoded schema. If the schema does not capture the correct logical connections between the different legal criteria and legal conclusions, the order of the applied criteria, and which criteria are applied by a judge, would not match the encoded schema.

¹⁰⁸² Julie Pomerleau, “Punaises de lit au logement”, (19 April 2018), online: *Blogue SOQUIJ - Actualités juridiques et judiciaires du Québec* <<https://blogue.soquij.qc.ca/2018/04/19/punaises-de-lit-logement/>>.

Likewise, if the legal conclusions are not correctly specified in the schema, the expert annotator will notice that the outcome ordered by the judge frequently does not match what is expected according to the schema.

Noticing these discrepancies can lead to the discovery that judges reason differently about legal areas than expected. The expert annotator can update and refine the schema based on this information, in order to more accurately capture the steps judges in reality apply to reason about cases.

In encoding cases into the system, the expert annotator thus empirically validates that the schema works and refines the schema. As more and more cases are added, and are coherent with the schema, the likelihood that the schema accurately captures the reasoning steps performed by judges increases. Instead of being based exclusively on the interpretation of a legal area by an expert annotator, the schema is thus also verified by constantly applying it to encode legal cases. This increases the legitimacy of the system.

After and during the cases are used to capture and verify the reasoning schema, they also serve another purpose. For each legal criterion that is treated in a case (e.g. “frequent lateness”), each case serves as an illustration of how judges tend to reason around this criterion. The next section looks at this purpose of case law in the JusticeBot methodology.

7.3.2.2 Illustrating the application of legal criteria

In each judicial case, a judge has to apply the legal rules to arrive at an outcome, based on the facts present in a case. In doing so, they can be seen to choose a path through the rules governing a certain legal area, by assessing whether each criterion applies, and if so move on to the next legal criteria. They will do so until they have established whether a certain claimed outcome can be awarded, and the case succeeds, or whether the claim should be rejected.¹⁰⁸³

¹⁰⁸³ See 4.5.

In doing so, the judge applies the criteria as they relate to a single unique case. However, as described above,¹⁰⁸⁴ there is likely to be a form of consistency in cases that deal with similar issues. We expect cases with similar facts to also have similar outcomes. For example, if an apartment that has issues with heating for five days is considered unfit for habitation in one case, we would in general expect subsequent cases where there is a similar heating issue to also be seen as being unfit for habitation.

Based on this insight, it would be useful to provide the user with a summary of the reasoning of a judge relating to a legal criterion, such as “unfit for habitation”. The user can then read this summary of the case, decide whether the facts of their case match the previous case, and if so hypothesize that the judge would apply the legal criterion in the same way in their case.

For these purposes, for each case we care about the following information:

- Which criterion did the judge assess?
- Did they find the criterion to apply in the case?
- Why did they come to this decision?

The case can thus be abstracted as shown in Table 4.

Table 4 - Way of abstracting cases by capturing assessed criteria and reasoning

Criterion	Does it apply?	Why?
Criteria 1	Yes/No	The factual situation that the judge saw as relevant to determine whether criteria 1 applies or not.
Criteria 2	Yes/No	The factual situation that the judge saw as relevant to determine whether criteria 2 applies or not.

¹⁰⁸⁴ See 4.9.3.

The summaries of the cases in the JusticeBot are thus not general summaries of the entire case. Rather, they are summaries that focus on understanding why a specific legal criterion was applied in a certain way by a judge. Each case can have a number of different summaries, each relating to different legal criteria.

7.3.2.2.1 Example case 1

Let us take a look at two example cases relating to the requirements around frequent lateness of rent. For reference, here is the graph representation of the reasoning steps a judge must take relating to this issue:

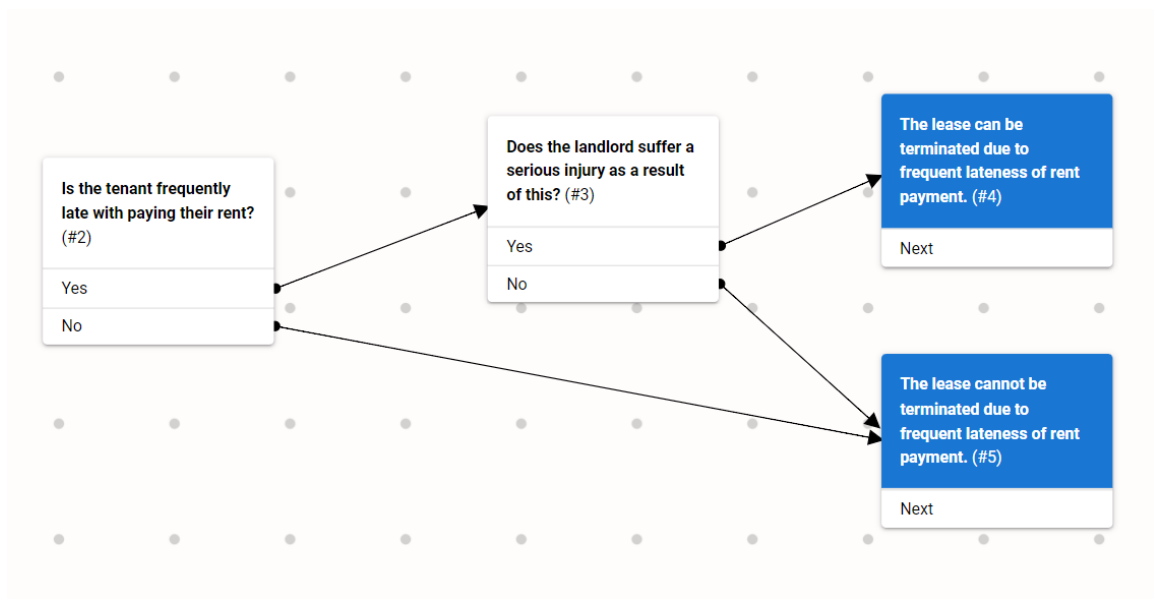


Figure 30 - Graph form of relevant legal criteria for determine lease termination due to frequent lateness of rent payment

Table 5 shows how an abstracted version of a case relating to the frequent lateness of rent payment.

Criteria	Does it apply?	Why?
Is the tenant frequently late with paying their rent?	Yes	The tenant had paid their rent late on 7 occasions in the past 9 months.

Does the landlord suffer a serious injury as a result of this?	Yes	The landlord was unable to carry out necessary repair work due to lack of money.
--	-----	--

Table 5 - An example of a case abstracted in JusticeBot

This can be seen as the judge “traversing” the logical path of the graph in the following way:

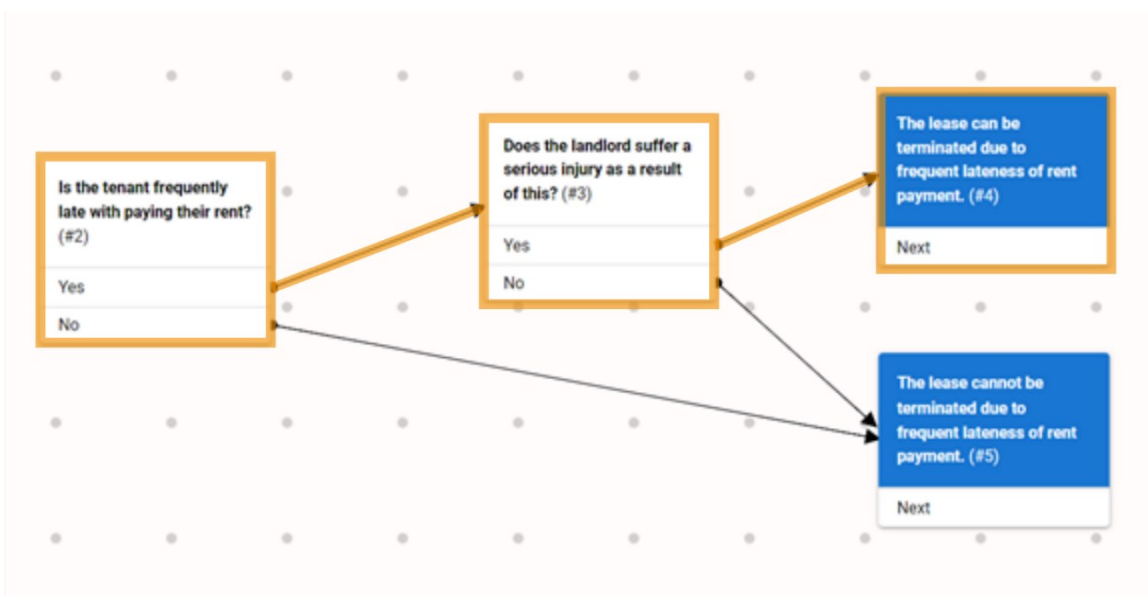


Figure 31 - The reasoning path by a judge taken in Table 5

7.3.2.2.2 Example case 2

On the other hand, a judge might find that the tenant is, in fact, not frequently late with paying their rent. This case could be represented in the following way:

Criteria	Does it apply?	Why?
Is the tenant frequently late with paying their rent?	No	The tenant had paid their rent late only twice in the past seven months.

Table 6 - An example of a case abstracted in JusticeBot

Since the tenant is not seen as being frequently late with paying their rent, the serious injury caused to the landlord does not have to be assessed. The lease cannot be terminated

based on the frequent lateness of paying rent by the tenant. In the graph structure, this reasoning looks like this:



Figure 32 - The reasoning path by a judge taken in Table 6

7.3.2.2.3 Which cases should be picked?

The purpose of these case law representations is to illustrate how legal criteria are applied in real-world cases. These summaries will then be shown to the user in order to support them in assessing their own case. Since the user is unlikely to read through many case summaries, and in the interest of keeping the information provided brief, we will need to pick a few case summaries to show to the user for each individual legal criterion. This makes it necessary to select the most helpful cases for each criterion. In the JusticeBot TAL, we decided to show a maximum of five cases where the criterion was found to apply, and five cases where the criterion was found not to apply.

Exactly which cases should be picked is likely to differ between the individual instances of legal decision support tools. However, there are a few points that may be relevant in selecting cases to summarize:

- Cases should be commonplace cases. Above, we discussed the difference between commonplace and landmark cases.¹⁰⁸⁵ Landmark cases are very important, as they establish new precedents or deal with novel factual situations. In the JusticeBot framework, they may therefore be included in the description below the question. However, the case summaries are used to illustrate how the judges *typically* apply the legal criteria. Therefore, they are more likely to be commonplace cases, that show how the criteria is generally applied to typical situations.¹⁰⁸⁶
- Cases are likely to cover common situations. Even for relatively typical legal criteria, there may be an immense number of possible factual situations that may lead to the decision that a criterion applies or not. Due to the limited number of cases that we can show to the user, the selected cases should ideally correspond to common situations. This makes it more likely that the situation of the user is similar to one of the presented case summaries.¹⁰⁸⁷
- Individual cases should be based on varied sets of facts. If the user is presented with five case summaries that essentially say the same thing (e.g. that a cold apartment is unfit for habitation), the summaries are less likely to correspond to the situation of the user. Instead, one of these cases should be presented, and other cases focusing on other factual situations, to maximize the chance that a situation similar to that of the user is present in the list of case summaries. In essence, we are trying to cover the greatest possible area of a "decision space", i.e. the factual situations that can affect a certain legal criteria one way or the other.
- Cases that are used to illustrate a certain legal criterion should contain an explanation of why the judge considered this criterion to apply or not. Cases where the judge merely confirms that a certain criterion applies or not, without elaborating why that is the case, are less useful to summarize, since they cannot help the user in understanding their own situation. This may be the case if the

¹⁰⁸⁵ See 4.9.3.

¹⁰⁸⁶ Compare 5.4.5.1.

¹⁰⁸⁷ Compare 5.5.2.

situation is obvious, or the applicability of a criterion is undisputed by both parties.

7.3.2.2.4 How should the cases be summarized?

Another important consideration is *how* the cases should be summarized with regard to the individual criteria. As we saw above, the summaries for each criterion are triplets, containing the following three datapoints:

- Which criteria did the judge apply?
- Did they find it to apply in the case?
- Why did they come to this conclusion?

The criterion itself, and whether the judge found that they applied are relatively obvious to summarize. The criterion stems from the schema, as the judge reasons about certain legal criterion imposed by the law. Usually, the decision will show whether they found a criterion to apply or not, although as we have seen this is not always clear.

The third element, *why* a judge came to a certain conclusion regarding the criterion, is a bit trickier to annotate. The annotator must decide how to summarize the reasoning of the judge in a way that makes the summary useful to see for the end-user of the system.

While, again, the specific way may depend on the specific use-case of a system and the legal area, I will discuss some general tips on how cases should be summarized. The created summaries should follow a few important rules:

- Summaries should be short. The user will be presented with up to ten summaries and may not spend a lot of time on reading the individual summaries. Therefore, the easier a summary is to skim, the better.
- Summaries should summarize the relevant factual situation that led a judge to decide whether the criteria applied or not. For the summary to be useful to the user, it should aim to summarize the facts that the judge found relevant in determining whether a criterion applied or not. As discussed, the layperson user is

likely to think of their situation in terms of facts.¹⁰⁸⁸ The summary must thus contain the facts that were important for an outcome, so that the user can compare these facts to their own situation.

In general, the following schema seems to correspond to the aforementioned requirements:

The judge found that the {criteria} {applied/not applied}, since {summary of the facts}.

For example, the following summaries could be useful for a user:

- *The judge found that an apartment was unfit for habitation, since the heating of the apartment was broken for five days in winter.*
- *The judge found that the tenant was not frequently late with paying their rent, since they had only been late two times in the past twelve months.*
- *The judge found that the landlord did not deal diligently with the infestation of bedbugs, since they only responded to the situation two weeks after being notified and tried to combat the infestation themselves instead of hiring a professional.*

7.3.2.2.5 What happens if cases do not match the schema?

In some instances, the reasoning process followed by the judge might not match that encoded in the schema. Perhaps, a party raises a defense that stems from another legal rule. For example, a tenant may argue that they were only late with paying their rent since the landlord had stolen their money. The judge may then consider whether this is accurate, which does not match the "expected" reasoning path, according to the schema encoding of the cases.

In some cases, this may mean that a schema needs adjustment. If a criterion often comes up in certain cases, the schema should be updated to include it as an additional criterion, as discussed in 7.3.2.1.

¹⁰⁸⁸ See 5.2.3.

Other times, the case may simply not correspond to the normal reasoning patterns in an area. Dealing with these kinds of cases is tricky and will be discussed below in future work. However, such divergent cases may be less common in areas of high-volume, low-intensity disputes.

7.3.2.2.6 Conclusion

I have described how cases can be captured in order to illustrate the application of legal criteria. Once enough cases are captured in this way, they can be shown to the user to support them in understanding how judges might apply legal criteria to their cases.

However, cases can also have another use – they can demonstrate the outcome of cases that are similar to that of the user. In the next section, I will explore this use of case law.

7.3.2.3 Illustrating the outcome of previous cases

In the JusticeBot, cases can also serve as an illustration of the outcome that a judge ordered in previous cases. As described above, this can help the user assess the potential outcome of going to court and provide a BATNA for negotiations.¹⁰⁸⁹

In prior work, we saw projects that aimed to predict the discretionary outcome that a judge awards based on the factors present in a case. Cases are seen as similar if the facts that appear in a case are similar. The JusticeBot takes a different approach. Here, cases are seen as similar if the judge finds that the same legal criteria apply. The user is asked which legal criteria they believe apply in their case, supported by the case law summaries described above. At the end, they are shown cases where the judge decided that the corresponding legal criteria applied. This comparison between cases will be described more in-depth below in 7.5.2. For now, let us analyze how cases can be represented in order to illustrate the outcome of cases.

¹⁰⁸⁹ See 5.4.1.2.

7.3.2.3.1 Why use case law to illustrate outcomes?

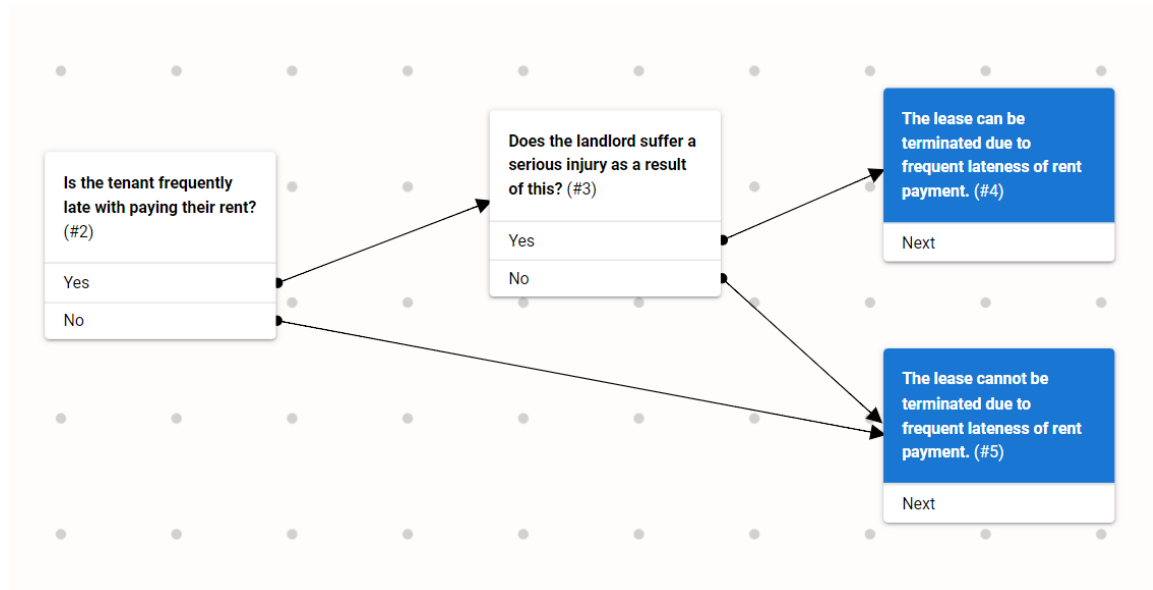


Figure 33 - Reasoning schema for cases regarding frequent lateness of rent payment

The first question we should explore is what the usefulness is of using cases to understand outcomes at all. Let us turn once more to the reasoning graph regarding frequent lateness, shown in Figure 33. Depending on whether the judge finds that the criteria apply, they will always arrive at a blue block, that contains information about potential outcomes, i.e. whether the lease can be terminated or not. Perhaps, providing this information to the user would be enough for them to understand the outcome?

I believe that merely providing this answer, as provided by the rules encoded, is not enough. I will discuss three arguments for why case law additionally needs to be provided to the user in order for the system to be the most useful.

7.3.2.3.1.1 Divergence from the expected outcome

A first reason for supplying the user with case law is that the judge may diverge from the outcome that is expected according to the schema. In some cases, they may be able to change the outcome in a discretionary manner. For example, article 1973 of the Code Civil du Québec allows the judge to order the party in breach of the legislation to perform their duty by a certain date, instead of terminating a lease. Even if the criteria for lease termination are fulfilled, the judge may thus choose not to terminate the lease after all.

Merely providing the user with the information that a lease can be terminated may therefore not capture the situation fully.

Instead, providing the user with previous case law gives them the possibility to evaluate what outcomes were *actually* achieved, rather than which outcomes may be awarded according to the law. The former is likely to be more interesting to a user aiming to determine a BATNA.

7.3.2.3.1.2 A discretionary outcome

In some cases, the resulting legal conclusion may include a discretionary component. For example, a user may be awarded damages for an event, the assets may be split between the parties in a divorce, or a defendant may be sentenced to prison for a certain number of years.

The legislation encoded in the schema does not provide information about how these decisions will be made – they are discretionary decisions taken by the judge.¹⁰⁹⁰ Instead, showing examples from previous case law to the user can help them get an overview over the ranges of outcomes that they can expect.

7.3.2.3.1.3 Augmented intelligence and legal information

Providing the user with case law instead of outcomes determined solely by the schema is also more well-aligned with the notion of augmented intelligence instead of artificial intelligence and providing legal information instead of legal advice.

Only providing information from the schema relies on the legal expert user interpreting the legal rules in an area, and encoding it correctly into a computer system. This step is not always easy, as we have seen.¹⁰⁹¹ Further, such a system conceptualizes the system as an artificial intelligence, that takes the information of the user and gives them a definitive answer for what the judge will say in their case. As we have seen in the previous chapters, reality often does not neatly fit into digital boxes, and the reasoning performed

¹⁰⁹⁰ See 4.6.1.

¹⁰⁹¹ See 4.5.3.3.4.

by judges in determining the outcome of a case may be beyond the current scope of artificial intelligence.

Providing the user with case law can lessen these concerns. The system will still contain information about the potential outcome as described by the law. However, this information is complemented by the encoding of case law. Instead of telling the user what the system believes their outcome could be, they are merely provided with a list of previous cases and their outcomes. The user can draw their own conclusions from this presentation. Thus, the system works more like an augmented intelligence system, providing the user with the tools to enhance their understanding of their own situation.

This way of dealing with cases changes the entire focus of the system. It can be seen as a sophisticated search engine, that searches previous case law for cases that are similar to that of the user and provide them with this information. Merely surfacing these cases is unlikely to be considered giving legal advice.

Now that we have determined that case law is useful in illustrating the outcome of cases, I will describe how specifically the outcome of the cases is encoded in the JusticeBot methodology.

7.3.2.3.2 How to encode the outcome of the cases

In the FactorBot, cases were encoded in terms of the applied rent reduction and moral damages. Encoding cases in this way turned out to be somewhat difficult, as it was not always clear whether the monetary amount awarded was a rent reduction or other types of damages. Further, in general, the outcomes awarded in cases may be divergent and depend on the legal area of the case. The outcomes of a single case may contain a mix of monetary awards and other orders for the parties. The outcome often depends heavily on the claim of the plaintiff – the judge will not award more damages than claimed by the plaintiff, for example.

Building a system to capture these outcomes in a quantitative way, that is also generalizable to multiple case types, would require significant work.

Therefore, in the JusticeBot, the outcomes of cases are represented as brief textual phrases detailing the outcome of the case. This maintains the flexibility of plain text in describing the outcome of the case, while providing the user with a useful way to examine the possible outcomes of cases.

Below are a few examples of how the outcomes of cases could be represented:

- The judge terminated the lease and expelled the tenant.
- The judge ordered the landlord to pay the tenant damages of 300 CAD, and ordered a rent reduction of overall 500 CAD.
- The judge did not terminate the lease, but ordered the tenant to pay their rent on the first day of the month in the future.

7.3.2.3.3 Global or local outcomes?

An important question is if outcomes of cases should be tracked on a global or a local case basis. These have different implications for how the outcomes of cases can be compared to the case of the user.

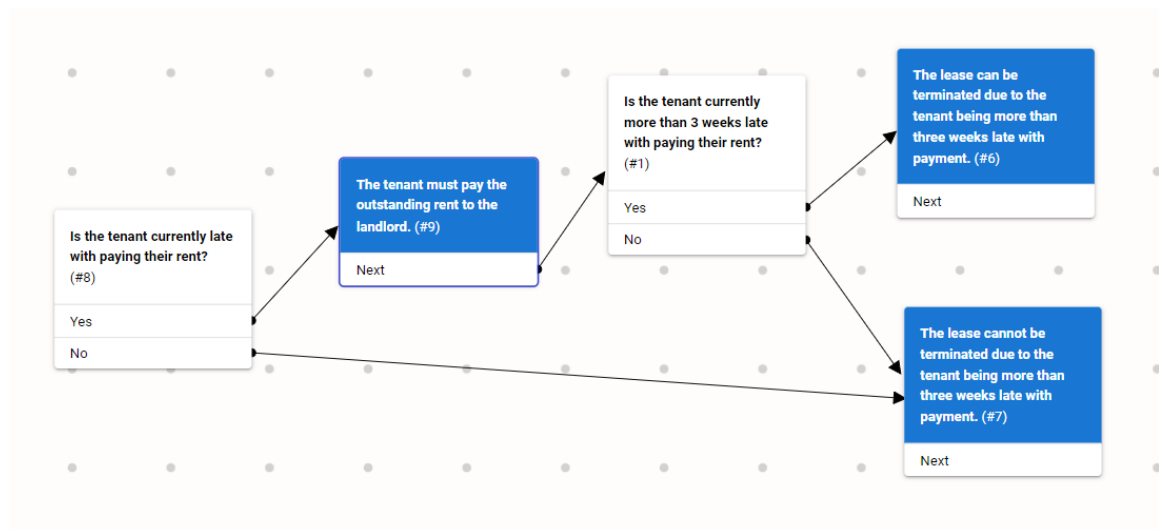


Figure 34 - Path with intermediary conclusion

By encoding the outcome *globally*, I here refer to attaching the encoded outcome to the entire case. In the case of a decision around the criteria presented in Figure 34, the decision together with its outcome would be represented in this manner:

Criteria (excluding summaries)

Criterion	Does it apply?
Is the tenant currently late with paying their rent?	Yes
Is the tenant currently more than three weeks with paying their rent?	Yes

Outcome

The judge ordered the tenant to pay the rent, and terminated the lease.

This kind of abstraction captures the reasoning steps of the case, and the overall outcome of the case.

However, some interesting information is lost in capturing cases in this way. It does not capture whether any of the outcomes are linked to any specific criterion. As we can see in Figure 34, there are two legal conclusions in this pathway. Each of these may have individual outcomes attached to them. For example, the judge ordering the tenant to pay the rent in this case stems from the fact that the tenant is late with paying their rent at all. In the JusticeBot, therefore we capture cases in the following *local* manner:

Criteria (excluding summaries)

Criterion	Does it apply?	Outcome
Is the tenant currently late with paying their rent?	Yes	The judge ordered the tenant to pay the rent
Is the tenant currently more than three weeks with paying their rent?	Yes	The judge terminated the lease.

The outcomes are here attached to the individual legal conclusions. As we will see below, this can enable some interesting ways to compare cases that are partially matching, thereby decreasing the number of cases that need to be annotated.

We have now seen how the JusticeBot methodology represents case law. Below, we will explore how these representations can be analyzed to support the user in understanding their case. First, let us explore how the JusticeCreator can be used to practically encode cases into the system.

7.3.2.4 Encoding cases in the JusticeCreator

The JusticeCreator contains the necessary functionality to encode cases in the aforementioned format. Here, I will briefly describe the steps to select relevant cases and encode them in JusticeBot format. This assumes that the schema itself has already been created, as described in 7.3.1.5. Let us assume that we want to add cases to the schema displayed below in Figure 35.

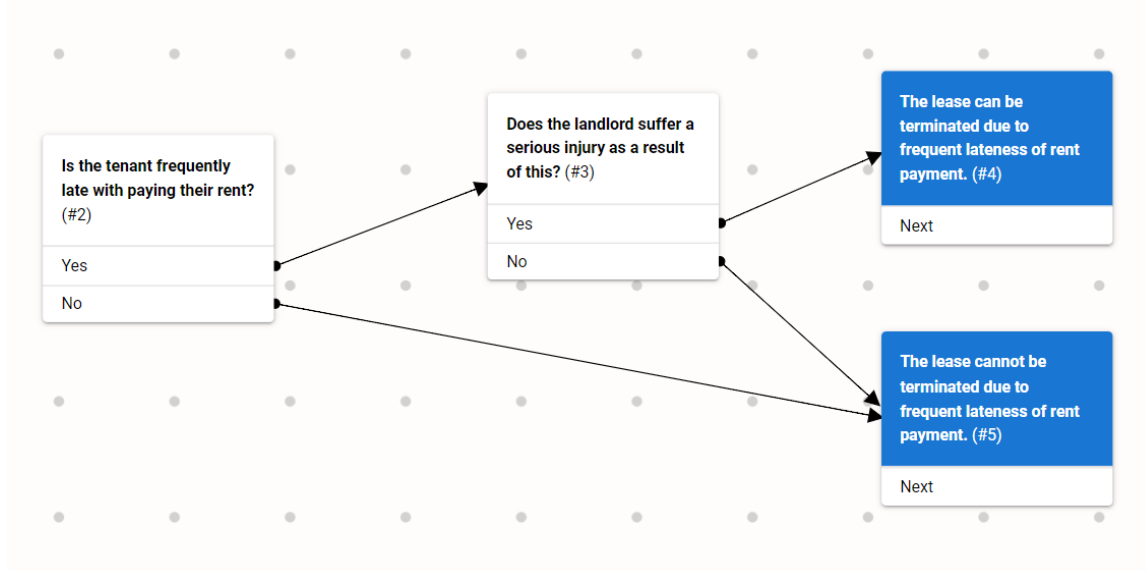


Figure 35 - Reasoning schema for cases regarding frequent lateness of rent payment

7.3.2.4.1 Identification of relevant cases

The first step to encode cases is to identify the cases to encode. One simple way to accomplish this is to search for the case in a publicly available database to find the cases.

In this case, relevant search terms could be "frequently late" or "serious injury late". This will result in a list of cases that deal with issues relating to these criteria.

7.3.2.4.1.1 *Automatic retrieval of relevant cases*

Searching for cases and identifying the correct search terms to find cases relating to a specific issue can be a significant time sink. Therefore, the JusticeCreator offers a functionality that uses artificial intelligence to support this task.

This system works by analyzing case law in bulk, and building an index capable of surfacing cases that contain sentences that are semantically similar to a certain legal criterion. The methodology is based on research I conducted with an international group of research collaborator that enabled us to achieve the best score of the case retrieval task in the Competition on Legal Information Extraction/Entailment (COLIEE) 2020,¹⁰⁹² and win the "Best Paper Award" at JURIX 2020.¹⁰⁹³

I implemented this system for cases dealing with landlord-tenant disputes. The program works by splitting the cases that we obtained from the TAL (see 8.3.2) into sentences. I selected the cases from 2017 to 2020, and used a regular expression algorithm to split these into individual sentences. In total, this results in around 1.8 million sentences.

Each of these sentences is then embedded into a vector format, using a pre-trained deep learning encoder known as the Universal Sentence Encoder (USE). This model creates embeddings of sentences that capture the semantic meaning of a sentence.¹⁰⁹⁴ For example, sentences that have similar meanings but use different vocabulary would result in similar embeddings using this system.

¹⁰⁹² Hannes Westermann, Jaromir Savelka & Karim Benyekhlef, "Paragraph Similarity Scoring and Fine-Tuned BERT for Legal Information Retrieval and Entailment" (2021) *New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science)* 269–285.

¹⁰⁹³ Westermann et al, *supra* note 195.

¹⁰⁹⁴ Yinfei Yang et al, "Multilingual Universal Sentence Encoder for Semantic Retrieval" (2019) arXiv, online: <<http://arxiv.org/abs/1907.04307>> arXiv:1907.04307 [cs]; "Universal Sentence Encoder Multilingual", online: <<https://tfhub.dev/google/universal-sentence-encoder-multilingual/3>>.

Then, I trained a fast nearest neighbor search algorithm to identify sentences that are similar to a given sentence. For this, I used the Spotify Annoy library.¹⁰⁹⁵ This system builds a search tree that allows us to give it the embedding of a sentence. Then, it returns the most similar sentences across the corpus of 1.8 million sentences in a few milliseconds.

This functionality can be used to surface cases that may be relevant to a certain legal criterion, to support the annotator in identifying cases. To achieve this, the title in each legal criterion is embedded into the same vector embedding as the sentences from all of the cases. Then, the nearest neighbor search is used to retrieve the 100 sentences from the cases that are the most similar to the title of the legal criterion. These sentences, and the cases that contain the sentences, are suggested to the legal expert as cases that could be relevant for encoding into the system.

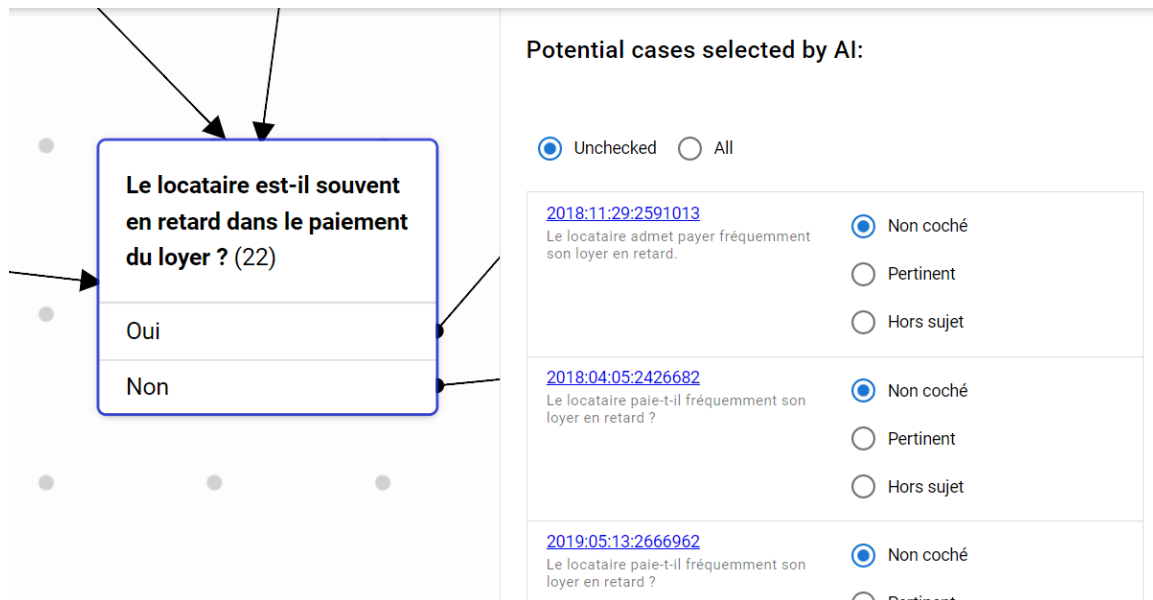


Figure 36 - AI case suggestion view in the JusticeCreator

Figure 36 shows how this is presented in the JusticeCreator interface. Upon clicking on the relevant question block (in this case, “Is the tenant frequently late with paying their

¹⁰⁹⁵ *spotify/annoy* (Spotify, 2022).

rent?”), the user is presented with a list of cases retrieved by the AI methodology described above. Each link contains a URL to read the case, and the sentences from the case that was matched. As we can see in the image, the algorithm has correctly retrieved cases mentioning “frequently”, even though the term used in the criterion title is “often”, showing the strength of the pre-trained model to understand the meaning of a sentence, rather than the specific vocabulary.

The expert user can choose to read the case and see whether it is relevant for adding it as a summary to the legal criterion. Further, the expert user can specify in the system whether the case is relevant or not. This information can be used to evaluate and improve the system.

As we can see, this kind of system is a clear example of augmented intelligence. It instantly surfaces potentially relevant cases from hundreds of thousands of documents, making the task of selecting cases much more efficient. However, the final decision of whether a case should be annotated, and how to annotate it, remains with the legal expert.

7.3.2.4.2 Encoding cases

Once the cases for annotation have been selected, the next step is to perform the annotation itself. This is done via the JusticeCreator interface, which allows the addition of case law summaries to both question blocks and information blocks.

In order to encode a given case, the legal expert opens the case in parallel with the relevant pathway in the JusticeCreator on their computer. They then read the case, aiming to identify the relevant reasoning steps carried out by the judge. Then, they follow the reasoning by the judge, simultaneously in the case and in the schema. This way, the relevant reasoning process, and outcomes tied to legal conclusions, can be added to the schema.

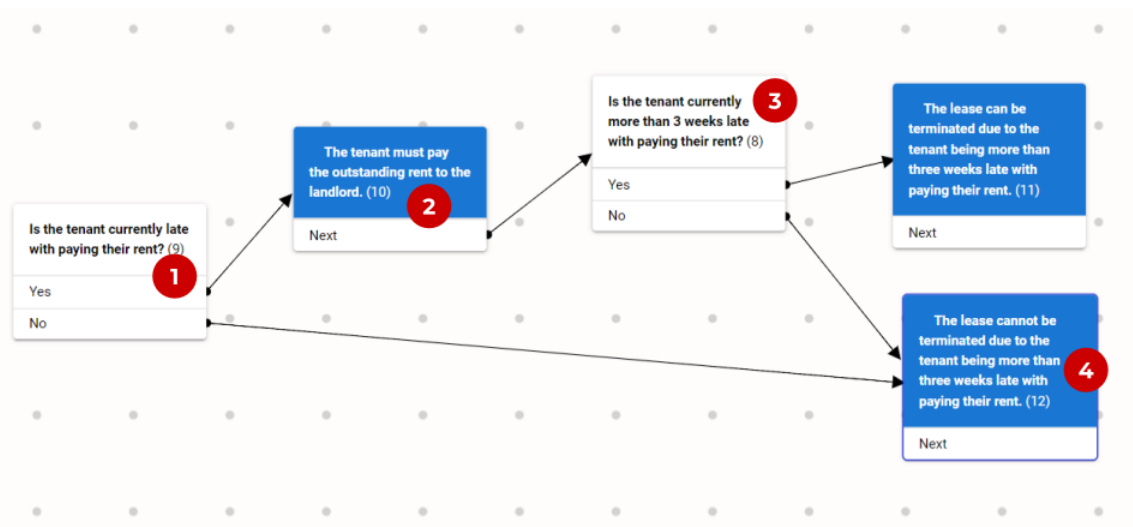


Figure 37 - Schema related to lateness of rent of more than three weeks

I will illustrate this process with the help of the schema presented in Figure 37. Let us imagine that we want to encode the following case into the system:

A tenant is currently 21 days late with paying their rent. The judge does not consider this to be more than three weeks late. Thus, they order the tenant to pay the outstanding rent, but they do not terminate the lease. They additionally add an order that the tenant must pay their rent on the first day of the month in the future.

We will now step through the different criteria in the schema and case and add the relevant annotations to the schema.

7.3.2.4.2.1 Legal Criterion (1) – Is the tenant currently late with paying their rent?

The first criteria presented in the schema is the question whether the tenant is late with paying their rent or not (See (1) in Figure 37). This seems to be the case here. However, the case is not very useful to illustrate the application of the legal criterion, since the judge does not explain how they reasoned, and both parties agree that the tenant is late. Instead, since the judge found the criterion to apply, the legal expert follows the arrow connected to the “Yes” option, arriving at the information block at (2).

7.3.2.4.2.2 *Legal conclusion (2) – The tenant must pay the outstanding rent to the landlord*

Are there any legal consequences that follow from the tenant being late with paying their rent? In the case, we can see that the judge ordered the tenant to pay the outstanding rent, and to pay the rent on the first of the month in the future. This outcome is a direct consequence of the tenant being late with paying their rent. In the schema, the legal expert has encoded the tenant being ordered to pay their rent as a possible consequence in the form of an information block (see (2) in Figure 37). Let us see how the outcome from the case can be added as an example to this information block.

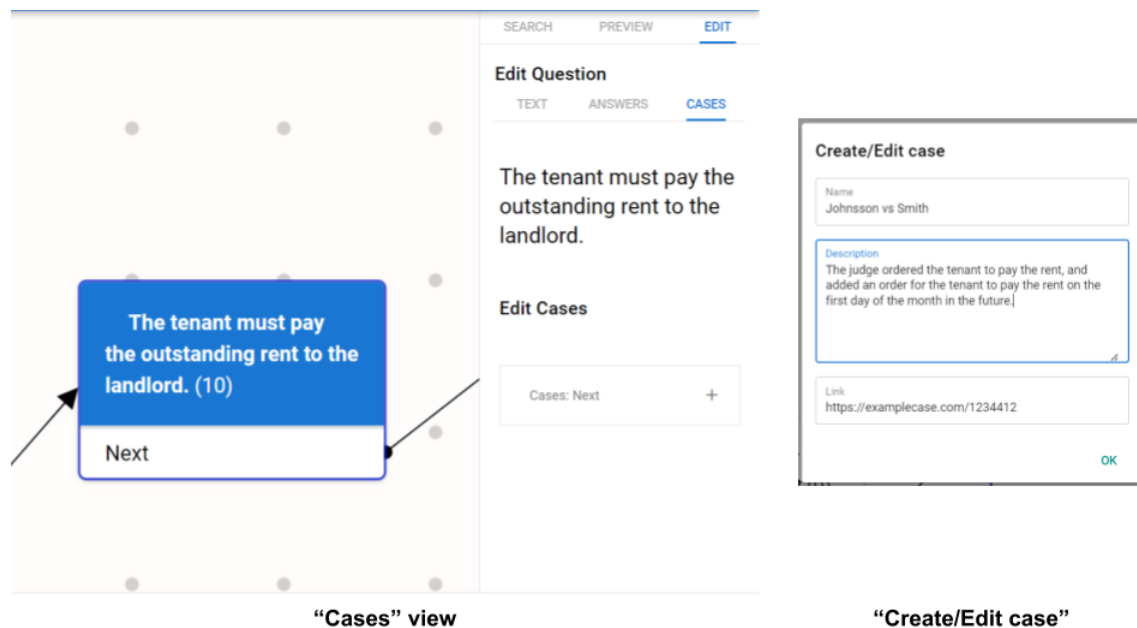


Figure 38 - Encoding case law outcomes in the JusticeCreator

After selecting the relevant information block in the JusticeCreator, the legal expert is given the option to add cases to the block by pressing the little “+” icon (see Figure 38). They can then add the case as shown in Figure 38, by adding the name of the case, summarizing the relevant outcome, and adding a link to the case so that the user can read it for themselves.

7.3.2.4.2.3 *Legal criterion (3) – Is the tenant more than three weeks late with paying their rent?*

Next, the judge discusses whether the lease can be terminated due to the tenant being more than three weeks late with paying their rent. This criterion can be seen in (3) in Figure 37. In our case, the judge decided that the tenant being 21 days late did not qualify as them being more than three weeks late that the legal rules require. This may not be obvious for everyone, therefore, let us add it as an illustration of previous reasoning to the legal criterion, so that users of the system can read a summary of this decision to help them understand their case.

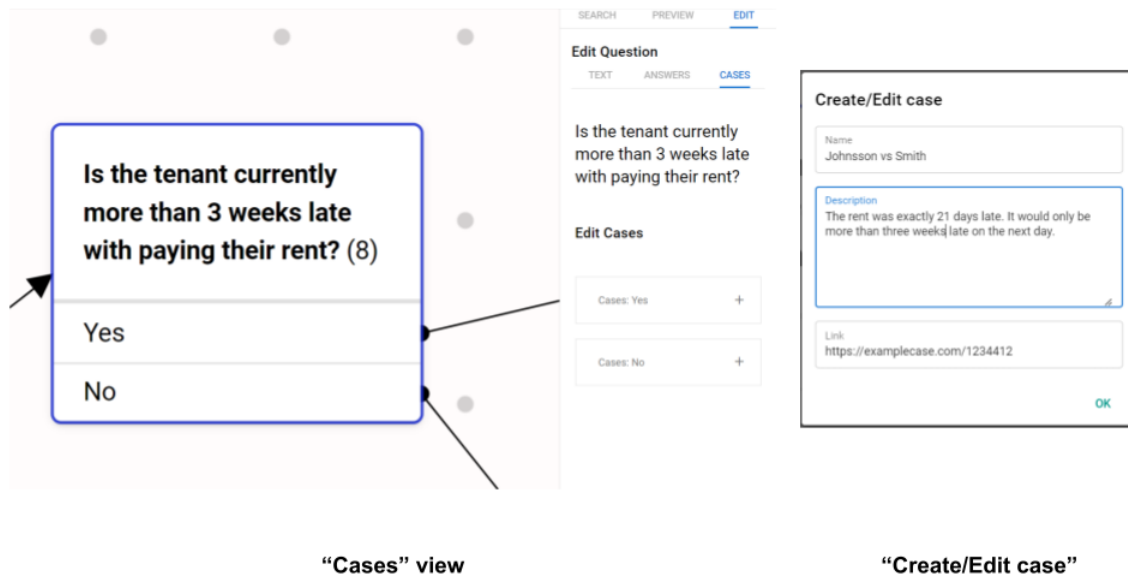


Figure 39 - Adding case law summaries to legal criteria

Figure 39 shows how a summary of this reasoning can be added to the criterion. By selecting the relevant criterion, and navigating to the “Cases” tab, we can see that the summaries that can be added to the case are separated by whether the judge decides that the criterion applies or not. In our case, the judge finds that the criterion does not apply. By clicking the “+”-sign next to the “No” button, the “Create/Edit case” dialog is opened, which allows us to describe the reasoning of the judge. In this case, we can add the information that the rent has to be 22 days late to be considered more than three weeks late, which could be helpful for future users of the system.

Of course, this is merely an example. These case summaries are likely to be even more helpful when criteria are based on open-textured legal concepts, such as “reasonable”, “frequently late”, “unfit for habitation”, etc. Here, the summaries can provide information about which specific factual situations fell under the criterion, and which ones did not.

7.3.2.4.2.4 Legal conclusion (4) – The lease cannot be terminated

Finally, following the schema, since the judge found that the tenant was not more than three weeks late in paying their rent, the final legal conclusion would be that the lease would likely not be terminated. This matches the outcome in our example case, where the lease is not terminated – this portion of the case is rejected. Just as described above in 7.3.2.4.2.2, the expert user would add this outcome as an example to legal conclusion 4, saying that the lease was not terminated. This empirically validates the fact that the lease not being terminated is a possible outcome for the rent payment being less than three weeks late.

The relevant features of the case have now been encoded into the JusticeCreator system. Below in 7.4, we will see how these will be displayed to support the user in assessing legal criteria for their own case and understanding the potential outcomes of their case.

7.3.2.5 Conclusion

We have seen how cases can be captured in the JusticeBot methodology. Each case can have three purposes. First, it can be used to discover and validate the legal reasoning schema. Second, cases can be used to illustrate how judges have previously applied legal criteria. Second, it can be used to exemplify the outcome that a judge ordered in a certain case.

To be used for these purposes, the cases are encoded in terms of how they relate to the schema that was created to capture the reasoning in a certain legal area. The JusticeCreator provides the tools to concretely encode the cases. It has features to automatically retrieve cases that may be useful for summarization, and further allows the adding of case law summaries to the different legal criteria and legal conclusions, using a simple interface.

Once a certain number of cases have been captured in this manner, they can be used to support the user in exploring their own case. This will be described below. First, let us examine some best practices in creating a JusticeBot decision support tool.

7.3.3 A manual for building JusticeBot tools

I have presented the tools that can be used to build a JusticeBot decision support tool in a new legal area. These tools give the legal expert a significant amount of flexibility in how they want to shape the tool, which is important to allow the JusticeBot methodology to adapt to new legal areas. In this section, I will briefly present a few potential guidelines, in the form of a possible manual of how a new legal area could be encoded. This is largely based upon our experiences with building the JusticeBot TAL, which will be described below in Chapter 8, and a few other application areas that are currently under development. While the manual worked well in our case, it is likely that it needs to be adapted to fit the particularities of new legal areas.

7.3.3.1 Identify a suitable legal area

The first step in building a JusticeBot tool is selecting an appropriate legal area for the tool to be built in. In some cases, this may be predetermined by the scope of the pursued project. In other cases, there may not be such a requirement. In this case, it is up to the team to determine the legal area. Above, in Chapter 5, I have described criteria that can help determine whether a legal area is well-suited for a JusticeBot system, such as a legal area being high-volume and low intensity. In Chapter 9.3, I discuss a few relevant legal and administrative areas that may be relevant targets to build JusticeBot systems.

It is also important to consider the target user of the decision support tool. This is important in order to determine the purpose of the tool, and the level of legal and reading comprehension that can be assumed. In this thesis, I mostly focus on building JusticeBot systems for laypeople. In 9.2, I explore how Justicebot systems could be built to target other target users.

Establishing institutional links can be a very important step in building legal decision support tools. In the case of the JusticeBot TAL, we were very lucky to work with the Tribunal Administratif du Logement, with help from Aide Juridique Montreal.

In order to benefit from the AI-case suggestion features described above, it is useful to have access to a corpus of case decisions, that can be scanned and analyzed using artificial intelligence. However, this is not a requirement, as cases can be identified on publicly available legal sources.

7.3.3.2 Build a scaffold schema

Some JusticeBot systems may target a single legal issue. In this case, the pathway can be built immediately. In other cases, there may be multiple different pathways that a user may want to explore, in the same legal area.

In these cases, it is useful to create a scaffold that introduces the user to the system and triages the specific pathway they may want to explore. We have found it useful to introduce an initial screen that introduces the user to the system, explains what they can expect, and informs them that the system merely provides legal information. Upon agreeing with these terms, the user enters a few initial, high-level question blocks. These are not pathways imitating legal reasoning, but rather pose a few initial questions, such as “How can we help you today?”, aiming to allow the user to select their specific issue.

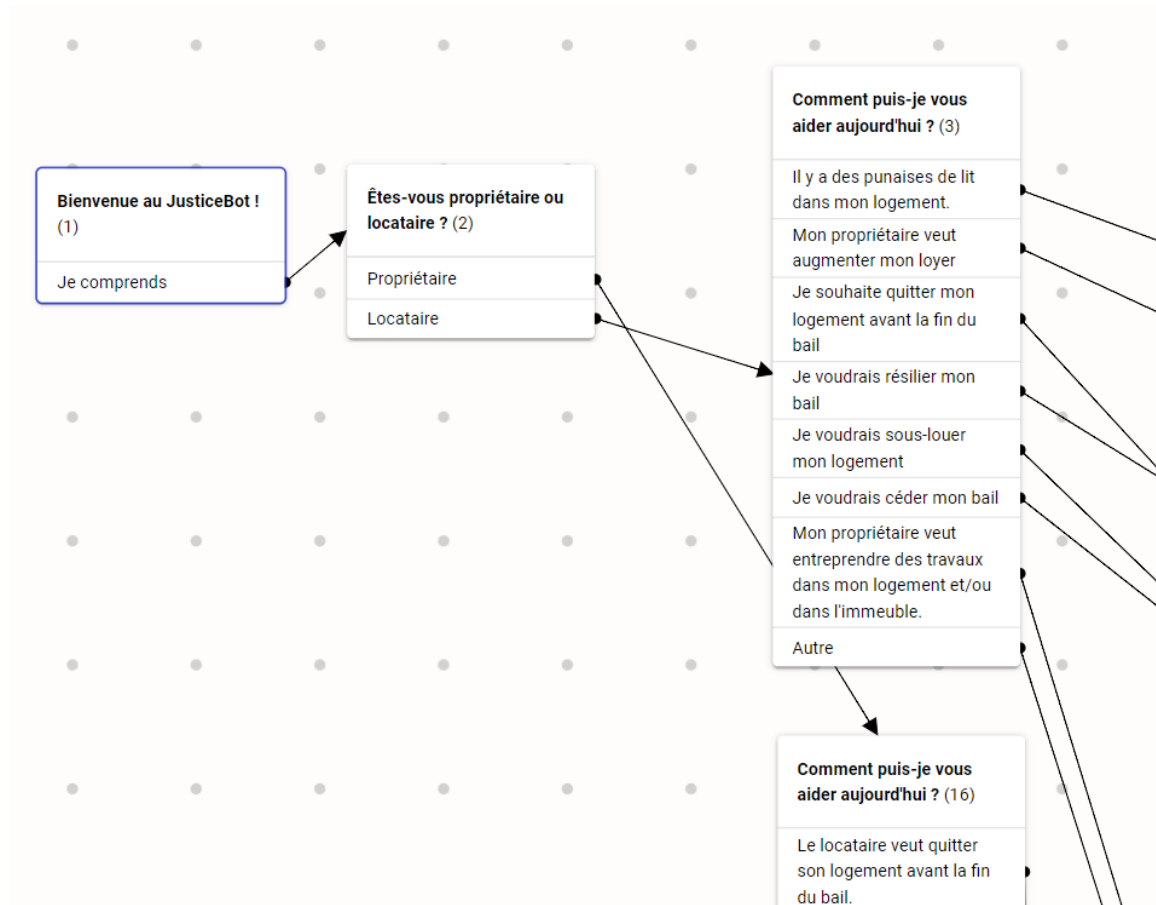


Figure 40 - Initial blocks in the JusticeBot TAL

Figure 40 shows these initial blocks in the JusticeBot TAL. Here, after being provided with information regarding the functioning of the JusticeBot, the user is asked whether they are a landlord or a tenant, since the issues these users face are likely to be very different. Then, the user is presented with a question titled “How can I help you today?”. This is a sort of portal page with a range of different answers, each bringing the user to different legal guided pathways. The answer texts range from factual situations (“There are bedbugs in my apartment”, “My tenant has stopped paying their rent”) to desired outcomes (“I would like to transfer my lease”). The user is also given an option to select “Other” if their case is not covered. Below, we will discuss the outcome of selecting this item.

Structuring a JusticeBot tool in this way makes it possible to develop multiple, independent legal guided pathways, in a modular fashion. It is thus a useful first step.

7.3.3.3 Identify a frequent type of case

We have identified a legal area, and built a scaffold schema that allows the creation of independent sub-paths for each specific legal issue. The next step is to identify a particular legal issue that is frequent enough to be useful to encode in the system. As discussed above, starting with frequent case types can increase the immediate usefulness of the created JusticeBot decision support tool.¹⁰⁹⁶

There are multiple ways of going about this step. If the tool is being built together with an institution, speaking to this institution to learn about the cases that are most frequent can be an important step. Annual reports may also provide this information. Finally, reading the case law from a court or tribunal and noting down the type of case can give an overview of which type of cases are frequently dealt with at that institution. If a corpus of decisions is available, performing data analysis on this corpus to discover topics could also be a useful way of identifying prevalent topics. Methods such as Latent Dirichlet Allocation can be employed to discover clusters of frequently co-occurring words.¹⁰⁹⁷ Salaün *et al* used a number of models to discover why tenants sue landlords.¹⁰⁹⁸

7.3.3.4 Building a schema

Once the specific type of case has been identified, the next step is to start building the schema covering the legal reasoning pathway around this legal issue in the JusticeCreator. This process was described above in 7.3.1 and 7.3.2.1.

There are multiple sources that can be relevant for this encoding. The law, of course, gives an important outline over the structure of the reasoning, and is likely to be at the base of the schema representation of legal reasoning. Likewise, books and legal doctrine can be an important source in understanding how the reasoning in a legal area functions. However, the goal with the JusticeBot is to encode the rules as they are practically

¹⁰⁹⁶ See 5.5.2.

¹⁰⁹⁷ David M Blei, Andrew Y Ng & Michael I Jordan, “Latent dirichlet allocation” (2003) 3:Jan Journal of machine Learning research 993–1022.

¹⁰⁹⁸ Olivier Salaün et al, “Why Do Tenants Sue Their Landlords? Answers from a Topic Model” (2022) Legal Knowledge and Information Systems 113–122.

applied by legal decision makers. Therefore, reading legal decisions to understand the important legal criteria is a crucial step in building the schema.

The schema also contains simplified explanations of the legal criteria. Here, the expert user will have to write simplified explanations of what the criteria mean. Once again, legal doctrine can be an important source for these explanations. Likewise, websites that aim to provide simplified legal information can be useful sources in elaborating these explanations. The explanations should include links to these sources to allow the user to explore the issues in-depth. If the tool is built together with an institution, they may allow the inclusion of information from their websites in the JusticeBot tool, which can greatly reduce the effort required to write the explanations.

7.3.3.5 Encode cases

The next step in building the JusticeBot decision support tool is the encoding of case law. This process is described above in 7.3.2. The cases serve to illustrate how legal criteria are applied, and to inform the user of the outcomes of cases similar to theirs.

Identifying cases to annotate is an important first step. The JusticeCreator includes functionality to automatically do this if we have access to a corpus of court cases.¹⁰⁹⁹ Otherwise, public case repositories can be used to find cases for annotation.

Once the case has been identified, the expert user follows the steps described above in 7.3.2.4 to add the case to the schema in the JusticeCreator. In doing so, they may realize that there is an issue with the schema. In this case, the expert user can go back to the previous step, and adjust the schema to better reflect the way courts reason around these types of cases.

Cases give an empirical verification of the schema created for a legal issue, thereby giving legitimacy to the system. They further provide important information to the user regarding how legal criteria are interpreted and the outcomes they can expect. Therefore,

¹⁰⁹⁹ See 7.3.2.4.1.

it is important to add a sufficient number of cases to the system. What exactly counts as sufficient will depend on the circumstances.

7.3.3.6 Verify the content

Once one or multiple legal issues have been encoded in the way above, the JusticeBot legal decision support tool should be verified, to make sure that the content is accurate. If the tool is built together with the institution, they may be able to read and verify the content. Otherwise, outside legal experts may be helpful in this task.

One way of allowing these individuals to review the schema is to send them the JSON-encoded version of the schema, which they can then import into the JusticeCreator to view and verify the content. The JusticeCreator also allows the export of an entire schema in the form of a Microsoft Word document, which can then be read and commented on by external collaborators.

7.3.3.7 Monitor feedback and adjust

Once the JusticeBot decision support tool has been launched, it is important to monitor the feedback of the users and adjust the system.

As we will see below, the JusticeBot frontend contains multiple ways for the user to provide feedback about individual pages or the entire process. Monitoring this feedback is important to make sure that the system serves the users in the best way possible.

Feedback may be related to the quality of individual pages, the overall process or legal issues that are not covered by the JusticeBot tool. This information can inform changes and adjustments to the JusticeBot schema and case law.

Of course, there may be other reasons to change the content in the JusticeBot. In some instances, legislation or court precedent may change, requiring updates to the schema. In other cases, the creator can continue adding more legal pathways to the system, in order to cover more possible situations. The system can be edited and updated by repeating the steps described above.

7.3.4 Conclusion

In this section, I discussed how legal information is encoded in the JusticeBot system. The methodology consists of encoding both legal rules and legal case law. The rules are encoded in a directed acyclic graph, containing the legal criteria and legal conclusions that judges typically apply to solve cases, as well as the logical links connecting the two. These elements can be intuitively and graphically encoded in a program called the JusticeCreator, allowing anyone to build legal decisions support tools.

The methodology also includes the addition of case law to the schema. Cases are read in light of the encoded schema and encoded in the JusticeCreator. These cases can be used to validate and expand the legal reasoning schema, to illustrate how judges apply legal criteria, but also to capture the outcome of cases that judges tend to order for certain types of cases.

Now that we have seen how legal data is encoded into the system, let us peruse the steps required for the JusticeBot to interact with the user and provide them with relevant and useful information.

7.4 Obtaining user input

We now have a schema containing legal rules and associated case law. The next step is to capture the situation of the user. Perhaps, the user is faced with a certain factual situation, and wishes to explore the legal implications of this. Otherwise, the user may want to achieve a certain goal, and wish to explore whether their circumstances fulfill the requirements of this goal. Therefore, the interface needs a way to capture the specific situation of the user, in a way that allows the comparison between the user case and previous cases.

In the JusticeBot schema, this is done by exposing the legal reasoning schema that we encoded above to the user. The user is thus asked to assess the same legal criteria that a judge would assess in solving their case. For each criterion, the user indicates whether they think that a judge would find that criterion to apply to their situation or not, aided by

plain-text descriptions and previous case law. By following the reasoning schema, the interface adapts to the answers given by the user to only ask relevant questions.

It might seem odd to ask the user to assess legal criteria. However, in my opinion, this is one of the few ways to capture the case of a user in a neutral state. By neutral, I here refer to a representation that can be translated between the yet hypothetical dispute a user faces, and the previously decided decisions. Only if these representations can be compared can the cases that are relevant and helpful to the user be identified.

The FactorBot research shows how a fact-based representation may not always be neutral. Capturing the facts from a previous case can be arduous and subjective. Further, the understanding of what constitutes a fact may be different between a judge and a layperson user.¹¹⁰⁰ Even if the comparison succeeds, as we have seen, it is not clear whether an AI system can accurately model whether a legal criterion applies or not based on a potentially unbounded set of facts.¹¹⁰¹

This is why the user is asked to apply the legal criteria in the JusticeBot. In essence, the user creates a *hypothesis* of the different legal criteria that a judge would find to apply or not apply. This hypothesis can then be compared to previous cases, to show the user the possible outcome of their case, should they have accurately predicted how the judge will assess the individual criteria. In order to support them in building this hypothesis, the user is provided with information and legal summaries.

7.4.1 The JusticeBot front-end

In order to record the details of their case, the user will interact with the JusticeBot front-end. The graphical aspects of this system were illustrated above in 7.2. Here, I will describe the technical and logical components of how this system works to capture the situation of the user.

¹¹⁰⁰ Compare 6.4.2.7.

¹¹⁰¹ See 4.4.3.4.3.2 and 6.4.2.5.

7.4.1.1 Technical background

Just like the JusticeCreator, the JusticeBot front-end is conceptualized as a single page application (SPA). This means that the entire system, including the logic and the content, is downloaded to the user of the computer before it is executed. This enables the system to respond very quickly, and to be very easy to host. The front-end itself is further independent from the content – it can load and run any schema file created using the JusticeCreator. It is thus ready to support the creation of many different legal decision support tools using the JusticeBot methodology.

The JusticeBot front-end is built using the Quasar framework.¹¹⁰² I built it in conjunction with the JusticeCreator in Spring 2021. It is now being further developed and improved at the Cyberjustice Laboratory. The JusticeBot front-end has been used to build the JusticeBot TAL, discussed below, and has thus been accessed over 17k times (see 8.4).

7.4.1.2 Logical functionality

The JusticeBot front-end helps the user logically traversing the schema encoding the legal reasoning in a certain area. Due to the way the schema is structured, this reasoning process is very simple.

The interface starts with the first block in the schema, as defined by the user. The system then traverses the blocks in the schema one by one. Depending on the type of the block, it performs a different action:

- **If the block is a question block:** Show the question to the user, together with the case summaries illustrating how judges have previously applied the criterion. Store the answer the user selects and move on to the block the answer is linked to.
- **If the block is an information block:** Store the information block in a stack for later. The information block will be part of the information presented to the user at the end. Immediately move on to the next block, defined by the "next" connector of the information block.

¹¹⁰² note 1080.

- **If there is no more block:** The pathway is finished. Move on to the analysis step (see 7.5) and show the results to the user (see 7.6).

This algorithm is very simple and can thus be executed on any device. Further, it is very easy to reason about, which helps in the creation of the system.

7.4.1.3 Example

Let us explore how the user case is captured based on the schema we used above in 7.3.2.4.2, replicated below in Figure 41. Note that the system is kept barebones for the sake of the example. Normally, it would contain explanations of the legal criteria, and more cases to illustrate how they are applied.

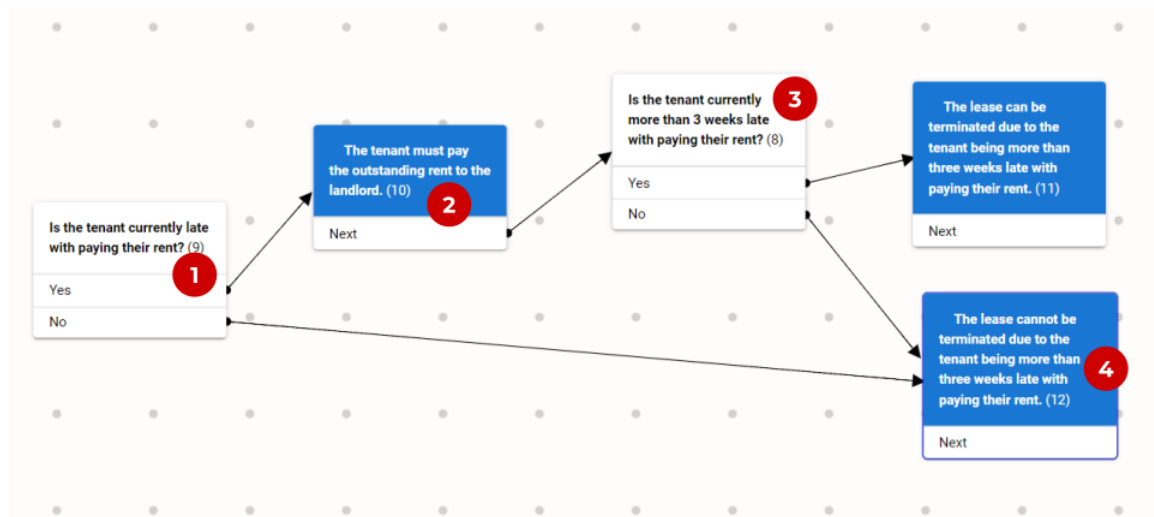


Figure 41 - Reasoning schema regarding lateness of rent

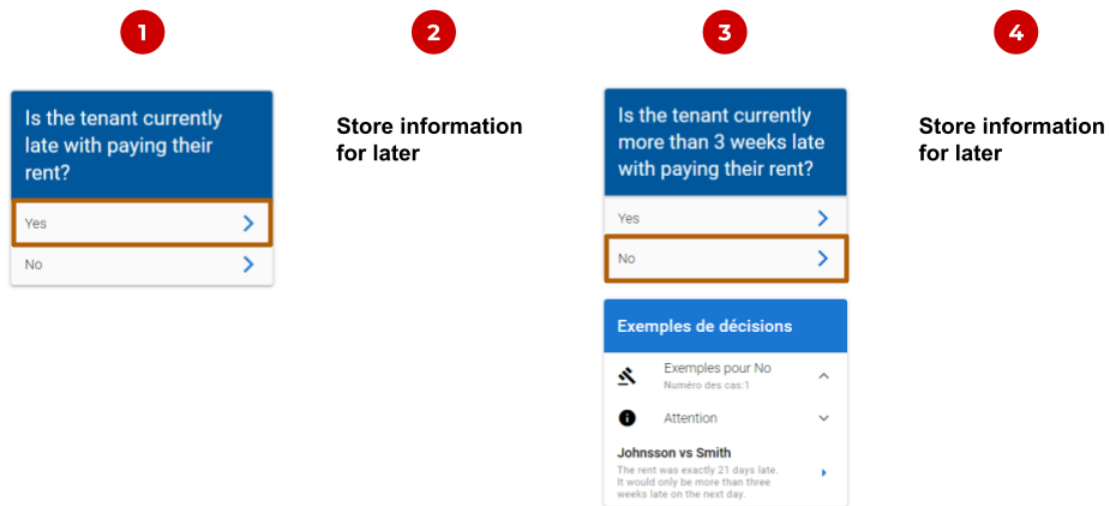


Figure 42 - A possible user path through the schema in Figure 41

Figure 42 displays how the system could interact with a user.

- The system starts with the question block at (1). This is displayed to the user, who chooses "Yes" as the answer. The system follows the connection (see Figure 41), arriving at the information block in (2).
- Information block (2) is stored, and the system advances to the question block at (3).
- The user is shown question block (3). The user is also shown the case that we previously encoded in 7.3.2.4. This can help them determine whether the criterion applies in their case. In this case, they picked the answer "No", which leads them to the final information block at (4).
- The information block at (4) is stored, and the system tries to advance.
- Since there is no more block, the system moves on to the analysis stage. The relevant details of the case of the user have been captured.

Based on the user traversing the schema, we have captured the details of the situation of the user in the system. They have created a hypothesis over which criteria a judge would

find to apply in their case. Figure 43 shows a visual representation of the path the user followed.

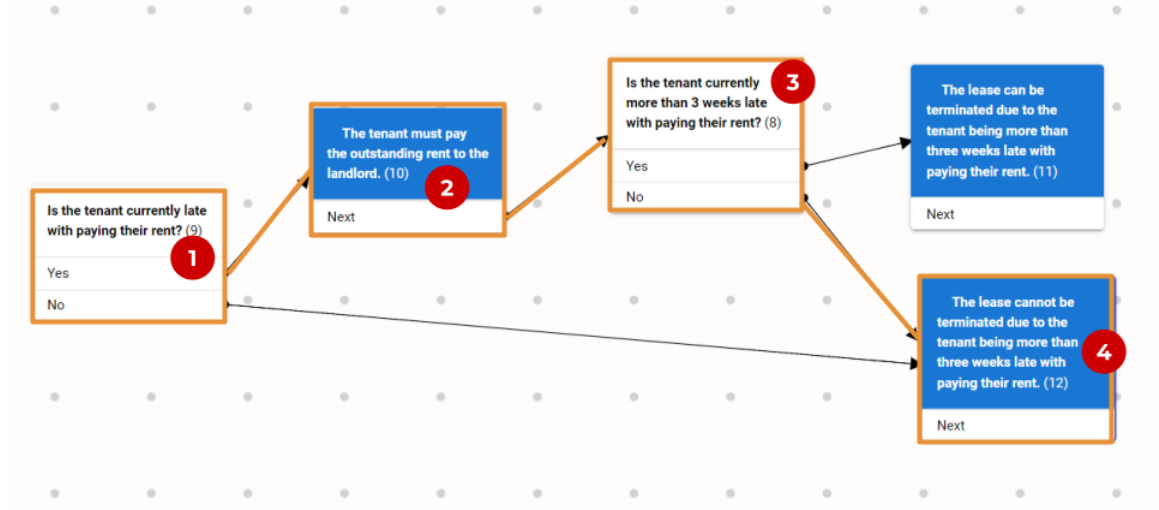


Figure 43 - Hypothetical path through legal reasoning steps

This path can also be represented in terms of the answers to the individual question blocks:

Criterion	Does it apply?
Is the tenant currently late with paying their rent?	Yes
Is the tenant currently more than three weeks late with paying their rent?	Yes

Or, it can be represented by the stored information blocks that we passed in traversing the graph:

Passed information block
The tenant must pay the outstanding rent to the landlord.
The lease cannot be terminated due to the tenant being more than three weeks late with paying

their rent.

The attentive reader will notice that the way we captured the case of the user corresponds exactly to the way we encoded previous cases into the system in 7.3.2. This is a very desirable property, as we can now compare the encoded user case to the previous cases we have stored in the system. In the next section, we will explore how this comparison is made.

7.5 Analysis of the case of the user

We now have a representation of the hypothetical case of the user, stored in the same format as previous cases encoded into the system. The next step is to analyze the user case, in light of the encoded data, to surface relevant information and case law. In this section, I explain how this analysis is made.

7.5.1 Surface relevant information

The JusticeBot schema contains information that can be shown to the user, in the form of information blocks, representing legal conclusions. As we saw above, the information blocks that are traversed by the user in answering the questions are stored in the system. These are dependent on the answers selected by the user and can thus be displayed to give information about their legal rights.

7.5.2 Surface relevant case law

The answers of the user can also be used to surface the outcome of previous cases. In order for this feature to work, the hypothetical case of the user is compared to the previous cases encoded in the system. We therefore need a measure of similarity – which cases are similar enough to that of the user to warrant the display to the user?

Through the research in the JusticeBot project, three such ways have been explored. I will briefly elaborate on these ways, and describe why I selected the one I did.

7.5.2.1 Similarity assessment 1 – matching all questions

The first way cases can be compared in is by considering the individual answers to every single question. Cases are thus considered to match if every single question block is answered in the same way between the hypothetical user case and the previous case.

This way of comparing cases works well, but I found it to be unnecessarily restrictive. For example, let us consider schemas with cumulative requirements (see 7.3.1.4.4). Here, multiple legal criteria need to be fulfilled in order to achieve a certain outcome. If one of the criteria is not fulfilled, the conclusion cannot be achieved. For the outcome, it does not matter which of the legal criteria is not fulfilled. Therefore, matching by all questions likely excludes many cases that are relevant to each other.

7.5.2.2 Similarity assessment 2 – matching all legal conclusions

Another possible way of comparing a user case to previous cases is to match cases if all of the same legal conclusions were passed in the user path versus the saved path of the previous case.

This is another interesting way of comparing cases. However, it also excludes certain cases that may be relevant. This is especially in the case of intermediary conclusions. For example, a case where a tenant is late with paying the rent is relevant to all other cases where the tenant is late with paying the rent, irrespective of whether the tenant is also three weeks late with paying the rent or not.

7.5.2.3 Similarity assessment 3 – matching individual legal conclusions

Therefore, the JusticeBot works by comparing cases by the individual legal conclusions that are passed when traversing the legal reasoning pathway. As described in 7.3.2.3.3, cases in the JusticeBot are encoded in terms of the local outcomes, that are directly tied to legal conclusions. Thus, the outcomes of cases are linked to the individual decisions a judge takes and can be shown to the user whenever they pass the same legal conclusion block as previous cases have passed in the schema.

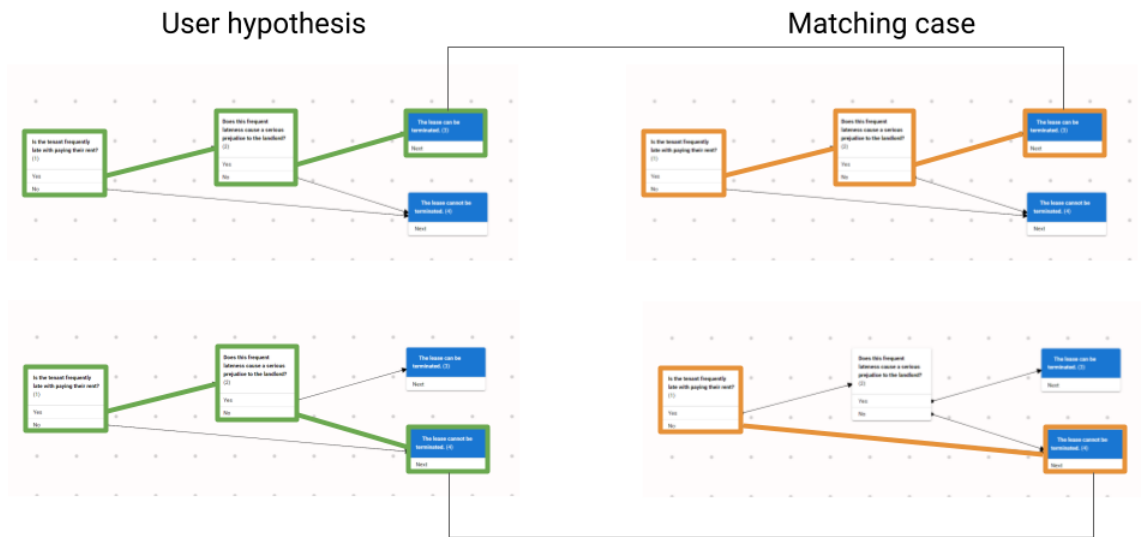


Figure 44 - A user case can be matched to previous cases, even if not all answers are answered in the same way

Figure 44 shows how user hypotheses can be matched to previous cases, based on the information blocks the user passes through. In the first example, the same criteria were found to apply in the user case and the previous case. Since both the previous case and the user case pass through the same information block, the previous case is found to be relevant to the user case.

In the second example, the user answers the questions slightly differently from the way the judge did in a previous case. However, since both the user case and the previous case end up passing through the same information block, the previous case is still found to be relevant. Using this type of matching, the user can be shown a larger number of relevant cases, compared to the approaches discussed above.

7.5.3 Example

Above, in 7.4.1.3, I showed a possible way for a user to traverse the example JusticeBot schema. The user claimed that the tenant was late in their situation, but not more than three weeks late. In doing so, they traversed the following information blocks in the schema, representing legal conclusions:

Passed information block
The tenant must pay the outstanding rent to the landlord.
The lease cannot be terminated due to the tenant being more than three weeks late with paying their rent.

Previously, in 7.3.2.4, we had added outcome from a case to these legal conclusions. Since the user passed the same legal conclusions, we can assume that the outcomes are relevant for them, and that they would benefit from knowing the outcomes that were previously awarded in these cases. The information combined with the case law that should be provided to the user is thus as follows:

Passed information block	Case law example outcomes
The tenant must pay the outstanding rent to the landlord.	<i>Johnsson vs Smith</i> - The judge ordered the tenant to pay the rent, and added an order for the tenant to pay the rent on the first day of the month in the future.
The lease cannot be terminated due to the tenant being more than three weeks late with paying their rent.	<i>Johnsson vs Smith</i> – The lease was not terminated.

The information thus contains legal information, along with examples from real-world cases where the judge reached the same legal conclusions. In this case, only one case is encoded. For real-world systems, we would want multiple cases, so that the user can get a representative overview over how judges decide these types of cases.

Now that we have collected the relevant information, let us explore how this information is presented to the user.

7.6 Providing results to the user

We have selected information from the schema, and cases that are relevant to the user. The only remaining step is determining how this information should be shown to the user. This step is relatively simple in the JusticeBot methodology. In the step above, we collected the relevant information and case law. Next, the JusticeBot front-end shows this information to the user.

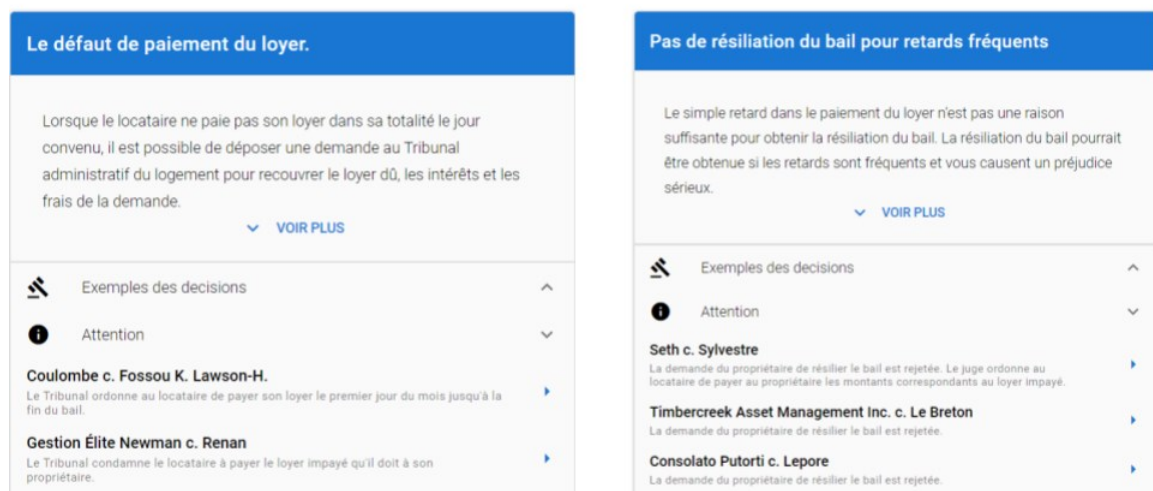


Figure 45 - Information blocks presented at the end of a pathway

Figure 45 shows how the selected information blocks appear in the JusticeBot front-end. Each information block that was passed by the user is showed with the title, legal information, and outcomes of previous cases. These cases lend legitimacy to the system, by showing that the results are empirically backed, and could be very helpful for the user in understanding possible outcomes.

The information provided further serves as an explanation for the possible outcome suggested by the system. Each information block can be given an explanation, that provides details of the possible outcomes of the cases. This explanation should contain information about why this outcome could be awarded (i.e. the tenant is late with paying

their rent, therefore they may be ordered to repay that rent) or references to the relevant legislation. Writing these explanations is simple for the legal expert, since they are able to examine the schema to see exactly which answers an individual chooses to arrive at a certain information block.

Finally, the user has the option to examine the answers they provided, in order to go back to any stage and explore what happens should the judge settle on a different answer.

We have now seen the entire functioning of the JusticeBot methodology, from the encoding of legal information to the capturing of the user information, to the provision of legal information and relevant cases based on the user situation. Next, let us explore some mechanisms in the JusticeBot methodology that can be used to capture feedback from the user, and how this can be used to improve and expand the tools.

7.7 Feedback mechanisms

The JusticeBot methodology contains a number of feedback mechanisms designed to capture the behavior of the users with regards to the system. These mechanisms are very important, as they allow the understanding of how users interact with the system, which allow the ongoing evaluation of the system, as well as inform decisions such as which areas to change or improve, and how to expand the system in a meaningful way. Let us briefly examine these mechanisms.

7.7.1 Web analytics

The JusticeBot platform can be connected to a web analytics provider, such as google analytics. This allows the examination of real-world usage patterns. For an analysis of this kind of data for the JusticeBot focused on landlord-tenant disputes, please see below at 8.4.1. Here, I will introduce the general kinds of data such a system allows us to collect.

Information that can be gathered includes the number of users that access the system, the general location that they access the JusticeBot system from, which device they use and how much time they spend on the platform. This can give an insight into the behavior of the users, and allow the adjustment of the platform, for example by focusing attention on a certain user type or platform used to access the system.

The analytics platform can further provide information about how many of the users were able to complete the pathway and receive an analysis, versus how many fell outside of the scope of the system. This can be an important metric to monitor, in order to maximize the impact of the system over time and understand how useful the platform is.

Further, the analytics platform can give in-depth information about the interaction of the users with the different elements of the platform. For example, the analytics allow us to see which pathways the users access, and which specific answers are chosen at any point in the pathway. This can guide the expansion of the platform. For example, in the JusticeBot TAL, the majority of the users turned out to be tenants – in order to maximize the impact of the tool, focusing on tenant-pathways could therefore be desirable.

Finally, individual interactions with different on-screen elements can be tracked. For example, the system can track how many people access the case summaries regarding a certain question, and how many followed the links to read the individual cases.

These types of analytics can be a very important tool in understanding overall user behavior. The user does not expend any effort for their feedback to be captured. However, web analytics may not give high-fidelity insights into the subjective experience of the user. To capture these, more intentional feedback mechanisms have also been included into the system.

7.7.2 Star rating on each page




Figure 46 - Star rating of individual pages

Figure 46 shows a design element that is present on every single page in JusticeBot based utilities. It allows the user to rapidly evaluate the page they are currently on, by assigning a value of between one and five stars. If the rating is at 3 stars or below, they are asked to optionally provide written feedback on the reason for the rating.

The rating is stored together with a reference to the individual page the user was on when rating the page. This allows the in-depth understanding of issues that are specific to certain pages. For example, if a certain question is not clear, this may be reflected in the ratings of that page, allowing the adjustment of this particular page. While the star rating page provides an insight into the subjective experience of the user, it also requires effort on the part of the user, which may make the data scarcer.

7.7.3 Survey at the end

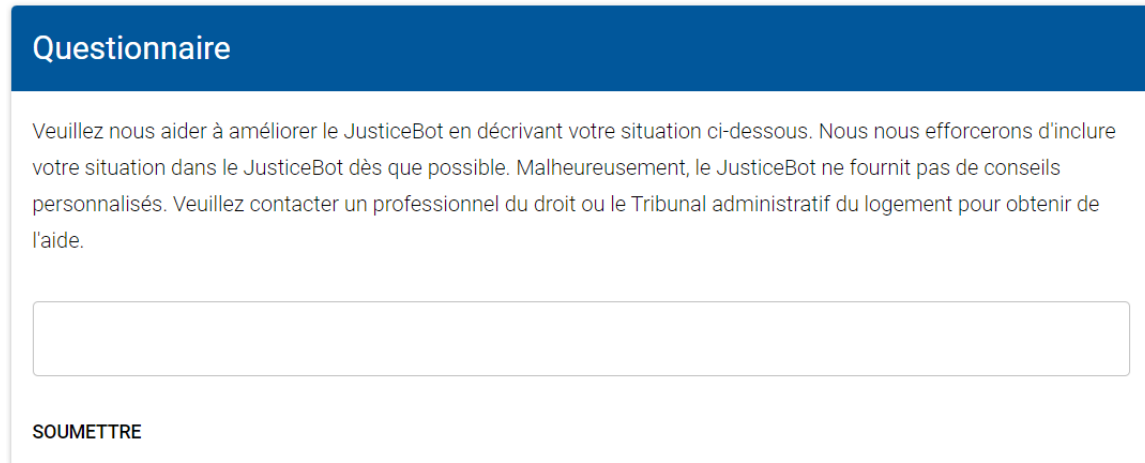
Another mechanism for collecting feedback from the user is a survey presented at the end of a JusticeBot pathway. Which specific questions are included of course depends on the nature of the specific utility built using the JusticeBot methodology.

In the JusticeBot TAL, the survey asks for detailed feedback regarding the different steps of the JusticeBot, including the overall experience and whether the user would recommend the system to a friend. While the number of people that fill out the survey may not be significant, it can give an important insight into the overall subjective experience of the user.

7.7.4 Missing issue screen

A powerful tool for user feedback in the JusticeBot methodology is the “missing issue” screen. Due to the wealth of particular legal issues that may arise in a legal area, it is unlikely that every conceivable issue is covered. Therefore, the JusticeBot methodology focuses on frequent legal issues, that can help a large amount of people. The process for selecting such areas is described in 7.3.3.3.

If the case of a user is not covered by the JusticeBot based tool, they will reach a screen that informs them of this. This may occur early at the stage of the portal page, that allows them to enter the different legal pathways, where the user is given the choice “Other” if none of the suggested pathways apply to the user. It may also be at a later stage, for example if a specific question determines that a particularity of the user case takes it out of the scope of the JusticeBot.



The image shows a web form titled "Questionnaire" with a blue header. Below the header, there is a paragraph of text in French: "Veuillez nous aider à améliorer le JusticeBot en décrivant votre situation ci-dessous. Nous nous efforcerons d'inclure votre situation dans le JusticeBot dès que possible. Malheureusement, le JusticeBot ne fournit pas de conseils personnalisés. Veuillez contacter un professionnel du droit ou le Tribunal administratif du logement pour obtenir de l'aide." Below this text is a large, empty rectangular text input field. At the bottom left of the form, there is a button labeled "SOUMETTRE".

Figure 47 - Form to ask user to describe their uncovered issue

At this stage, the user is taken to the “missing issue” screen. They are informed that the JusticeBot does not yet cover their issue and asked to fill out a text field (shown in Figure 47) describing their issue. Their description is then saved to a database.

This mechanism is very useful. It allows the creator of the JusticeBot tool to get an overview over the issues that users face that are not yet covered by the JusticeBot. If certain issues appear frequently, this is an important indication that the issue is an important real-world concern, and that covering it in a legal guided pathway would be a viable way to increase the usefulness of the system.

The data is likely to be more representative than other methods of determining which legal areas to cover, since it stems from real-world users accessing the system. It thus corresponds directly to the target user of the system, which may not be the same as the type of cases that go to court.

Beyond showing areas that should be covered by the JusticeBot, the data collected on this screen can also inform the creator of the system of potential discoverability issues. If a user describes a situation that is, in fact, covered by the JusticeBot based tool, this means that they did not manage to find the relevant path. Based on this insight, the creator may wish to reformulate the entry point of the pathway.

7.7.5 Conclusion

Collecting real-world feedback from a deployed JusticeBot instance is a crucial step in monitoring how users interact with a system, and how to make it more useful. The JusticeBot offers multiple such mechanisms, allowing the collection of different types of user statistics. Some of these, such as web analytics, are very low effort for the user, but may not give in-depth information about their subjective experience. Others, such as ratings and surveys, require more effort on behalf of the user, which means that they may not be employed as often. However, they can give qualitative insights into user experience and behavior. Combining these tools can thus provide an overview of how users interact with the system, and how it can be improved.

7.8 Discussion

Now that I have described the JusticeBot methodology, let us discuss some aspects of the methodology. I will first describe how it compares to previous work, and then discuss how well it corresponds to the design criteria set out in Chapter 5.

7.8.1 Comparison to prior work

7.8.1.1 Encoding of legal rules

The JusticeBot methodology uses directed acyclic graphs to encode the legal reasoning steps in a legal area. This consists of a flowchart like representation that encodes legal criteria and legal conclusions in the same graph.

This way of representing legal rules has a lot in common with the systems described in 4.5.3.2, using rule-based reasoning. These systems used logical representations of legal rules to be able to give legal advice or information. Many of these systems rely on the encoding of logical legal connections in a programming language such as prolog, which can then forward-chain from facts to outcomes, or backward chain from outcomes to the required facts.

I believe the system used in JusticeBot, namely relying on directed acyclic graphs to encode legal criteria and conclusions, has several interesting and potentially novel aspects.

7.8.1.1.1 Intermediary conclusions

First of all, the format used to encode the schema allows for several intermediary legal conclusions. This lends itself to allowing the user to explore many possible consequences of a factual situation. Being able to produce multiple independent results, based on a single legal path, can be very helpful in these situations.

7.8.1.1.2 Non-requirement of a technical background for building tools

A second important feature of the schema representation used in the JusticeBot is the simplicity of the representation. The system is designed to not require any technical background to use and create, allowing legal experts to create JusticeBot tools in their domains of expertise, without relying on programmers. There are a few properties of the schema representation used that make the representation suitable for use by non-programmers.

The way schemas are represented in the JusticeCreator is inherently visual. Schemas are created in the JusticeCreator, which allows the WYSIWYG (What You See Is What You Get) creation, editing and connection of the different legal criteria and legal conclusions.¹¹⁰³ This way of editing does not require the knowledge of any logical concepts or programming languages, as legal experts can design and understand the system by merely watching the schema itself. Hopefully, this way of editing and creating schemas can decrease the effort of encoding legal rules.¹¹⁰⁴

The visual nature of the schema, combined with the simple way that the system traverses the schema, makes it very easy to reason about what the user will see. Akin to a flowchart, the system follows the arrows in the schema to arrive at the next block, which is then shown to the user in case of a question block or saved for display at the end in case of an information block. This system can be learnt in a few minutes by anyone. It is also fully deterministic – the expert user is able to precisely verify the content the user

¹¹⁰³ See 7.3.1.5.

¹¹⁰⁴ See 4.5.3.3.1.

will be exposed to. Despite this, as we have seen, the system is flexible enough to represent complex legal reasoning paths.¹¹⁰⁵

In previous work, visualizations of the logical patterns were often shown.¹¹⁰⁶ Usually, these visualizations are generated post-hoc in order to visually explain the reasoning of a system, while the content itself is encoded in the form of a logical language such as prolog, or other text-based tools. These systems may require technically trained individuals to encode the rules, unlike the JusticeBot schema, where the visualization itself is the logical content.

Further, the systems often use more complex reasoning tools, such as forward-reasoning or backward-reasoning, to traverse the encoded rules. While these may be understandable for a technically trained audience, they may be difficult to understand for lawyers, making it difficult to encode and verify the content for legal experts. The JusticeCreator, on the other hand, uses a simple but effective system of implementing legal reasoning. Being equivalent to a flowchart, it allows easy understanding of which questions and which information the user will see at which point.

7.8.1.1.3 Creation and verification using case law

Another interesting aspect of the JusticeBot is the use of case law to capture the legal rules. This could overcome the issue regarding syntactic ambiguity of legal rules, which has often been discussed in prior work.¹¹⁰⁷

In the JusticeBot methodology, the rules governing a legal area are discovered by reading legislation and case law in conjunction. The aim is to discover the criteria that judges typically apply in the real world, in order to decide upon the outcome of certain cases. Thus, the encoded version of the rules is based on the interpretation of a statute made by

¹¹⁰⁵ See 7.3.1.4.

¹¹⁰⁶ See e.g. Allen & Engholm, *supra* note 727 at 390–391; Walker, *supra* note 579 at 240; Thomasset, Blanchard & Paquin, “Loge-expert”, *supra* note 749 at 384.

¹¹⁰⁷ See e.g. 4.5.3.3.4

a court, rather than the interpretation of the legal expert that builds the system. The rules are discovered from reading case law.

Especially in areas of high-volume, low-intensity cases, each rule may have been applied many times, allowing a consensus to emerge on how it should be applied, from a syntactic standpoint. The JusticeBot methodology should be well suited to infer this consensus by reading and encoding a number of cases.

Further, as new cases are encoded in the JusticeCreator, the schema is continually verified. Encoding case law requires the expert annotator to step through cases and associate them to steps in the schema. In doing so, they are able to verify that the schema accurately captures the reasoning steps typically undertaken by a judge. Thus, as the system is updated, the schema is empirically verified.¹¹⁰⁸

While the use of case law as a source to capture the legal reasoning steps performed in a legal area is not new,¹¹⁰⁹ I believe the formalized methodology used in the JusticeBot to empirically verify the schema to be quite novel.

7.8.1.2 Case law to illustrate legal criteria

The JusticeBot uses summaries of case law to illustrate how legal criteria are applied. This can help the user understand whether a certain legal criteria applies in their case or not.

This approach is quite different from the use of cases to predict the applicability of legal criteria employed in the FactorBot approach.¹¹¹⁰ As we have seen, this approach had difficulties with the high number of cases needed for encoding,¹¹¹¹ and the difficulty of

¹¹⁰⁸ See 7.3.2.1.

¹¹⁰⁹ See e.g. Thomasset, Blanchard & Paquin, “Loge-expert”, *supra* note 749 at 382; Walker, *supra* note 579 at 197.

¹¹¹⁰ See 6.4.1.2.1.

¹¹¹¹ See 4.4.3.4.3.3 and 6.4.2.2.

predicting situations with new or complicated factual situations.¹¹¹² Further, the approach may be seen as giving legal advice.¹¹¹³

The approach taken in JusticeBot is very different. Here, the user is given a few illustrative examples of how a criterion is applied, and then asked whether they believe that a judge would see a certain criterion to apply or not. This overcomes the issue of needing many cases, as cases are not used for prediction but rather as examples. Further, the use of plain language instead of specified factors to describe cases overcomes the difficulty of fitting cases into neatly defined factors, as the annotator can use language to encode the cases.

This idea in itself is not new – several projects suggested the use of case summaries to overcome open-textured terms in rule-based reasoning systems.¹¹¹⁴ Further, there have been several systems, such as CATO and IBP, that arrange factors from cases by the legal issues that they discuss.¹¹¹⁵

However, I believe the use of cases in the JusticeBot methodology to be very developed and specified, by providing an end-to-end, implemented methodology. The JusticeBot methodology includes an automated system to identify cases that could be relevant to annotate, based on semantic similarity to the legal criteria.¹¹¹⁶ The cases used in the JusticeBot approach are commonplace cases, which may correspond the most closely to the situation of a layperson user.¹¹¹⁷

The JusticeBot methodology can further be compared to GREBE, which uses semantic nets to capture and compare cases.¹¹¹⁸ In a way, the cases in the JusticeBot are also encoded in a network. However, what specifically is captured differs between the methodologies. While GREBE captures individual facts and their relationships as nodes

¹¹¹² See 4.4.3.4.3.2 and 6.4.2.5.

¹¹¹³ See 6.4.2.1.

¹¹¹⁴ See 4.5.3.3.3.

¹¹¹⁵ See 4.4.3.4.2.2.

¹¹¹⁶ See 7.3.2.4.1.

¹¹¹⁷ See 7.3.2.2.3.

¹¹¹⁸ See 4.4.3.4.2.3.

and edges in the semantic network, the JusticeBot methodology uses legal issues as nodes, and then attaches the cases to these nodes by summarizing the facts as they relate to a certain legal issue. Likewise, the purpose of the systems is different – while GREBE aims to generate explanations, JusticeBot targets laypeople and aims to inform them of their rights right after a situation has occurred.

Once the cases have been identified, the annotator steps through each legal criterion in a pathway. The annotator is given guidelines to determine whether the case should be summarized or not.¹¹¹⁹ If so, they will add a summary written in a certain way to the criterion,¹¹²⁰ and otherwise move on to the next criterion.

Following this methodology would add relevant, real-world summaries to legal criteria, which can help the user determine whether a specific criterion would apply in their case.

7.8.1.3 Case law to illustrate outcomes

One of the most interesting features of the JusticeBot is the comparison of user cases to give information about the outcomes of previous cases.

In prior work, projects that use case law to predict outcomes of future cases include Dahan *et al*, who predicted the length of notice period, and Stranieri and Zeleznikow who predicted the division of assets in divorces.¹¹²¹ The FactorBot aimed to predict the outcome of landlord-tenant dispute cases based on facts that appeared in a case.¹¹²² In these projects, the researchers built machine learning models to predict the outcomes of the cases based upon the facts of a case. Cases are seen as *similar* (i.e. the outcome of a previous case is relevant to a future case) if the facts of the case match. We ran into issues applying this approach in unbounded domains, where the facts of a case may be very varied, and thus difficult to capture in terms of categorical representations. Further,

¹¹¹⁹ See 7.3.2.2.3.

¹¹²⁰ See 7.3.2.2.4.

¹¹²¹ See 4.6.3.2.

¹¹²² See 6.3.3.2.

as we saw in the Factorbot research, it was difficult to annotate the required number of cases, and the predictions were not very accurate.¹¹²³

The JusticeBot system also uses case law to illustrate potential outcomes. However, it uses a different approach. Instead of building a model of potential outcomes based on the facts of a case, it uses the legal reasoning schema to index cases, that can then be shown to the user.

This implies a completely different notion of similarity than most of the previous systems. While the previous systems saw cases as similar if the facts were overlapping, the JusticeBot methodology sees cases as similar *if the judge came to the same legal conclusions in previous cases as the user expects them to do in their case*.

Of course, this approach requires the user to be more active than by using facts. The user needs to predict how a judge would apply a legal criterion in their case, to enter their situation into the system. However, the methodology provides the user with example cases that support them in making this assessment.

Additionally, the approach is very transparent. As we have seen, predicting the case of a user based on only facts does not always work well, as every situation is unique, and new factors may impact the situation.¹¹²⁴ Further, comparing the facts as a layperson sees them to the facts as a judge sees them is not trivial, as users may see their situation on a different level of abstraction than the judge. In trying to capture the case of a user in a neutral way, it is very difficult to remain purely on the side of facts, as even many base-level facts require some level of judgment to establish.¹¹²⁵

The JusticeBot approach acknowledges these difficulties. It does not try to accurately predict what a judge will say based on purportedly neutral representations of a case. Rather, it gives the user the tools to support them in building a hypothesis about how

¹¹²³ See 6.4.2.

¹¹²⁴ See 6.4.2.5.

¹¹²⁵ See 6.4.2.7.

their case is likely to be seen by a judge, and allows them to explore the outcomes that judges *previously* awarded, in cases where they came to legal conclusions matching this hypothesis. In a way, the JusticeBot system can be seen as a sophisticated case retrieval system, that indexes cases by the legal reasoning path of a judge and allows the user to search the database by supplying a new potential reasoning path. This is in line with the goals for the JusticeBot methodology of being an augmented intelligence system and giving legal information rather than legal advice. Further, as far as I have seen, it is quite unique in the scope of previous work in the area.

One may further question the need for case law at all. The JusticeBot contains a rule-based system, that is able to provide information on the legal conclusions and outcomes that the rules specify for certain legal conclusions. Why then, do we need encoded case law to give these outcomes?

Section 7.3.2.3.1 explains why cases are necessary, in addition to information provided through the rule-based system. Cases illustrate the outcomes that judges *factually* order in certain cases, rather than what they in theory should decide on.¹¹²⁶ They can also illustrate the quantitative decisions that judges make based on certain legal conclusions, such as the amount of damages they may award. Finally, the cases give legitimacy to the information provided by the schema. Providing case law examples of outcomes thus significantly enhances the usefulness of the JusticeBot system, in a way that is hopefully a useful contribution to previous work with expert systems.

7.8.1.4 Fully implemented toolchain

Another unique aspect of the JusticeBot methodology is that it goes beyond a research project and has been fully implemented. This includes both the tool to create JusticeBot decision support tools, and the tools to deploy JusticeBot based tools directly to the public. While a lot of research has mentioned the need for such implementations, it is relatively rare to see systems that are implemented to this extent.

¹¹²⁶ Compare 1.3.2.1.

7.8.1.4.1 The JusticeCreator

The JusticeBot methodology has been implemented in the JusticeCreator tool, which allows the quick and intuitive creation of JusticeBot based decision support tools. The tool can be used by anyone to build such tools, without the need for technical assistance. It comprises methods to retrieve cases from a database that may be suitable for annotation, and a fully fleshed out methodology for encoding legal rules and case law.

In previous research, there have been a number of projects discussing this kind of interface. Thompson, for example, argues for the need to implement an interface for the simplified entry of information into a system by legal experts.¹¹²⁷ Susskind argues for the creation of carefully and clearly formulated standard practices, to help the creation of expert systems.¹¹²⁸ Hopefully, this research can contribute the start to such standard practices.¹¹²⁹

There have also been a number of projects that do implement concrete methodologies and tools. For example, Al-Abdulkarim *et al* built a support environment for their ANGELIC methodology.¹¹³⁰ This is a web-based tool that allows the visualization of rules, input of new cases and other interactions with a database supporting the creation of tools in the ANGELIC methodology.¹¹³¹ This approach uses prolog rules for encoding tests, which may be aimed at a more technical audience.¹¹³²

There are also publicly available, production-ready tools that aim to support the development of legal support tools. One such tool is DocAssemble, which is an open-source system leveraging the programming language Python to create sophisticated systems able to assemble legal documents.¹¹³³ Another such tool is A2J Author.¹¹³⁴ Both

¹¹²⁷ Thompson, *supra* note 75 at 42.

¹¹²⁸ Susskind, “Expert systems in law”, *supra* note 598 at 4.

¹¹²⁹ See 7.3.3.

¹¹³⁰ Al-Abdulkarim *et al*, “Factors, issues and values”, *supra* note 702.

¹¹³¹ See 4.4.3.4.2.7.

¹¹³² Al-Abdulkarim *et al*, “Factors, issues and values”, *supra* note 702 at 5.

¹¹³³ note 456.

¹¹³⁴ note 455.

of these tools are very valuable resources and have been used to create important tools aiming to support access to justice¹¹³⁵ and are in many ways much more sophisticated than the JusticeCreator. They encompass powerful and flexible logical reasoning frameworks and/or computer code. Furthermore, they have been used to create many practical tools accessible to the public.

These tools provide features for helping individuals fill out court forms. The JusticeBot methodology, on the other hand, focuses on providing individuals with legal information regarding their specific issue. Further, the JusticeBot tools uses the methodology described in this thesis to capture rules and cases, by encoding the criteria that judges tend to assess when dealing with certain types of cases, and then stepping through these together with the user at runtime. The previously mentioned tools likely use other ways to encode legal information, or leave the user the choice on how a legal area should be captured. Finally, the small and focused feature set of the JusticeCreator may allow legal experts to quickly learn and use the JusticeCreator in a productive manner, together with the JusticeBot methodology. This is a powerful combination that has already led to the successful deployment of a legal decision support tool,¹¹³⁶ and will hopefully be used to create many more tools in the future.

7.8.1.4.2 Implemented front-end system

The JusticeBot methodology also encompasses a front-end system, which allows the user to access the legal decision support tool from any web-capable device such as a smartphone or a laptop. This tool has been deployed to the public in the scope of the JusticeBot TAL, see Chapter 8.

The use of web-based tools is, of course, nothing new. While a lot of legal reasoning research does not contain a full implementation of a front-end of the system, there are some projects that do. For example, Zeleznikow discussed such a system in 2002.¹¹³⁷

¹¹³⁵ See e.g. Salter, “Online dispute resolution and justice system integration”, *supra* note 488.

¹¹³⁶ See Chapter 8

¹¹³⁷ Zeleznikow, *supra* note 67.

Likewise, many systems developed by legal aid groups, based on projects such as the A2J author, expose their information through a web-based interface, such as the Civil Resolution Tribunal in British Columbia and Community Legal Education Ontario.¹¹³⁸

There are a few differences between the JusticeBot front-end and these systems. First of all, the JusticeBot frontend is based upon the JusticeBot methodology, described in this chapter. It thus exposes both legal rules and case law in order to provide useful information to increase access to justice. Evidently, the other systems are based on other methodologies.

The JusticeBot front-end further benefits from a high degree of generality. The JusticeBot front-end derives almost all of its content from a schema file, exported by the JusticeCreator. The work is thus not limited to a single decision support tool but can rather be used to directly publish any legal decision support tool created with the JusticeCreator. Improvements made to the JusticeBot frontend benefit all of these tools, several of which are under development at the Cyberjustice Laboratory. For a discussion of more potential tools, see section 9.3.

Finally, the technical implementation of the JusticeBot frontend allows it to very easily scale to an enormous number of users. Since it is conceptualized as a static page, that executes on the device of the user, it is trivial to host, and near impossible to hack. Many free or very cheap services exist to host static pages. Even the “free” tier some of these hosting services could support millions of monthly users of the platforms.

7.8.2 Discussion

Now that we have compared the JusticeBot methodology to prior work, I will discuss how well the methodology corresponds to the specific goals set out in Chapter 5, namely:

- The system should target laypeople

¹¹³⁸ Salter, “Online dispute resolution and justice system integration”, *supra* note 488; “Home”, online: CLEO (*Community Legal Education Ontario / Éducation juridique communautaire Ontario*) <<https://www.cleo.on.ca/en>>.

- The system should target high-volume, low-intensity legal areas
- The system should give specific and useful legal information
- The system should be practical.

I will go through these criteria one by one, and discuss how the specifics of the JusticeBot methodology correspond to these criteria.

7.8.2.1 The system should target laypeople

The first criteria, described in 5.2, targets the fact that the system should be useable by laypeople. Since laypeople cannot be assumed to have any legal knowledge, such a system should start with a goal or factual situation (5.2.3), the system cannot rely on knowledge of legal concepts (5.2.4) or the user's ability to evaluate the quality of the provided information (5.2.5), it should avoid complex language (5.2.6) and it should respect the technological modalities of laypeople (5.2.7).

I believe that the JusticeBot methodology corresponds to these criteria. The system allows the user to select an option that corresponds to their situation on the portal page, which can be referred to either by a certain factual situation that has occurred (such as “there are bedbugs in my apartment), or a goal that they wish to achieve (such as “I wish to terminate my lease”). Thus, the user is not required to understand the legal situation underlying a situation or goal, but is instead guided towards the correct legal rule by the system.

The methodology further does not rely on an understanding of legal concepts – with a caveat. The system does not assume that the user understands legal concepts (such as “habitable condition”) before using the system. However, the system does require the user to assess whether the legal concepts apply to their situation or not. To their aid, the user is given summaries of relevant commonplace legal cases, that can inform them understand how the legal criterion was applied in previous cases. While this does assume a certain level of literacy on behalf of the user, it is also a very transparent approach. It explicitly assumes that AI-based prediction of legal cases is tricky and risky when exposed to laypeople, and that the indexation of cases is more practical and realistic. By

answering each question, the user builds a hypothesis of how a judge may see their situation, that is then used to surface relevant information and previous cases. In order to allow the user to understand the caveats of the information (such as the fact that the court may diverge from the previous practice, and that every case is unique), the user is given disclaimers and clarifications with every step.

The system is further conceptualized to not require complex language or long texts. Of course, in the end the creator of each individual JusticeBot tool is responsible for writing the individual language. However, the system at each step selects only the relevant information, allowing the user to only need to understand the currently relevant information to respond to each question.

The system also relies on a simple and intuitive interface that responsively adapts to mobile devices.

7.8.2.2 The system should be able to handle areas of high-volume, low-intensity legal problems

Another criterion I specified was that the system should be able to deal with high-volume, low-intensity legal issues. This is important in order to have a significant impact on access to justice, and for the practical feasibility of the system.

I also believe that this criterion is fulfilled by the JusticeBot methodology. The methodology encodes the legal reasoning steps performed by judges in a schema representation. In order to build this representation, reading legislation and case law that is covered by this schema is a crucial step. Therefore, covering issues with many legal cases is advantageous, as the cases can be read to get an accurate overall view of the legal reasoning steps that are usually performed. Further, the methodology benefits from a number of cases being available for encoding, in order to provide illustration of the application of legal criteria, and outcomes tied to specific reasoning paths.

The methodology further works best when the syntactic ambiguity of a legal area is low, i.e. the steps performed by judges are relatively consistent between cases. This may be more likely in areas that are seen as high volume and low intensity, as the syntactic

structure in these areas are likely to have been applied many times and are thus likely to be clear. Further, the cases are likely to differ in the facts that are raised and the evidence provided, rather than legal arguments about specific criteria or the interpretation of complicated rules.

Of course, the real test of whether the methodology can handle areas of high-volume, low-intensity legal cases is practically implementing such an area. The implementation of a JusticeBot tool in the area of landlord-tenant disputes is described below in Chapter 8.

7.8.2.3 The system should give specific and useful information

The third design criterion is that the system should be able to give specific and useful legal information, based on legislation and case law. Let us explore whether the JusticeBot methodology can fulfill these criteria.

The JusticeBot can give three types of legal information: Information about the legal rights of the user, information about outcomes of previous cases, and information about possible next steps.

The system provides the user with information about their legal rights, as encoded by the creator of the system. This information can inform them, for example, that the legislation gives them the right to terminate their lease due to a certain factual situation. Or, it can inform them that their goal of subleasing their apartment is legally permitted. This information is specific to the situation of the user, since it relies on the questions that they answered. However, this also means that the accuracy of the information depends on the answers given by the user, and the hypothesis posed by the user of how their case will be seen by the judge. This, of course, can depend on a lot of factors, such as evidence, the specific situation etc. The user is reminded of this in the disclaimer. In a way, the user can thus be seen to *explore* the possible legal rights, based on certain legal criteria applying or not. Since laypeople are often not aware of the legal factors of a situation, this information could be very helpful in understanding what they can do and how to proceed.

The user is also given the outcomes of previous cases. These are selected by categorizing the cases based on the same legal criterion applying in previous cases, as in the current case. Again, the previous caveat applies – the outcomes will only be relevant if the user correctly responds to the questions. Further, showing the outcome of previous cases of course carries the caveat that judges may decide differently in future cases. However, the information can be important to contextualize the possible outcomes a user may face when going to court, and thus help them decide how to proceed.

Finally, the user receives a list of possible next steps that they can undertake. These are also defined in the JusticeCreator. The creator of the system can define which next steps should be integrated in the system, and how specific they should be. For example, some information may always be relevant (e.g. “speak to an attorney”), while some next steps may be relevant only in specific situations (e.g. “contact a support group for your specific issue”, such as a tenant rights organization).

All of the information above is to some extent dependent on the user correctly answering the questions posed by the system. Does this decrease the usefulness of the system? I believe that it does not, as long as this caveat is explained to the user. As we saw in the FactorBot, AI systems may not be able to accurately predict the application of legal criteria, based on a set of facts. The JusticeBot approach acknowledges this difficulty, and instead provides the user with a way of *exploring* their potential legal rights, previous outcomes and possible next steps. The system provides information augmenting the intelligence of the user.

This information may be relevant even if the assessment of the user does not match that of the judge. As we have seen, laypeople users often are not aware of the legal consequences of their situation. The JusticeBot system can provide them with this information, by informing the user of the potential legal rights that they have. The user can then rely on this information in settling their dispute with the other party, or deciding which court procedure to initiate and what to claim from the court. Likewise, the information regarding the next steps may inform the user of how they can enforce their rights and possible ways of proceeding with their case. Since this kind of information is a

necessity in enforcing their rights, it is useful for the user even if the judge in the end denies their claim (e.g. due to lacking evidence), since without the information they may not even have arrived at the point where their claim can be rejected based on the merits in the first place.

The same applies to the information about previous case law. Here, the user is provided with the outcomes of previous cases that were similar to that of the user. This information is not meant to inform the user of the result that they may obtain if they go to court. Rather, it is meant to provide them with relevant information that can help them in making better decisions. Thus, the case outcomes (such as the amount of damages awarded in previous cases) can serve to inform the user of the generally awarded outcome, given that they are able to prove the relevant facts and have correctly assessed the applicable legal criteria. While the system, again, does not attempt to predict the individual user case, this information can be useful in contextualizing the situation of the user, in order to give them a BATNA for settling their case, and a general estimation of the kind of outcomes that cases like theirs generally arrive at, which can be useful to decide whether to risk going to court or not.

Another possible criticism of the JusticeBot approach is that the cases whose outcomes are shown to the user lack specificity. For example, the user may be shown outcomes of cases that relate to mold assessed with regards to “peaceable enjoyment” of an apartment, while the user faces a situation relating to a water leak. Again, this approach makes sense in the context of exploring rights rather than predicting an outcome, as the user can be informed of the general outcomes that have been awarded based on a certain legal conclusion. Informing the user of outcomes specifically focused on their situation may run into issues such as the ones discussed above relating to the FactorBot. That said, exploring how the cases shown to the user at the end of their pathway could be made more specific is a fascinating question for future work.

In conclusion, I do believe that the JusticeBot methodology is able to provide useful and specific information to the user. The information is tailored to the situation of the user, allowing them to explore their potential rights, possible next steps and outcomes of

previous similar cases. Further, the information could be useful to the user, providing valuable information and context to help them choose how to proceed with their situation. Next, let us explore whether the system is practical.

7.8.2.4 The system should be practical

The final criteria set out in Chapter 5 is that the system should be practical. This includes that the methodology should be possible to apply to many legal areas (5.5.1), that it should be able to start by focusing on frequent types of cases (5.5.2), that it should focus on practical approaches to encoding legal information (5.5.3), that it should be implemented in intuitive interfaces, both for the creation of legal decision support tools (5.5.4) and the end-user of such tools (5.5.5), and that it should focus on giving legal information, rather than legal advice (5.5.6). Let us explore whether the JusticeBot methodology fulfills these criteria.

7.8.2.4.1 Can the methodology generalize?

I believe that the JusticeBot methodology can be applied to many legal areas. In Chapter 8, I detail the first deployed JusticeBot version, focused on landlord-tenant disputes. There are also several other versions under development at the Cyberjustice Laboratory. In 9.3, I explore a number of other legal areas that could be amenable to implement JusticeBot tools.

In essence, the JusticeBot methodology relies on encoding legal reasoning steps, as they are in reality carried out by judges or other legal decision makers. This kind of reasoning works well when there are structured steps that are carried out by judges in assessing new legal situations. Therefore, the methodology benefits from areas where there is a high degree of “syntactic consistency”, i.e. where the paths of reasoning that judges follow are consistent between cases. This may not be the case in courts of appeal, where the question in dispute may involve the interpretation of the law itself. Rather, areas of high-volume, low-intensity disputes are likely more well-suited for the methodology, as the disputes are contained within legal criteria, e.g. arguing over whether a specific legal criterion is fulfilled or not, in a specific case. Such reasoning is also likely to be present in administrative decision making, e.g. in determining whether a person deserves social

aid. I believe that such areas exist across legal systems – see 10.5.6 for more discussion. Of course, such areas are also some of the most common areas of disputes for laypeople – even if not every conceivable legal area can be treated in the JusticeBot, there are an enormous number of areas that can be treated, where building legal decision support tools could significantly increase access to justice.

Currently, only a single implemented JusticeBot version has been implemented and deployed to the public. However, the tools have been used in the Cyberjustice Laboratory and by a team in Italy to start the development of further versions, focused on other legal areas. The methodology seems to work well in these areas, particularly in family law and consumer protection. In the field of data protection, a new law is about to enter into force in Quebec, meaning that there are less decisions available for adding to the system. As currently conceptualized, the system focuses more on providing general information to the user. The JusticeCreator has worked well to implement this approach as well. Once completed and deployed, these tools will practically demonstrate that the JusticeBot methodology can generalize.

7.8.2.4.2 Can the methodology focus on frequent cases?

The methodology is also well suited to start with certain, frequent types of cases. This can be accomplished by identifying a legal issue that frequently arises, mapping the legal reasoning schema, and integrating relevant cases in the schema. Thus, the frequent types of cases can be integrated initially, in order to build a tool that can be useful even before all possible issue types have been encoded. In this manner, useful Justicebot tools can be built relatively quickly, and expanded over time to cover more legal areas. The feedback mechanisms can inform the creator of the system which such legal areas could be appropriate targets for adding to the system.

7.8.2.4.3 Is the encoding method of the JusticeBot practical?

Another important design criterion was that the approach of encoding legal information should be practical, in order to make it quicker and easier to create tools in new areas.

I believe that the JusticeBot fulfills this criterion. Rules are encoded in the form of a directed acyclic graph, which mimics the fashion that judges reason. This graph can be built in the JusticeCreator, with a visual WYSIWYG interface, allowing people without technical experience to encode new legal areas. The encoding method is relatively simple and does not rely on complex logical structures – rather, it adopts the logic of a flowchart, making it easy to reason about and understand. The schema can be verified and adapted as case law is integrated into the system.

The encoding of cases has also been designed to be practical. Cases are encoded based on the legal reasoning schema. Each case is read in conjunction with the schema, and cases are summarized in terms of the assessment a judge performs regarding a specific criterion. Since cases serve as an illustration of how a criterion is assessed by a judge, rather than to generate a model of the reasoning, the required number of cases is relatively low. Further, cases are summarized in terms of short textual summaries, which maintains the flexibility of natural language, overcoming the potential difficulties of encoding cases into fixed categories, as evidenced in the FactorBot.

The same cases can further be used to exemplify the outcome of legal cases. Here, the individual cases are tied to legal conclusions, in order to illustrate the outcome of previous cases similar to that of the user. The encoding is simple and easy to grasp, but allows the useful comparison of previous cases to the situation of the user. Since the cases merely provide an exemplification of previous outcomes, a relatively low number of cases need to be annotated for the system to be feasible.

In conclusion, I do believe that the approach to encode cases is practical. The methodology to encode the legal reasoning schema is relatively simple, while still allowing the simulation of real-world legal reasoning. Further, the methodology is designed to require as few cases as possible.

7.8.2.4.4 Does the methodology include interfaces for the creation of the JusticeBot and for the end-user?

Another important design criteria for the real-world use of the methodology is the development of interfaces that allow the use of the methodology, both for encoding information and for the end-user to interact with the system.

In order to implement the JusticeBot methodology in a practical manner, I have designed and built the JusticeCreator. It is a tool that allows the creation of schema pathways with a visual, easy-to-use interface, the addition of content such as questions, explanations and information, the encoding of case law, and even the AI-assisted retrieval of cases that may be relevant for inclusion in the pathway. The entire content of a new JusticeBot can thus be created in the JusticeCreator. Likewise, as legislation changes or the system is adapted to new use cases, the system can be used to update the existing JusticeBot.

Further, I believe that the JusticeCreator does not require a technical background for the creation of JusticeBot tools. The interface is fully visual, and does not require the use of programming to create new pathways. Of course, using the tool requires the user to understand how legislation and case law works, and how to logically encode these systems. However, as described in 7.8.2.4.3, the logical encoding of rules and cases in the JusticeBot is aimed to be practical and easy to reason about. Thus, I believe that legal experts should be able to learn how to use the tool relatively quickly, allowing the tool to be used in many contexts and have a potentially significant impact on access to justice. The system has already been used to build a legal decision support tool in the domain of landlord-tenant disputes, as will be explored in-depth in Chapter 8. Overall, the legal experts involved in the project were quickly able to grasp the functionality of the tool and use it to build the JusticeBot TAL.¹¹³⁹

Of course, for a system to be viable, it further has to include the faculties to expose an interface to the public, and thus allow lay people to interact with the system in an

¹¹³⁹ See 8.3.3.

intuitive manner.¹¹⁴⁰ For the JusticeBot, I therefore developed a web-based interface that is accessible via smartphones or computers, and allows the user to navigate the encoded legal information. Individuals respond to questions, and are then shown legal information corresponding to their situation, as well as previous relevant case law and possible next steps.¹¹⁴¹ The system is fully implemented and production-ready, and has been deployed to the public in the version focused on landlord-tenant disputes. Thus, the JusticeBot methodology is ready to be used to build tools that can be deployed to the real world, increasing the potential impact on access to justice.

7.8.2.4.5 Does the tool focus on giving legal information?

Another important design criteria for the legal decision support methodology is that it should not give legal advice, but rather legal information. Sticking to legal information ensures that the tool can be deployed to the real world and does not infringe on the exclusive right of lawyers to give legal advice.

While not entirely clear, the rules in many jurisdictions seem to indicate that the dividing line between legal information and legal advice lies at whether the information is targeting a specific situation or applies more generally.¹¹⁴²

Under this rule, the JusticeBot should be seen as providing legal information. At no point does the system try to predict the case of the user. Instead, it steps through the legal criteria that a court is likely to apply to a case and provides the user with examples from previous cases where the criterion is applied. All of this is general information, describing a legal decision, rather than applying the law to the case of a user.

The *user* then applies the law to their own case, by considering which of the cases are more relevant to them, and answering the questions in the system. As such, it is clear that the system provides only general information, while the user is the party that provides an opinion of their own case.

¹¹⁴⁰ See 5.5.5.

¹¹⁴¹ See 7.2.

¹¹⁴² See 3.5.3.

Finally, once the user has entered their hypothesis, the system provides them with general legal information based on their hypothesis, and summaries of previous cases. One might argue that this could be seen as giving a legal opinion, since the system provides information about the analyzed case of the user. However, this viewpoint misses a crucial nuance of the system, namely that the system explicitly does not provide the user with information regarding their case. Instead, it provides the user with general legal information about legal conclusions that judges tend to come to when certain legal criteria are fulfilled, and a list of summaries of previous legal cases that are similar to the hypothesis entered by the user. This allows the user to explore possible outcomes in previous cases, and gain an insight into the typical outcomes of certain cases, without the system trying to predict the case of the user, which could be seen as giving legal advice.

Since the system does not provide the user with advice, or tell them what to do, I believe that the criteria of the methodology providing legal information is fulfilled. Therefore, the practical application of JusticeBot tools is allowed, ensuring the possibility of increasing access to justice today.

7.9 Conclusion

This concludes my description of the JusticeBot methodology. Informed by the insights discussed in Chapter 6, this methodology takes a radically different approach. It starts with the encoding of a legal reasoning schema into a computer system, which represents the logical steps judges in reality tend to follow when dealing with certain types of cases. Using this schema, individual cases are encoded into the system, in order to illustrate how judges have previously determined whether relevant legal criteria apply or not, and the outcomes judges have awarded.

The user interacts with the system by navigating the schema of encoded reasoning steps and answering the question of how individual criteria may be applied in relation to their hypothetical case. Based on these answers, the system can provide the user with information and previous case law. However, the system never tells the user how judges will see their case, what the outcome will be or what they should do – rather, it acts like

an augmented intelligence system giving the user the tools to better understand their own case and their possible legal rights.

As mentioned, I have fully developed and implemented the tools for the legal expert to build a legal decision support tool using the methodology, and for users to be able to interact with the resulting system via a website. In the next chapter, we will explore the experience of building and publishing such a tool, relating to disputes between landlords and tenants.

Chapter 8 Case study: JusticeBot TAL

Research Objective: Validating the resulting methodology (1.2.2.7)

Research Topics:

- Does the methodology allow the implementation of legal decision support tools?
- Does the created legal decision support tool address the issues with access to justice and legal information in an area?
- How was the user experience of individuals interacting with the system?

8.1 Introduction

In the previous chapter, we explored the JusticeBot methodology. In this chapter, I will give an overview over the first legal decision support tool created using this methodology, targeting disputes between landlords and tenants, here referred to as the JusticeBot TAL. It is accessible to the public at <https://justicebot.ca>.

The JusticeBot TAL was created at the Cyberjustice Laboratory, together with a team of legal experts, led by Me Mark Likhten. It was developed together with the Tribunal Administratif du Logement du Québec, with help from Aide Juridique Montréal et Laval. The implementation and evaluation of the project received a grant from the Ministère de l'Économie et Innovation Québec. It was launched to the public on 20 July 2021, and has since been used over 17,000 times.

In this chapter, I will describe the particularities of the JusticeBot TAL. I will start by giving a background of the area of landlord-tenant disputes in Quebec, including a description of the housing market and potential issues faced by individuals, the applicable rules, and existing mechanisms for individuals to receive information and advice (8.2). Then, I will describe the JusticeBot TAL, including an examination of why landlord-tenant disputes are an appropriate area for the application of the JusticeBot methodology, a description of the available data sources, and an overview of the development and launch of the tool (8.3). Next, I will describe the feedback and analytics that have been collected from the system (8.4). Finally, I will wrap up and summarize this chapter (8.5).

The JusticeBot TAL is an important validation of the JusticeBot methodology, and a demonstration that it can be used to develop concrete legal decision support tools in socially relevant areas. However, it is far from the only area that can be tackled with the methodology. In Chapter 9, I will explore some other areas that could be relevant target domains for the JusticeBot methodology.

8.2 Background

This section will describe the general background for the area of landlord-tenant disputes. It gives an overview of the rental market and the relevance of rental disputes. I will also describe applicable laws and the forum available to deal with these issues. Finally, I will describe the ways that individuals can currently obtain legal information in the area.

8.2.1 The housing market in Canada and Quebec

Canada and Quebec face a shortage of affordable housing. Between 2000 and 2019, the median price of houses and condos across Canada increased by 195%.¹¹⁴³ The median cost of buying a house or a condo in Montreal in 2018 was 500,000 CAD.¹¹⁴⁴

These increases have further been exacerbated by the pandemic – in 2020, real estate sales jumped by 28%, the highest increase in 18 years.¹¹⁴⁵

This rise far outpaced the increase in wages. Between 2008 and 2017, the nominal median wage increased only by 22%.¹¹⁴⁶ In 2017 in Montreal, the price of a house or condo represented 16.6 times the median annual household income – up from 9.8 times in 2002.¹¹⁴⁷

¹¹⁴³ Louis Gaudreau, Guillaume Hébert & Julia Posca, “Analyse du marché de l’immobilier et de la rentabilité du logement locatif” (2020) Institut de recherche et d’informations socioéconomiques 20 at 2.

¹¹⁴⁴ *Ibid.*

¹¹⁴⁵ Christopher Curtis, “Renters are getting burned in Quebec’s red-hot housing market”, (27 May 2021), online: *Ricochet* <<https://ricochet.media/en/3664>>.

¹¹⁴⁶ Michal Rozworski, “The roots of our housing crisis: Austerity, debt and extreme speculation”, (14 June 2019), online: *Policy Note* <<https://www.policynote.ca/the-roots-of-our-housing-crisis-austerity-debt-and-extreme-speculation/>>.

¹¹⁴⁷ Gaudreau, Hébert & Posca, *supra* note 1143 at 3.

For many individuals and families, renting apartments thus becomes the only viable way to obtain housing. However, rental prices have also increased significantly. Between 2002 and 2020, the price of housing situated in apartment blocks constructed for rental has increased by 53%.¹¹⁴⁸ The average cost of renting an apartment with two bedrooms in Montreal increased from 760 CAD in 2015 to 855 CAD in 2019, according to the Société canadienne d'hypothèques et de logement (SCHL).¹¹⁴⁹ A study conducted by the “Coalition of Housing Committees and Tenants Associations of Quebec” (RCLALQ), analyzed the price of apartments that appear for rent on the web platform Kijiji. It found that the average price for an apartment listed during 2020 was 1,302 CAD, and the price for an apartment with two bedrooms 1,349 CAD.¹¹⁵⁰ While the vacancy rate of dwellings across the island of Montreal was 2.7% in 2020, this number falls to under 1% when only considering affordable apartments costing under 925 CAD per month.¹¹⁵¹

These issues hit especially hard for immigrants and refugees, who may face various forms of discrimination and may not know their rights.¹¹⁵² Similarly, the indigenous population may face issues with discrimination and resulting difficulties with finding housing.¹¹⁵³

There are certain aid programs available to help people find a rental unit. For example, the government offers the so-called Low-Rental Housing program, which offers housing capped at 25% of the individual's income. To qualify, households have to have assets of under 50k CAD and an income under a certain level (for example, 32,500 CAD for an

¹¹⁴⁸ *Ibid.*

¹¹⁴⁹ *L'habitation en bref 2020* (Société d'habitation du Québec, 2020) at 1.

¹¹⁵⁰ *Les loyers explosent - Enquête sur le prix des logements à louer au Québec* (Regroupement des comités logement et associations de locataires du Québec, 2021) at 4–6.

¹¹⁵¹ Curtis, *supra* note 1145.

¹¹⁵² Chloe Reiser, “Migrants bear the brunt of Canada's worsening housing crisis”, (23 September 2021), online: *openDemocracy* <<https://www.opendemocracy.net/en/pandemic-border/migrants-bear-the-brunt-of-canadas-worsening-housing-crisis/>>.

¹¹⁵³ Curtis, *supra* note 1145.

individual or a couple).¹¹⁵⁴ There is also a Rent Supplement Program that caps the price of renting certain housing options at 25% of the income of certain individuals. In Montreal, these programs are managed by the Office municipal d'habitation de Montréal (OMHM).¹¹⁵⁵ In 2020, the OMHM housed 55,000 tenants in 20,810 low rental housing units and 14,000 rent supplemented apartments. However, the waiting list was quite long at 22,972 households in line to obtain access to low-rental housing, with an average wait time of 5,4 years.¹¹⁵⁶

8.2.2 Potential issues faced by tenants and landlords

The increase in real estate prices, and scarcity of housing may make it difficult for tenants to obtain reasonable housing and enforce their rights. For example, apartments in Quebec are subject to strict rent controls, capping the allowed increase of rent by the landlords. However, due to the scarcity of apartments, tenants may decide to accept an illegal increase just to have a place to live.¹¹⁵⁷ Further, while the rent can only be increased by a certain percentage each year, tenants may not be aware of how much the previous tenant paid, to know how much of an increase they are facing.¹¹⁵⁸

In some cases, it may be beneficial for landlords to evict the tenants, combine or renovate apartments, and rent the apartments out at significantly increased prices. Landlords have two ways of legally evicting their tenant – either by claiming the apartment for their own use, or by aiming to change the use of the building or increase the size or subdividing the apartment. Challenges to the latter type of eviction notices increased by 142% between 2018-19 and 2019-20 according to the tribunal which handles these cases.¹¹⁵⁹ A study

¹¹⁵⁴ “Eligibility criteria (Housing Application)”, online: *Office municipal d'habitation de Montréal* <<https://www.omhm.qc.ca/en/submit-application/eligibility-criteria>>; *By-law respecting the allocation of dwellings in low rental housing*, S-8, r 1 Article 14.

¹¹⁵⁵ “Types of housing”, online: *Office municipal d'habitation de Montréal* <<https://www.omhm.qc.ca/en/about-us/types-housing>>.

¹¹⁵⁶ *Rapport annuel 2020* (Office municipal d'habitation de Montréal, 2021) at 27; Reiser, *supra* note 1152.

¹¹⁵⁷ Curtis, *supra* note 1145.

¹¹⁵⁸ Katelyn Thomas, “Rental hell in Montreal: Caught in a renoviction nightmare”, (3 July 2021), online: *Montreal Gazette* <<https://montrealgazette.com/news/local-news/rental-hell-in-montreal-caught-in-a-renoviction-nightmare>>.

¹¹⁵⁹ *Ibid.*

conducted by the “Comité logement de la Petite Patrie” claimed that many of the evictions or retaking of apartments are done under fraudulent or malicious pretexts, where the landlord claims to want to make a certain change but then does not follow through, instead often selling the apartment or renting it out to other tenants.¹¹⁶⁰

These numbers may be just the tip of the iceberg, as most such evictions may be settled directly between landlords and tenants, sometimes under pressure, with monetary incentives, or with insidious strategies such as the “renoviction”, where the landlord commences large renovation projects on the buildings, claiming the buildings to be unsafe, even when this is not the case. The renovations may cause significant noise, discomforts or safety hazards to the tenants, forcing them to leave the building.¹¹⁶¹

Individuals that are able to find an apartment to rent may suffer from issues in that apartment. For example, apartments may face issues with cleanliness, heating, infestation of bedbugs or vermin, or other issues that may make them unsuitable for living in. For example, in 2019, 2.8% of all households on the island of Montreal faced issues with bedbugs. This especially affected rental households, where 4% were affected in 2019, and even more so households with low income. 8.9% of households earning under 20,000 CAD were affected by bedbug infestations in 2019.¹¹⁶²

Of course, landlords may also face issues, including having tenants that do not pay or keep their apartment in a clean state. In the summer of 2022, the CEO of the Quebec Landlord Association claimed that 1/3rd of apartments where tenants moved out on July 1st needed deep cleaning to be used.¹¹⁶³

¹¹⁶⁰ *Entre fraude et spéculation - Enquêtes sur les reprises et évictions de logements* (Comité de logement de la Petite Patrie, 2020).

¹¹⁶¹ Thomas, *supra* note 1158; Thea McLachlan, “The Montreal apartment building on the front lines of Canada’s rental crisis”, (10 May 2021), online: *Ricochet* <<https://ricochet.media/en/3642>>.

¹¹⁶² *Les punaises de lit : État de situation à Montréal* (2019).

¹¹⁶³ Matt Grillo, “Some Quebec landlords frustrated at having to clean disaster apartments post moving day”, (7 July 2022), online: *CTV News - Montreal* <<https://montreal.ctvnews.ca/some-quebec-landlords-frustrated-at-having-to-clean-disaster-apartments-post-moving-day-1.5979162>>.

As we can see, there is potential for conflict between landlords and tenants. Quebec has laws governing the relationship between the lessor and lessee of a dwelling. Let us explore which rights and obligations these rules impose on the parties of a residential lease contract, and how these might affect the aforementioned situations.

8.2.3 The Law

The relationship between a landlord and a tenant typically begins with the entering into a lease. In Quebec, leases are regulated in the Civil Code of Quebec (CCQ).¹¹⁶⁴ Lease is here defined as:

*Article 1851 - Lease is a contract by which a person, the lessor, undertakes to provide another person, the lessee, in return for a rent, with the enjoyment of movable or immovable property for a certain time. [...]*¹¹⁶⁵

There are a number of general rules that apply to leases, such as general contract law. However, there are also specific rules relating to leases of dwellings set out in Division IV of Chapter IV of the CCQ.¹¹⁶⁶ These rules apply to leases of rooms, mobile homes or the land intended for the placement of mobile homes. There are also a few exceptions, such as dwellings in vacation resorts, rooms in hotels and certain dwellings that are part of the principal residence of the lessor.¹¹⁶⁷

In order to conclude a lease in Quebec, the use of a certain form provided by the Tribunal Administratif du Logement (the tribunal with the exclusive jurisdiction over cases regarding leases, see below) is mandatory.¹¹⁶⁸ However, this is not a formal requirement for a lease contract to be established – even verbal leases are valid, and must be

¹¹⁶⁴ CCQ, *supra* note 557.

¹¹⁶⁵ *Ibid* Article 1851.

¹¹⁶⁶ *Ibid* Chapter IV Division IV.

¹¹⁶⁷ *Ibid* Article 1892.

¹¹⁶⁸ *Regulation respecting mandatory lease forms and the particulars of a notice to a new lessee*, CQLR c T-1501, r3 Article 1.

confirmed in writing within 10 days of the agreement.¹¹⁶⁹ In the lease, the parties can agree to shape aspects of their relationship in certain ways through the contract. However, there are many clauses in the CCQ that are mandatory, and any contractual clauses trying to affect these are invalid.¹¹⁷⁰ This is especially the case for certain articles that protect the tenant. For example, it is impossible to limit the liability of the lessor or make the lessee liable for damages caused without their fault.¹¹⁷¹

The lease gives rise to a number of rights and obligations on the side of both the landlord and the tenant. While rental law is a complex area, and I am not able to cover all the particularities, let us take a look at some of these rights and obligations.

8.2.3.1 The rights and obligations of the landlord

The landlord has a number of important obligations towards the tenant. These include:

- On the agreed upon date, give the tenant access to the leased property, in clean condition, good state of repair and habitable condition.¹¹⁷²
- During the time of the lease, to give peaceful enjoyment of the leased property to the tenant¹¹⁷³. This includes making sure that the noise levels are acceptable,¹¹⁷⁴ The landlord must also maintain the dwelling in good habitable condition throughout the lease.¹¹⁷⁵

¹¹⁶⁹ CCQ, *supra* note 557 Article 1895; *Regulation respecting mandatory lease forms and the particulars of a notice to a new lessee*, *supra* note 1168 Article 3; “What is a lease?”, (18 December 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/signing-a-lease/what-is-a-lease>>.

¹¹⁷⁰ See e.g. CCQ, *supra* note 557 Article 1893.

¹¹⁷¹ *Ibid* Article 1900.

¹¹⁷² *Ibid* Articles 1854, 1910, 1911; “Rights and obligations of the lessor”, (20 November 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/being-a-lessor/rights-and-obligations-of-the-lessor>>; “Responsibilities of Landlords”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/responsibilities-of-landlords/>>; Pierre Gagnon & Isabelle Jodoin, *Louer un logement*, 2e ed (Cowansville, Québec: Éditions Y. Blais, 2012) at 8.

¹¹⁷³ CCQ, *supra* note 557 Article 1854.

¹¹⁷⁴ “Noise”, (18 December 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/the-dwelling/noise>>.

¹¹⁷⁵ CCQ, *supra* note 557 Article 1910.

- To make sure that the leased property can be used for the purpose it was leased.¹¹⁷⁶
- To make all necessary repairs to the leased dwelling. The tenant must take care of minor maintenance repairs, however.¹¹⁷⁷

If the landlord fails to comply with their obligations, the tenant may apply for specific performance or the termination of the lease, if the nonperformance causes them serious injury. The tenant may also apply for a reduction of rent in these cases.¹¹⁷⁸

The landlord also has a number of rights. These include:

- The right to ascertain the condition of the leased property, to carry out work thereon and to have the dwelling visited by potential lessees or acquirers¹¹⁷⁹
- The right to repossess a dwelling to use it as a residence for themselves or close family.¹¹⁸⁰
- The right to evict the tenant in order to subdivide the dwelling, enlarge it substantially or change its destination.¹¹⁸¹
- The right to terminate the lease if the tenant is over three weeks late in paying their rent, or if the tenant is frequently late with paying and the landlord suffers serious injury as a result.¹¹⁸² Likewise, the right to apply for termination of a lease if an apartment becomes unfit for habitation.¹¹⁸³
- The right to increase the rent of the dwelling at lease renewal.¹¹⁸⁴

8.2.3.2 The rights and obligations of the tenant

Likewise, the tenant has a number of obligations. These include:

¹¹⁷⁶ *Ibid* Article 1854.

¹¹⁷⁷ *Ibid* Article 1864.

¹¹⁷⁸ *Ibid* Article 1863; Gagnon & Jodoin, *supra* note 1172 at 25–28.

¹¹⁷⁹ *CCQ*, *supra* note 557 Article 1857.

¹¹⁸⁰ *Ibid* at 1957.

¹¹⁸¹ *CCQ*, *supra* note 557 Article 1859.

¹¹⁸² *Ibid* Article 1971.

¹¹⁸³ *Ibid* Article 1972.

¹¹⁸⁴ *Ibid* Article 1942.

- To pay the rent¹¹⁸⁵
- To use the dwelling with prudence and diligence¹¹⁸⁶
- To keep the dwelling in clean condition¹¹⁸⁷
- To make minor repairs to the dwelling¹¹⁸⁸
- To grant the landlord access to the apartment in order to ascertain the condition, or to have it visited by a potential acquirer¹¹⁸⁹
- To not disturb the peaceable enjoyment of other tenants¹¹⁹⁰
- To inform the landlord about defects or deterioration of the apartment¹¹⁹¹
- At the end of the lease, to remove their movable effects and leave the dwelling in the condition in which they received it.¹¹⁹²

If the tenant breaches any of their obligations, the landlord can request specific performance of the obligation, or terminate the lease if they are caused serious injury by the nonperformance of the obligation.¹¹⁹³

The tenant also has a number of rights. These include:

- The right to maintain occupancy of the dwelling¹¹⁹⁴
- The right to sublease their dwelling, or assign the lease to someone else¹¹⁹⁵
- The right to challenge a rent increase¹¹⁹⁶
- The right to be informed of any changes to the lease and if the landlord wishes to access the dwelling or conduct any major repairs¹¹⁹⁷

¹¹⁸⁵ *Ibid* Article 1855.

¹¹⁸⁶ *Ibid* Article 1855.

¹¹⁸⁷ *Ibid* Article 1911.

¹¹⁸⁸ *Ibid* Article 1864.

¹¹⁸⁹ *Ibid* Article 1857.

¹¹⁹⁰ *Ibid* Article 1860.

¹¹⁹¹ *Ibid* Article 1866.

¹¹⁹² *Ibid* Article 1978, 1890.

¹¹⁹³ *Ibid* Article 1863; Gagnon & Jodoin, *supra* note 1172 at 25–28.

¹¹⁹⁴ *CCQ*, *supra* note 557 Article 1936.

¹¹⁹⁵ *Ibid* Article 1870.

¹¹⁹⁶ *Ibid* Article 1949.

As we can see, the relationship between the landlords and tenants underlies a number of rules. It seems like several the situations described above, such as the landlord increasing the rent to an unfair extent, not taking care of the apartment or unfairly evicting the tenant, should be covered by the legislation. Likewise, the tenant is bound to follow the legislation, and e.g. keep the apartment in a clean condition.

Of course, for the rights lined out in the Civil Code of Quebec to be effective, the parties must be aware that these rights exist. We will examine the potential for this below in 8.2.5. Further, there must be a mechanism that allows them to enforce their rights against the other party. For rental disputes in Quebec, this function is performed by the Tribunal Administratif du Logement. Let us examine the role and procedures of this tribunal.

8.2.4 The Tribunal Administratif du Logement

8.2.4.1 Jurisdiction & Procedures

Cases regarding leases are heard by the Tribunal Administratif du Logement du Québec (TAL). Let us briefly examine the particularities of the tribunal and its procedure.

The Tribunal Administratif du Logement has exclusive jurisdiction to hear cases “respecting the lease of a dwelling”, where the value of the case is below 85,000 CAD.¹¹⁹⁸ Beyond this, the TAL has exclusive jurisdiction for certain cases, such as setting rent, renewing a lease, repossessing a dwelling etc.¹¹⁹⁹

The TAL also has several other functions. For example, it must inform landlords and tenants about their rights, promote conciliation between landlords and tenants, compile

¹¹⁹⁷ *Ibid* Article 1942, 1931, 1922.

¹¹⁹⁸ *Act respecting the Administrative Housing Tribunal*, CQLR c T-1501 Article 28.1; *CCP*, *supra* note 374 Article 35; Regarding the monetary limit, see also *Reference re Code of Civil Procedure (Que)*, art 35, 2021 Supreme Court of Canada.

¹¹⁹⁹ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 28.2; “The Tribunal administratif du logement (TAL or rental board)”, online: *Educaloi* <<https://educaloi.qc.ca/en/capsules/the-tribunal-administratif-du-logement-rental-board/>>.

statistics about the general housing situation and publish compendiums of decisions rendered by the TAL.¹²⁰⁰

8.2.4.2 Rules of procedure

The rules of procedure at the TAL are established in Chapters IV of the Act respecting the Administrative Housing Tribunal.¹²⁰¹ They are further elaborated in the Rules of procedure of the Administrative Housing Tribunal.¹²⁰²

In broad strokes, the procedure works as follows. A plaintiff who wishes to file a claim regarding their lease against a defendant first has to fill out a form with the TAL.¹²⁰³ The motion includes information such as who the plaintiff and the defendant are, what outcome the plaintiff desires and on what grounds.¹²⁰⁴ The plaintiff notifies the defendant (e.g. via registered mail or bailiff),¹²⁰⁵ and attaches proof of the notification to the file.¹²⁰⁶

At this point, the tribunal may invite the parties to a conciliation meeting,¹²⁰⁷ to try to get them to settle their dispute. If the parties fail to find an agreement, the case will go on to hearing at the tribunal.¹²⁰⁸ The tribunal may invite the parties to a case management conference to plan the proceedings and clarify the questions at hand.¹²⁰⁹

¹²⁰⁰ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 5.

¹²⁰¹ *Ibid.*

¹²⁰² *Rules of procedure of the Administrative Housing Tribunal*, CQLR c T-1501, r5.

¹²⁰³ “Procedures for filing an application”, (4 January 2016), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/filing-an-application/procedures-for-filing-an-application>>.

¹²⁰⁴ *Rules of procedure of the Administrative Housing Tribunal*, *supra* note 1202 Article 3; note 1203.

¹²⁰⁵ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 56; *Rules of procedure of the Administrative Housing Tribunal*, *supra* note 1202 Article 7; “Notification of an application to the other party”, (4 January 2016), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/filing-an-application/notification-of-an-application-to-the-other-party>>.

¹²⁰⁶ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 56.2; note 1205.

¹²⁰⁷ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 56.5.

¹²⁰⁸ “Conciliation process”, online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/conciliation-between-lessor-and-lessee/conciliation-process>>; “Hearings at the Tribunal administratif du logement”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/hearings-at-the-tribunal-administratif-du-logement/>>.

¹²⁰⁹ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 56.5.

The Tribunal then invites the parties to a hearing, or, if the parties consent, proceed on the record.¹²¹⁰ During the hearing, the parties must be present or represented by a spouse or advocate, or another mandatary in certain conditions.¹²¹¹ Just like the Small Claims Division of the Court of Quebec,¹²¹² no advocate may act in certain cases of recovery of small claims.¹²¹³

During the hearing, the Tribunal member calls the case, acknowledges the presence of the parties and proceeds with the hearing. The parties state their pretensions and introduce their witnesses. The Tribunal member should give impartial assistance to the parties.¹²¹⁴

After the hearing, the Tribunal member will render a decision within 3 months and send this to the parties.¹²¹⁵ This decision can be executed.¹²¹⁶

8.2.4.3 Statistics

The annual report of the TAL can give us an overview of the significance of landlord-tenant disputes in Quebec.¹²¹⁷ In total, 51,748 claims were introduced in 2020-21.¹²¹⁸ This is down from almost 70,000 in the preceding year,¹²¹⁹ possible due to the impact of the covid-19 pandemic.

The TAL reports statistics regarding the specific case types by the subject of the claim. The categories are as follows:¹²²⁰

- **Non-payment of rent** - Applications to terminate the lease due to non-payment of rent

¹²¹⁰ *Ibid* article 60.

¹²¹¹ *Ibid* article 72.

¹²¹² See 3.3.1.3.

¹²¹³ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 73.

¹²¹⁴ *Ibid* article 63.

¹²¹⁵ *Rules of procedure of the Administrative Housing Tribunal*, *supra* note 1202 article 41.1, 41.2; *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 article 79.

¹²¹⁶ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 article 83.

¹²¹⁷ Simard, *supra* note 993.

¹²¹⁸ *Ibid* at 3.

¹²¹⁹ *Ibid* at 1.

¹²²⁰ *Ibid* at 44.

- **Fixation et revision** - applications to change the annual rent or other bail conditions
- **Urgent civil matters** - applications for specific performance of a party, access to a dwelling, urgent repairs, evacuation, expulsion of a person who illegitimately inhabits an apartment, lease transfer or sublease, restitution of overpayment, repossession of apartments or evictions of tenants, or cases that involve the health or security of a person
- **Priority civil matters** - Lease termination for other reasons that are not qualified as “urgent civil matters”, rent reduction, urgent civil matters where damages are also claimed.
- **General civil matters** - Claims where the dispute does not affect occupancy of the dwelling (e.g. damages, reimbursement of expenses), claims involving parties that are no longer tenants or landlords.

Below, you will see a list of statistics related to these categories of claims, for the 2020-2021 season.

Type of case ¹²²¹	Cases introduced and restarted in 2020-21 ¹²²²	Average wait time for first audience ¹²²³	Average time between the opening and closing of a file ¹²²⁴
Non-payment of rent	24,094	2,6 months	7,8 months
Fixation et revision	7,250	13,2 months	
Urgent civil matters	20,404	2,1 months	
Priority civil matters		7,9 months	
General civil matters		11,5 months	

¹²²¹ *Ibid.*

¹²²² *Ibid* at 45.

¹²²³ *Ibid* at 6.

¹²²⁴ *Ibid* at 9.

We can see that individuals that want to use the court system to resolve their issue will in some cases have to wait several months for a hearing. In 2019, a study analyzing cases related to mold in apartments claimed that the median case might take several years to be resolved. 90% of the claimants had in fact left their apartment before the final decision was rendered.¹²²⁵

8.2.4.4 Conclusion

As we can see, the TAL offers a path to enforce the rights set out in the Code Civil du Quebec regarding rental disputes. The tribunal has exclusive jurisdiction with regards to many rental disputes. It follows a procedure that focuses on simplicity, ease and speed, allowing tribunal members to overlook procedural mistakes and not allowing paid legal support in the court session.

The statistics presented by the TAL show the high demand the tribunal faces. Tens of thousands of cases are filed each year, and the average time to resolve a case ranges from months to years. Previously, we discussed how wait times in courts and the adversarial process can cause psychological stress to people, or cause them to skip the process altogether, meaning that a perceived injustice is not dealt with.

Of course, going to the TAL requires the individual to be aware of their rights with regards to their situation. Tenants may not know that they can obtain damages or terminate their lease due to issues with their apartment, or that they can contest an illegal rent increase or renoviction. In this case, they are unlikely to file their case with the TAL and are in a weak position to negotiate with their landlord. Even tenants who are aware of their rights may have difficulty interpreting exactly what those rights entail (i.e. should loud noises be seen as a loss of the “peaceable enjoyment” of an apartment?) and how to deal with the procedure at the TAL. Likewise, landlords may not know how to deal with tenants that refuse to pay or take bad care of the rented property.

¹²²⁵ Katia Gagnon, “Plaintes à la Régie pour des logements moisés: des années d’attente”, (12 March 2019), online: *La Presse* <<https://www.lapresse.ca/actualites/sante/201903/12/01-5217878-plaintes-a-la-regie-pour-des-logements-moisés-des-années-dattente.php>>.

Martine Dubé, an indigenous woman interviewed in the news paper “ricochet”, sums up the situation in the following poignant quote:

*“If you’re out there, looking for an apartment, you have to know your rights. There was a time where I didn’t. Now I have to. We all do.”*¹²²⁶

In Quebec, there are a number of ways for individuals to obtain legal support and information in the area of landlord-tenant disputes.

8.2.5 Legal support and information for parties in landlord-tenant disputes

As we have established, there is a significant need for legal support and advice for people who face potential rental disputes. In Quebec, there are multiple avenues for individuals to obtain this support and information. As described above, in Quebec only members of the Barreau du Québec are able to give legal advice and consultations, and draft documents that will be used in court.¹²²⁷ This limits the level of support certain sources can give to individuals.

8.2.5.1 Legal advice

The most traditional way of obtaining legal advice and support is through hiring a lawyer. There are many lawyers and law firms that specialize in rental disputes and are able to help the individual with explaining their rights, drafting legal documents and letters and helping them understand the court procedure.

However, hiring a lawyer can be very expensive, often costing hundreds of dollars per hour.¹²²⁸ This may be out of reach for a large part of the population. Further, since rental disputes often do not deal with huge sums, hiring the lawyer might cost more than one can expect to gain from the case. According to a study conducted in 2009, only 11,4% of individuals in Canada chose to confer with a lawyer when faced with legal issues.¹²²⁹

¹²²⁶ Curtis, *supra* note 1145.

¹²²⁷ *Act respecting the Barreau du Québec*, *supra* note 510 section 128.

¹²²⁸ “JuridiQC”, online: *JuridiQC* <<https://juridiqc.gouv.qc.ca>>.

¹²²⁹ Currie, *supra* note 297 at 56.

The government in Quebec offers free or subsidized legal support for individuals below a certain threshold of income.¹²³⁰ The threshold for free legal support is currently at 24,570 CAD for a single person, although the family situation and the value of different assets can also play a role.¹²³¹ Both employees of the legal aid network and lawyers in private practice can offer legal support through this system, and are able to give both advice and represent the clients in court. According to the Legal Aid website, around 250,000 applications for legal aid are processed each year.¹²³² The legal aid network brings together 11 community legal aid centers. An example of such centers is “Legal Aid Montreal | Laval”, which has 115 lawyers and 158 non-lawyer employees.¹²³³

Certain lawyers also support individuals for free for a certain number of hours per year (so-called pro-bono work). In Quebec, this is organized through the non-profit organization “Justice Pro-Bono”.¹²³⁴ There are also organizations that offer subsidized legal support for individuals that do not qualify for legal aid but still cannot afford to hire a lawyer, such as JuriPop, which offers services at 65 CAD per hour.¹²³⁵ Juripop claims to have supported over 4,000 individuals in court, and to have informed over 26,000 individuals of their rights.¹²³⁶

8.2.5.2 Legal Information

There are also a number of important sources for legal information for individuals. While these are not able to provide assessment of individual cases, they are nonetheless able to inform the individuals of their rights and let them make the decision on how to proceed, including hiring a lawyer to take the case to court if the outcome of the legal information warrants it.

¹²³⁰ note 365.

¹²³¹ “Legal Aid - Am I financially eligible?”, online: *Comissions des services juridiques* <<https://www.csj.qc.ca/commission-des-services-juridiques/aide-juridique/volet-gratuit-aj/en>>.

¹²³² note 365.

¹²³³ “What is the CCJM?”, online: *Legal Aid Montreal | Laval* <<https://www.aidejuridiquedemontreal.ca/en/legal-aid-montreal-laval/>>.

¹²³⁴ “Our Mission”, online: *Justice Pro Bono* <<https://justiceprobono.ca/en/aboutus/our-mission/>>.

¹²³⁵ “About us”, online: *Juripop* <<https://juripop.org/en/about-us/>>; “I Need a Lawyer - Eligibility”, online: *Juripop* <<https://juripop.org/en/i-need-a-lawyer-eligibility/>>.

¹²³⁶ *Rapport Annuel 2019-2020*, by Sophie Gagnon (JuriPop).

As mentioned, the TAL has an important role in informing citizens on their rights. In 2020-21, 285,309 phone calls were handled, and 140,159 correspondences per email, post or fax were treated.¹²³⁷ In addition, 625,490 calls were treated by an interactive call response system installed by the TAL.¹²³⁸ This shows the enormous desire for information in the area of landlord-tenant disputes. Of course, as the tribunal, the TAL has to be neutral in providing information.

There is also a network of community-run legal aid clinics, that provide free legal support for individuals and can help tenants understand their rights. The RCLALQ lists 20 such associations on the island of Montreal.¹²³⁹ Unlike the previously mentioned institutions, they are often limited to provide legal information, but cannot give legal advice or represent clients in a court. An example of such a clinic is Project Genesis, which assists thousands of people yearly with often poverty-related issues.¹²⁴⁰ In 2021, Project Genesis made a total of 8,457 interventions (including in-person and via phone), 57% of which were related to housing issues.¹²⁴¹ This shows the prevalence of these issues among people needing support. Other legal aid organizations include “Community Justice Centers”¹²⁴², a call-in legal clinic offered by the Young Bar of Montreal,¹²⁴³ the Mile End Legal Clinic¹²⁴⁴ and the Mobile Legal Clinic.¹²⁴⁵

There are also online resources to explain legal rights. The TAL itself has explanations of legal rights of landlords and tenants on their website.¹²⁴⁶ Éducaloi is a neutral and independent organization that aims to “Explain the law to Quebecers in everyday

¹²³⁷ Simard, *supra* note 993 at 3.

¹²³⁸ *Ibid* at 50.

¹²³⁹ *Les comités logement et associations de locataires de l’île de Montréal* (Regroupement des comités logement et associations de locataires du Québec, 2019).

¹²⁴⁰ *Project Genesis Annual Report 2020-2021* (Project Genesis, 2021) at 5.

¹²⁴¹ *Ibid* at 9.

¹²⁴² “Centres de justice de proximité”, online: <<https://www.justicedeproximite.qc.ca/en/>>.

¹²⁴³ “Call-in Legal Clinic”, online: *Jeune Barreau de Montréal* <<https://ajbm.qc.ca/en/public-services/annual-legal-helpline/>>.

¹²⁴⁴ “Clinique juridique du Mile End | Mile End Legal Clinic”, online: <<http://justicemontreal.org/>>.

¹²⁴⁵ “Mobile Legal Clinic”, online: *cji-mlc* <<https://www.cji-mlc.org>>.

¹²⁴⁶ “Questions fréquentes”, online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/fr/questions-frequentes>>.

language and enhance their legal competencies”.¹²⁴⁷ *Éducaloi* employs 31 professionals and a number of students,¹²⁴⁸ who write articles explaining the law in simple and clear language. Among other areas, *Éducaloi* features a significant section on landlord-tenant disputes.¹²⁴⁹

Of course, just like other sources of legal information, web-based sources are limited to provide legal information rather than legal advice. They cannot give advice for specific cases. Further, they rely on users being able to read the provided texts and perform the logical thinking necessary to understand how the information may apply to an individual case.

8.2.6 Conclusion

In this section, I have given an overview over the rental market in Quebec and Canada, examined some possible issues that may arise for tenants and landlords, and described the avenues available to solving these issues.

As we have seen, there are many issues that can arise in rental situations. For example, tenants may be unduly evicted, have their rent increased illegally or be forced to live in unmaintained apartments.

The law provides protections for both tenants and landlords. Enforcing these protections requires a forum to raise and resolve these issues, and knowledge that the rules exist and how they apply.

For rental disputes, the Tribunal Administratif du Logement has exclusive jurisdiction in Quebec. It offers a procedure for people to file a claim with the tribunal and obtain a judgment on the dispute.

However, without the knowledge of the rights accorded to the parties and knowledge of the procedure, exercising their rights, or the decision to go further in exercising their

¹²⁴⁷ “About Us”, online: *Éducaloi* <<https://educaloi.qc.ca/en/about/>>.

¹²⁴⁸ “Our Team”, online: *Éducaloi* <<https://educaloi.qc.ca/en/our-team/>>.

¹²⁴⁹ “Renting”, online: *Éducaloi* <<https://educaloi.qc.ca/en/categories/renting/>>.

rights, can be a difficult process for laypeople. The most traditional way of obtaining support is, of course, to hire a lawyer. However, this may be out of reach for a lot of individuals. There are various programs and services to provide cheaper legal support, but these may have limited capacity and low cut-off limits. There are also various community driven legal aid clinics that offer legal information for free. These are, however, constrained by what kind of information they can give, and by the need for qualified individuals and volunteers. Of course, using any of these sources pre-supposes that the individual is aware that their situation has a legal character and how to find more information, which may not always be the case.

As we can see, there is a need for novel ways of giving legal information to individuals involved in landlord-tenant disputes in Quebec. Therefore, we decided to build a JusticeBot-based decision support tool in the area. Such a tool can ask the user questions and give them specific information, using the methodology described above. Let us take a look at the JusticeBot TAL.

8.3 The JusticeBot TAL

As we have seen, there is a significant need for legal information in the domain of landlord-tenant disputes. The JusticeBot methodology, described in Chapter 7, can be used to build such tools. In this section, I will describe the development of the JusticeBot TAL, which uses the JusticeBot methodology to increase access to justice in the domain. First, I will discuss why landlord-tenant disputes are an appropriate area for implementing such a tool (8.3.1). Then, I will describe the data that is used for the development of the system (8.3.2), the development process (8.3.3), the resulting decision support tool (8.3.4), and the public launch of the tool (8.3.5).

8.3.1 Appropriate area for application of methodology

As discussed in 5.3, the JusticeBot methodology is designed to be able to handle areas of high-volume, low-intensity disputes. These requirements are helpful for the following reasons:

Cases should be high volume – For an area to be ideal for the creation of a legal decision support tool using the JusticeBot methodology, it should be an area with a high volume of cases. First of all, this means that many individuals will have to deal with this kind of cases, and that building tools to support them can increase the access to justice for a lot of individuals.¹²⁵⁰ Further, there are practical reasons for targeting areas with a high volume of cases. In such areas, semantic ambiguities in the rules are likely to have been resolved by repeated application in court. Further, larger number of cases mean that cases that are similar to that of the user are likely to have arisen, making it feasible to provide similar cases to the user and accurately assess their situation.¹²⁵¹ Also, the higher the number of cases in which users receive support for pretrial or out-of-court resolution, the more positive impact this support has on the court system's workload.

Cases should be low-intensity – Another important factor is that the area should contain cases of low intensity. These are the type of everyday legal cases that are likely to affect laypeople, and often lead to people representing themselves due to the comparatively low claim values and potential lack of resources. Building decision support tools that target these areas can thus give an important improvement to the status quo.¹²⁵² Further, targeting areas of low-intensity cases is important from a practical perspective, since these cases are less likely to involve complicated legal interpretation and policy reasoning, which may be intractable using current artificial intelligence approaches.¹²⁵³

Let us examine whether the area of landlord-tenant cases fulfills these criteria.

¹²⁵⁰ See 5.3.1.1.

¹²⁵¹ See 5.3.2.3 and 5.3.2.4.

¹²⁵² See 5.3.1.2 and 5.3.1.3.

¹²⁵³ See 5.3.2.2.

8.3.1.1 High volume of cases

In 8.2.4.3, I reviewed the statistics of the cases heard at the Tribunal Administratif du Logement. In 2020-2021, 51,748 claims were introduced.¹²⁵⁴ Other years have seen even more cases introduced.¹²⁵⁵

This certainly fulfills the requirement of being a significant number of cases. Many individuals are likely affected by issues in rental relationships, as can be seen by the high number of cases. Further, the number does not even cover cases that do not go to court, possibly due to the lack of understanding of their legal rights on behalf of the parties. Therefore, the area fulfills the requirement of having a high volume of cases.

However, for the use of artificial intelligence to be viable, the data further has to be accessible in a computer-readable format. Here, we are very lucky to be collaborating with the Tribunal Administratif du Logement, which has granted us access to a collection of 899,522 unique cases. A more in-depth description of the content of this data follows in 8.3.2. Having access to almost 900k decisions is very helpful in building legal decision support tools. While there, of course, is a wide array of types of cases covered by the TAL, the high number of accessible cases also likely means that there will be clusters of cases for some common individual legal questions. This should give us a good source of material to extract the patterns from the individual legal questions.

Based on the presented data, it seems like rental disputes in Quebec is a high-volume area of law. Thus, it should be well-suited for the creation of legal-decision support tools using the JusticeBot methodology.

8.3.1.2 Low intensity of cases

Another important issue to consider is whether the cases are of low intensity. The types of cases that are handled at the TAL are described in 8.2.4.3. Based on the case description, it seems like the cases may often be what can be referred to as low-intensity

¹²⁵⁴ Simard, *supra* note 993 at 3.

¹²⁵⁵ *Ibid* at 1.

cases. While they can of course be very important for the parties involved, the high number of cases and the type of cases may indicate that the cases are more focused on factual situations, rather than requiring complicated legal reasoning or policy decisions.

Another indication that supports this assessment is the rules regarding the Tribunal Administratif du Logement. The amount of damages claimed at the tribunal is capped at 85,000 CAD.¹²⁵⁶ Further, in some case, where the sole aim of a claim is the recovery of a debt below 15,000 CAD, the individual cannot be represented by an advocate.¹²⁵⁷ Explicitly excluding complex cases and the option of legal representation in certain cases could indicate that the issues dealt with at the TAL tend to not involve complex situations or high-value claims.

Further, the procedure is designed to furnish a process that is “simpler, easier and faster, while still respecting the basic principles of justice and equality for both parties.”¹²⁵⁸ It gives the member of the tribunal the possibility to overlook errors in procedure or delays by the parties as long as these are remedied while still possible.¹²⁵⁹

These statistics and rules seem to indicate that the cases heard at the TAL are typically low-intensity cases, and thus well suited for the application of the JusticeBot methodology.

Now that we have determined that the area of landlord-tenant disputes is well suited for the building of a JusticeBot based tool, let us analyze the corpus of cases. This will provide important background information as for how the tool should be built, such as which types of cases frequently appear in the tribunal.

¹²⁵⁶ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 Article 28.1; *CCP*, *supra* note 374 Article 35; Regarding the monetary limit, see also *Reference re Code of Civil Procedure (Que.)*, art. 35, *supra* note 1198.

¹²⁵⁷ *Act respecting the Administrative Housing Tribunal*, *supra* note 1198 article 73; *CCP*, *supra* note 374 article 536.

¹²⁵⁸ *Rules of procedure of the Administrative Housing Tribunal*, *supra* note 1202 Article 1.

¹²⁵⁹ *Ibid* Article 2.

8.3.2 The data

This section will briefly describe the dataset that was used in the construction of the JusticeBot. It consists of around 900k decisions furnished by the Tribunal Administratif du Logement, covering cases between 2001 and 2020. I am very grateful for being given access to this dataset. The cases were provided in the format of Words documents. For this analysis, I have extracted the text content of these files into a database.

In general, the documents contain the following elements:¹²⁶⁰

- A header containing metadata such as the date, the court and the names of the parties
- Procedural background
 - A list of the demands of the plaintiff
 - The procedural history of the case
- The relevant facts:
 - The factual circumstances the plaintiff bases his demand on
 - The competing version of the factual circumstances as claimed by the defendant
- Analysis of the facts:
 - An overview over the applicable legislation
 - The establishment of which factual version the tribunal sees as proved
 - The application of the law to the established factual version.
- Outcome:
 - A list of remedies ordained by the tribunal

The data allows us to understand important specifics of the area of landlord-tenant disputes. This can help us assess the field in an empirical way and make decisions on

¹²⁶⁰ Westermann et al, *supra* note 536 at 4.

how the JusticeBot should be built. I will briefly give an overview of some insights I gained from the data.

8.3.2.1 Constellations statistics

The first question the data allows us to understand is which types of party constellations are generally involved at the TAL. In order to gain this insight, I created a computer program that analyzes the header of the decision, to understand who is listed as the “partie demanderesse” (plaintiff) and “partie defenderesse” (defendant). The algorithm uses regular expressions and other search methods to identify the parties and add it to a database for analysis.

For this analysis, I decided to focus on only constellations involving landlords and tenants, and exclude constellations that involve roommates, occupants or co-tenants. This, together with the fact that formatting can differ between different cases, means that a total of 1.2% of the cases are excluded from the analysis. This shows how difficult the analysis of legal texts can be – even the formulaic field of who is the plaintiff and who is the defendant can be more complex than expected.

The results of the analysis can be seen in Table 7.

Plaintiff	Defendant	Number (%)
Landlord	Tenant	782,895 (89,1%)
Tenant	Landlord	93,818 (10,6%)

Table 7 - Distribution of landlords suing tenants vs tenants suing landlords

As we can see, out of the cases I analyzed, almost 90% are feature a landlord suing a tenant. This matches the finding of Salaün *et al.*¹²⁶¹

8.3.2.2 Claim statistics

Now that we have analyzed the constellations, let us explore the claims by the parties, i.e. which outcome they ask for. Since the claim is usually described in the beginning of the case, I analyze this information by searching the first 400 characters of each case for

¹²⁶¹ Salaün et al, *supra* note 621.

certain keywords that were empirically determined. In total, I identified 13 common claims, each with 1-6 keywords that identify them. Luckily, the language used here is quite formulaic, although for 12k cases (around 1.4%), no claim could be identified. While it is difficult to determine the accuracy of the analysis, it should suffice for an overview of the types of claims introduced.

In total, the algorithm identified 2,399,214 claims in 888,201 analyzed cases, bringing the average number of claims per case to 2.7. Let us look at the top 5 claims for landlords and tenants.

Claim	Frequency
Recovery of unpaid rent	579,231
Termination of lease	574,151
Expulsion of tenant	542,487
Damages	224,276
Reimbursement of judicial costs	78,840

Table 8 - Top 5 most frequent claims where the plaintiff was a landlord

As we can see in Table 8, the overall majority of cases filed by landlords involved recovery of unpaid rent, termination of lease and the expulsion of the tenant. These three claims seem to occur in cases where the tenant has stopped paying the rent or is currently or frequently late with payment, causing the landlord to want to expel them. This being the most frequent claim type matches the statistics in the annual report of the TAL.¹²⁶²

Claim	Frequency
Damages	39,441
Retraction of previous decision	31,031
Rent reduction	27,274
Order of specific performance	21,553
Termination of lease	12,064

Table 9 - Top 5 most frequent claims when the plaintiff was a tenant

The statistics for tenants (shown in Table 9) show a different set of claims. The most common claim is that of damages, followed by the retraction of a previous decision and rent reduction. Tenants also frequently want the court to order the landlord to do

¹²⁶² Simard, *supra* note 993 at 45.

occurring topics in the cases. However, future work is needed to more reliably understand the important topics in the decisions.¹²⁶³

We have determined that landlord-tenant disputes are an appropriate area for implementing a JusticeBot-based tool and explored the data to gain some initial insights. Next, I will describe the development process that led to the creation and implementation of the JusticeBot TAL.

8.3.3 Development process

The JusticeBot TAL was developed at the Cyberjustice Laboratory at Université de Montréal, in collaboration with the Tribunal Administratif du Logement. It is the results of a multi-year research project started in the summer of 2017.

The project was awarded a grant by the Ministère de l'Économie et de l'Innovation du Québec, which focuses on the evaluation and integration of the system with the procedures of our partners, the Tribunal administratif du logement du Québec (the court with jurisdiction for disputes regarding leases in Quebec) and Aide juridique de Montréal et Laval.

The development procedure of the JusticeBot TAL largely followed the steps outlined above under 7.3.3. A team of legal experts used the JusticeCreator system to encode the legal rules that judges tend to apply to legal issues, as described above in 7.3.3.4. The pathway was developed by consulting the law, online sources and case law. We were very lucky to have the permission of the TAL to include content from their online sources in the pathway, which provided an efficient way to integrate accurate simplified legal information.

Then, case law was integrated into the system, benefiting from the large number of cases provided to us by the Tribunal Administratif du Logement. The machine learning case retrieval method implemented in the JusticeCreator (see 7.3.2.4.1) proved useful to

¹²⁶³ Compare Salaün et al, “Why Do Tenants Sue Their Landlords?”, *supra* note 1098.

identify cases that are relevant for the individual pathways. The cases were used to add summaries and outcomes to the different pathways, as described above in 7.3.3.5.

The process was generally very smooth. After a brief introduction, the legal experts were quickly able to learn the features of the JusticeCreator, and how to use it to create legal decision support tools. Being accessible to non-programmers was one of the key goals of the JusticeCreator. Based on this experience, the goal seems to have been achieved.

As described in 7.3.3.6, it is important to verify that the content is accurate before launching a decision support tool. In the case of the JusticeBot TAL, this process consisted of exporting the content in the schema (including questions, descriptions, case law and the connections between the elements) into a word document. This document was sent to the Tribunal Administratif du Logement, who graciously agreed to read and verify the accuracy of the legal information. They commented and adjusted certain sections, which were re-incorporated into the JusticeBot TAL schema.

8.3.4 Pathway

The development process described above resulted in a schema that contains a total of 127 questions and 146 information blocks. Figure 49 contains an overview over the schema, as seen in the JusticeCreator. For the purposes of visualization, all pathways have been placed on the same page – normally, it would be split into multiple pages each covering a single pathway. Even so, the tools in the JusticeCreator make it trivial to navigate the pathway and find individual questions or preview how the resulting system will work when deployed.




Figure 49 - An overview over the JusticeBot TAL schema

A path of the JusticeBot TAL was used as an example above in 7.2.2. The schema starts with an introduction screen, after which the user is asked whether they are a landlord or a tenant. Depending on the choice, they are given a menu of possible legal pathways that they may want to explore. Figure 16 showcases these menus. Each of the options points

to a more in-depth pathway. Multiple options may point to the same pathway, but be formulated in a different way, or contain a more specific formulation that points to a further point in the pathway. For example, the option “I would like to leave my apartment before the end of the lease” gives the tenant multiple options to pick regarding whether to assign the lease or to terminate the lease. The user may also directly pick the option “I wish to assign my lease”, which would take them directly to the path regarding the lease transfer. In this manner, the user of the system can find the correct path even though they think of their issue in different ways.

In total, the following pathways are included in the JusticeBot TAL:

- Tenant:
 - The apartment is infested with bedbugs
 - The landlord wishes to raise the rent
 - Can I terminate my lease?
 - Can I sublet my apartment?
 - Can I transfer my lease?
 - My landlord wishes to undertake work in my apartment or building.
- Landlord:
 - The tenant wishes to terminate their lease.
 - The tenant wishes to sublet their apartment.
 - The tenant wishes to transfer their lease.
 - The tenant is late with paying their rent.

As you can see, many common issues are considered from both the perspective of the tenant and the landlord. Many of the issues further correspond to the common types of issues at the TAL, as determined by the statistics in 8.2.4.3 and the data analyzed in 8.3.2.2.

Several of the paths have further sub-paths or are inter-connected. As described above in 5.2.3, layperson users can think of their situation in terms of facts or in terms of outcomes that they would like to achieve. The aforementioned pathways contain both these modes.

An apartment being infested with bedbugs, or the landlord attempting to raise the rent focus on a *situation* and allow the exploration of possible rights in these situations. Other questions, such as “Can I terminate my lease?” focus on an *outcome* the user wishes to achieve, and perform a check to see whether this may be a possibility.

After clicking one of the options, the user is taken through a number of questions to better understand the situation. Many of the questions in the pathway have summaries of how the legal criteria were previously reasoned about attached to them. In total, 46 of the questions have a total of 157 such summaries linked to them. Overall, the questions in the pathway contains a significant amount of content – the explanations of what a question means contain enough content for almost 50 pages of text.

Then, their situation is analyzed to provide relevant information about their legal situation to the user, by selecting the information blocks integrated in the pathway the user traversed.¹²⁶⁴ An example of how this can look can be seen above in 7.2.3. The information provided informs the user of what their rights might be, and links to the relevant legislation. Many of the information blocks have the outcomes of previous cases attached to them. In total, 44 of the information blocks have a total of 165 case outcome examples linked to them. The explanations contained in the information blocks correspond to around 93 pages of written text.

8.3.5 Launch

The JusticeBot TAL was launched to the public on the 20th of July 2021 on <https://justicebot.ca>. It was featured and shared on the website of the Cyberjustice Laboratory¹²⁶⁵ and featured on the Tribunal Administratif du Logement.¹²⁶⁶ It was further

¹²⁶⁴ See 7.5.

¹²⁶⁵ “Communiqué de presse - Lancement du JusticeBot”, (21 July 2021), online: *Laboratoire de cyberjustice* <<http://web.archive.org/web/20210721123006/https://www.cyberjustice.ca/2021/07/20/communique-de-presse-lancement-du-justicebot/>>.

¹²⁶⁶ “JusticeBot – Interactive legal information tool”, (12 August 2021), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/justicebot-interactive-legal-information-tool>>.

shared on the faculty of law of Université de Montréal page¹²⁶⁷ and Université de Montréal news page.¹²⁶⁸ The launch received significant media attention, with articles appearing among others in the Journal de Montreal,¹²⁶⁹ 24 heures,¹²⁷⁰ droit inc¹²⁷¹ and Radio Canada.¹²⁷² The strong launch lead to a significant influx of curious users, as can be seen below in Figure 50. After this, the user number settled on a stable level, as will be further explored in 8.4.1.1.




Figure 50 - Statistics JusticeBot uses from July 2021 to March 2022

Since the launch, we have been working on improving and expanding the JusticeBot TAL. In September 2021, we added a pathway covering renovations.¹²⁷³ Further, the website

¹²⁶⁷ “JusticeBot, une application au service des locataires et propriétaires québécois!”, (16 November 2021), online: *Université de Montréal* <<https://droit.umontreal.ca/en/faculty/communications/news-details/news/detail/News/justicebot-une-application-au-service-des-locataires-et-proprietaires-quebecois/>>.

¹²⁶⁸ Virginie Soffer, “Un robot répond à vos questions sur le droit du logement”, (26 July 2022), online: <<https://nouvelles.umontreal.ca/article/2021/07/26/un-robot-repond-a-vos-questions-sur-le-droit-du-logement/>>.

¹²⁶⁹ François Carabin, “JusticeBot, un nouvel outil web pour aider les locataires et les propriétaires à connaître et à faire respecter leurs droits”, (20 July 2021), online: *Le Journal de Montréal* <<https://www.24heures.ca/2021/07/20/nouvel-outil-en-droit-du-logement-la-jurisprudence-au-bout-des-doigts/>>.

¹²⁷⁰ *Ibid.*

¹²⁷¹ Audrey Bonaque, “JusticeBot : un nouvel outil utile et intelligent”, (26 July 2021), online: *Droit-inc* <<https://www.droit-inc.com/article29094-JusticeBot-un-nouvel-outil-utile-et-intelligent/>>.

¹²⁷² *Lancement de la plateforme JusticeBot : Entrevue avec Karim Benyekhlef* (Radio Canada, 2021).

¹²⁷³ “JusticeBot répond maintenant à vos questions sur les « renovations » !”, (17 September 2021), online: *Laboratoire de cyberjustice*

has received adjustment and updates to address small issues and make the user experience smoother.

8.4 Feedback and Analytics

As described above in 7.7, the JusticeBot platform contains a number of mechanisms that allow for the collection of statistics and feedback. These can provide an important indication of how users interact with the system and which areas work well to increase access to justice, and which areas might need improvement. Let us analyze the analytics and feedback collected through the JusticeBot TAL.

8.4.1 Analytics

One of the most important feedback mechanisms in the JusticeBot is Google Analytics. It is set up to collect anonymous information about user interactions with the system.

8.4.1.1 General statistics

Google analytics allows us to see statistics for one year back. Therefore, most of the statistics in this chapter will correspond to a period of 365 days (in this case, between the 2021-09-09 and 2022-09-09). This omits many users that initially came from the media coverage after the launch. However, it allows us to understand the statistics for real users that seek out the system in order to receive legal information, rather than visit out of curiosity after the media launch.

Overall, in this time period, the JusticeBot was accessed over 11,000 times. In total, users responded to 48,000 questions, and a total of 77,000 individual pages were seen by users. In the total time period the JusticeBot has been available, (July 2021 to February 2023), it has been accessed over 17,000 times.

Unsurprisingly, 94% of the users came from Canada. More surprisingly, 66% of the users accessed the tool via a computer, compared to 31% on a phone and 3% on a tablet. I would have expected more users to access the tool from their phone.

<<https://web.archive.org/web/20210917130831/https://www.cyberjustice.ca/2021/09/16/justicebot-repond-maintenant-a-vos-questions-sur-les-renovictions/>>.

8.4.1.2 User source

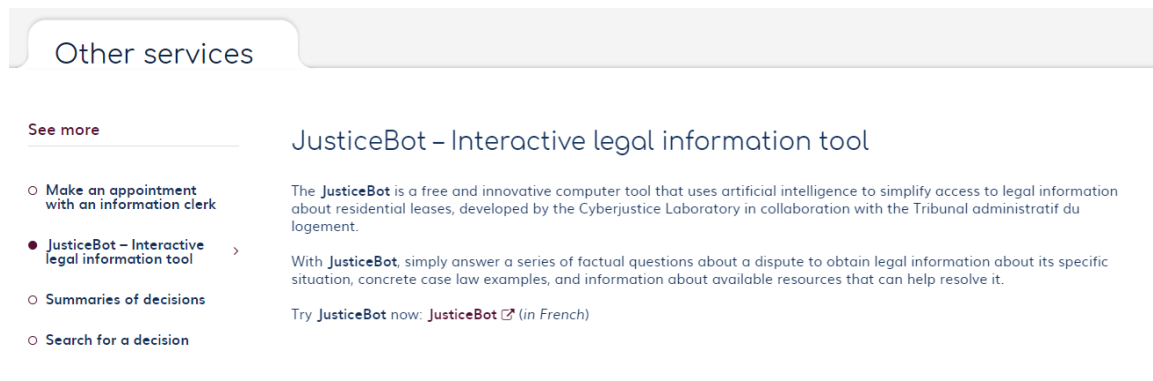


Figure 51 - JusticeBot as featured on the website of the Tribunal Administratif du Logement¹²⁷⁴

The TAL links to the JusticeBot TAL on their website, in the section listing different ways to obtain information online. 80% of the total users to the JusticeBot TAL follow this link to the JusticeBot. This indicates that the users are real users, that seek out the tool to obtain legal information. It also shows the importance of working together with a partner institution, as it creates a very natural access point where individuals go to obtain legal information about their issues. Other sources include users visiting the JusticeBot site directly (9.4%), via Google (3.3%) or via the website of the Cyberjustice laboratory (0.4%).

8.4.1.3 Outcome statistics

The analytics solution also allows us to track whether the user reaches the analysis screen, which gives them legal information, or ends up on the “missing question” screen, which informs them that their issue is not yet covered. Over the past year, 3,210 (32.4%) of the users went through an entire pathway to arrive at the “Analysis” screen. However, 6,714 (67.6%) of the users instead ended up on the screen that tells them that their issue is not covered.

This is surprisingly high. Of course, if the 32% of the users that do end up receiving information can be helped, this is an important and significant contribution to access to

¹²⁷⁴ note 1266.

justice. For the users that do not find help, the time investment of trying the JusticeBot is very small.

There could be multiple reasons for users reaching the “missing question” screen rather than the analysis screen:

- The user is unable to correctly identify the path that corresponds to their issue. Perhaps they do not see the correspondence between the title in the portal screen and their own situation. A possible way of addressing this issue is discussed below in 9.4.
- The issue of the user is not covered by the JusticeBot pathway. The number and variance of possible legal issues that could affect individuals, even in the relatively constrained space of rental disputes, is enormous. The JusticeBot TAL was developed by focusing on the most frequent issues, not perfect coverage of the entire domain. The data submitted by the users whose issue is not yet covered can be very helpful in determining which area to focus on next as we will see below in 8.4.4.

At the same time, being able to inform the user that their issue is not yet covered is a crucial feature of the JusticeBot TAL. Only so can it be ensured that people do not obtain irrelevant or erroneous information. Even in these cases, after the “missing question” screen, the user receives information on possible next steps, such as references to other sources of relevant information and the relevant institutions and tribunals. This way, even users that arrive at a dead end may benefit from using the JusticeBot.

8.4.1.4 Time spent

Google Analytics also allows us to see the time the users spend in the JusticeBot. This data is available from 2022-07-11 to 2022-09-10. On average, users are interacting with the system for 1m 35s. This may seem low. However, segmenting the users by outcome shows us a different picture, see Figure 52.

Totals		1m 35s Avg 0%
1	Get to Analysis	3m 37s
2	Missing question	2m 20s
3	Did not finish	0m 30s

Figure 52 - Time spent in JusticeBot depending on outcome

As we can see, the users that get to the analysis screen spend an average of 3m 37s on the JusticeBot. This is encouraging, as it shows that people interact attentively with the system. At the same time, it indicates that people are able to obtain the information in the system within a few minutes, which is also promising.

The users that reach the missing question screen spend on average 2m 20s on the screen, while users that do not arrive at either at 0m 30s. This is to be expected, since these users will have less content to peruse, or may simply be curious about the JusticeBot and leave after answering a few questions.

8.4.1.5 Pathway statistics

The JusticeBot sends anonymized events to Google Analytics whenever a certain answer to a question is chosen. This can allow us to understand the most popular pathways that users choose to explore. This data is available for the past 93 days, i.e. from 2022-07-10 to 2022-09-09.

The first question a new JusticeBot user is asked when entering the system is whether they are a landlord or a tenant. Over the analyzed time, 1,312 answers to this question were tracked. Out of these, 895 users (68%) indicated that they were tenants, whereas 409 users (31%) indicated that they were landlords. While we expected that the majority of the users would be tenants, it is good to see that landlords use the tool as well.

Pathway	Number of times	Percentage
Other	580	55.4%

My landlord wishes to undertake work in my apartment/building.	110	10.5%
My landlord wishes to increase my rent.	85	8.1%
I would like to leave my apartment before the end of the lease.	77	7.4%
I would like to terminate my lease.	72	6.9%

Figure 53 - Top 5 most clicked pathways for tenants

Pathway	Number of times	Percentage
Other	212	55.4%
The tenant does not pay their rent.	79	10.5%
The tenant is late in paying their rent.	46	8.1%
The tenant wishes to leave their apartment before the end of the lease.	32	7.4%
The tenant has left the apartment before the end of their lease.	27	6.9%

Figure 54 - Top 5 most clicked pathways for landlords

This feature also allows us to explore the pathways selected by the users, when faced with the menu page for landlords or tenants (see 8.3.4). In total, the menu page for tenants was accessed 1,047 times, while the page for landlords was accessed 449 times. These numbers are slightly higher than how often a user indicated that they were a landlord or tenant, likely due to individuals returning to previous pages to explore further pathways.

Figure 53 shows the most clicked pathways for tenants, while Figure 54 shows the most clicked pathways for landlords. As expected from the previous section, the most frequent option is “Other”, indicating that the user did not find a pathway matching their issue. As discussed, reasons for this may include users not finding the pathway corresponding to their issue, and the enormous variety of specific situations that landlords and tenants may encounter.¹²⁷⁵

Below that, however, there is a significant variance in the pathways that are chosen by the users. For tenants, the most frequently explored pathway was that of major renovations on an apartment. This pathway was added to the system last, after analysis of the currently relevant issues, showing the importance of updating and maintaining the JusticeBot.

For the landlord, the most important pathway is the one regarding tenants not paying their rent. This matches the expectations both from the statistics of the TAL and the data analysis. The fact that the real-world usage corresponds to the statistics to some extent indicates that real users access and benefit from the system.

Now that we have seen the statistics automatically collected by the JusticeBot, let us analyze the information explicitly provided by users, such as through ratings and a survey.

8.4.2 Ratings

Each single page in the JusticeBot has a rating section, which allows the user to rate that page (see 7.7.2). If the response is 3 stars or below, the user can further choose to describe with words why they were unhappy with their experience.

Overall, we received 349 such ratings. Figure 55 shows the distribution of the ratings. While many of the ratings are positive, there is also a substantial portion of negative ratings. Of course, one has to be cognizant of potential biases in who chooses to rate the system – people who are dissatisfied might be more likely to rate a screen, skewing the

¹²⁷⁵ See 8.4.1.3.

statistics. Overall, it seems like a majority of the users (52.5%) find the system to be useful, rating it 4 or 5 stars.




Figure 55 - Star ratings of pages in the JusticeBot

24 users made use of the possibility to describe their issue with the service. The most common complaint seems to be that the system is unable to answer the question of a user, which matches the findings above. Some users also asked for an English version of the tool.

We have not yet investigated the correlation of ratings to certain questions. We will do this in the future to potentially identify questions that are confusing or not useful to users of the system.

8.4.3 Survey

At the end of the JusticeBot TAL, the user is given access to a survey that allows them to give detailed feedback on the functionality of the JusticeBot. Overall, 28 users responded to this survey. While this is less than the other feedback methods, it allows for the qualitative understanding of the user experience. The survey contains a large number of questions – I will here explore a few interesting results.

8.4.3.1 User source

The large majority of survey respondents (71.4%) discovered the JusticeBot through the site of the Tribunal Administratif du Logement. This confirms the information presented above under 8.4.1.2. Other sources were the website of the lab and legal information clinics.

8.4.3.2 Welcome screen

A majority of the users agreed that the information on the welcome screen (describing the functionality and goals of the JusticeBot) was clear (88.5%) and allowed them to understand the goals of the JusticeBot (92.6%).

8.4.3.3 Questions

People were also generally happy with the questions they were asked. Of the respondents to the survey, 85.7% agreed that the questions were clear, 85.7% agreed that the order of the questions was logical, and 85.7% of the respondents agreed that the explanatory texts helped them understand the question better.

8.4.3.4 Case law summaries

People had generally positive reactions to the case law summaries. While 67.9% of the users thought that the decisions were easy to understand, and 64.3% of users thought that the decisions helped answer the questions, only 39.3% of people thought that the decisions were relevant to their own situation.

8.4.3.5 Analysis screen

With regards to the analysis page presented at the end of a JusticeBot parkour, 67.9% of individuals believed the information to be clear, and 60.7% of individuals agreed that they received an overview over the different relevant legal elements linked to their situation.

8.4.3.6 Overall experience

Overall, 57.2% of individuals agreed that they received the information necessary to understand their situation from the JusticeBot. 53.6% of individuals agreed that they gained a good idea of the next steps they could undertake from the JusticeBot.

This data serves two purposes for the JusticeBot TAL.

First, it allows us to understand the type of issues that people are confronted with in real life. The JusticeBot aims to cover areas that individuals frequently face. Therefore, it is built upon an analysis of case law and statistics, to identify issues that affect a lot of individuals. Of course, the choice of which areas to cover rests upon the assumption that the statistics are accurate and correspond to the issues that people are often faced with. However, people in the real world may face different issues than issues that people generally go to court with. Since the court decisions are the only data source we have, this may give us a skewed perception of which areas are important for users.

The data collected through the “missing questions” form are a much better source of data for this. It stems from real users that interact with the system, and do not find answers to their questions. By analyzing the responses, finding common clusters, and preparing pathways that address these clusters of questions, one can be sure that more users can be helped by the JusticeBot. This way, the more frequently the platform is used, the better it becomes, as the creator of the system can more accurately determine the necessary next steps. Analysis of the missing questions data has allowed us to determine that e.g. repossessions are a frequently occurring topic. The creation of a pathway for this area is currently under way.

Second, the data allows us to discover instances where the user may have an issue that is covered by the JusticeBot TAL but is not able to find the appropriate pathway in the menu. The pathway titles used to describe the pathways covered by the JusticeBot TAL may not correspond to how a user would describe their case, making them miss the fact that their situation is already covered.

In this instance, they would arrive at the “missing questions” page and enter their situation into the form. We can then spot the situations that are entered into the form but already covered by the system and adjust the names of the pathways to more closely

correspond to how individuals tend to think about their issues. There could also be other solutions to help the user identify the relevant path, which I describe below in 9.4.

8.5 Conclusion

In this chapter, I have given an overview of the JusticeBot TAL. It focuses on landlord-tenant disputes in Quebec, which is a highly important area with a significant need for legal information.

To respond to this need, we developed the JusticeBot TAL based on the JusticeBot methodology. The project was realized in collaboration with the Tribunal Administratif du Logement, with help from Aide Juridique Montreal et Laval. It further received financing from the Ministère de l'Économie et de l'Innovation du Québec.

In order to build the system, we started by analyzing the statutes, case law and statistics to identify suitable pathways to build. These were then encoded in the JusticeCreator by a team of legal experts and enhanced by adding summaries of case law. The experience of building the tool was very positive – the legal experts were able to quickly understand how the system works and start contributing content. This validates the JusticeCreator suite, and the capability of non-programmers to use it to construct legal decision support tools.

The process resulted in a legal decision support tool with almost 300 pieces of content. It covers a number of important issues that can affect both landlords and tenants. The JusticeBot TAL was launched on the 20th of July 2021 on <https://justicebot.ca> to significant media attention and continues to be visited by hundreds of users per month.

Overall, these users seem very happy with the tool. Most users indicated that they received the information necessary to understand their situation, and that they knew which next steps may be available to them. Further, over 85% of the users indicated that they would recommend the tool to their friends. This is a very positive reflection on the work we have done in building the JusticeBot TAL.

One important takeaway from the feedback received to the system is that it is currently able to handle under half of the situations that users have in mind when accessing the system. While helping any users is an important contribution to access to justice, increasing the number of users that can be helped is an important step in making the system more useful.

In conclusion, it does seem like the JusticeBot TAL can have an impact on access to justice for the user. By using the platform, they are able to gain important insights into their situation, that can be helpful no matter which path they choose to pursue.

Understanding these rights may allow them to take their issue and rights directly to the other party, thereby being able to build a constructive solution together. If the user wishes to go to the Tribunal, the JusticeBot offers an easy introduction to understanding their potential rights, and the ways they can pursue these rights. The tool is easy to use, and gives specific information based on the user answering questions, rather than having to search for general information in the law or on static websites.

The tool also demonstrates the importance of collaborating with a partner institution. Working with the TAL to build the JusticeBot has allowed us to have access to the necessary data in the form of legal decision documents and the professional verification of the content. Further, the website of the TAL is by far the biggest source of traffic to the JusticeBot. Since individuals that face rental issues frequently access the website of the TAL, being linked on this site has allowed us to significantly expand the reach of the platform, hopefully being able to help some of these users and thus increasing access to justice.

Chapter 9 Other Areas & Future Work

Research Objective: Discussing other application areas and future work (1.2.2.8)

Research Topics:

- Which areas of public law, administrative procedures or legal disputes may benefit from implementing legal decision support tools using the methodology?
- How can the methodology be further improved and expanded?

9.1 Introduction

I have described the JusticeBot methodology. I have also given an overview of the JusticeBot TAL, the first legal decision support tool that was built and deployed using the methodology.

In this chapter, I will explore what might be next for the JusticeBot. First, I will explore some other target users that the JusticeBot may serve (9.2). Until now, we have focused on layperson parties to conflicts. However, using the JusticeBot methodology, JusticeBot decision support tools could also be useful for other stakeholders, such as the public, legal aid clinics, government employees etc.

Next, I will focus on other areas where building a JusticeBot may be useful (9.3). This will include areas of legal disputes, public law and administrative areas. While the list will not be exhaustive, I will provide an overview of some specific areas that may be viable targets.

Finally, I will discuss some future work on expanding and improving the JusticeBot methodology and tools (9.4). This will include the addition of natural language processing to make it easier for users to interact with the JusticeBot. Further, I will describe future work aiming to make the creation of new pathways more efficient. I will also discuss work that can expand the capabilities of the platform and the decision support tools created with the platform, including using the platform to generate documents or as part of an online dispute resolution system.

9.2 Other target users

9.2.1 Introduction

The JusticeBot TAL, described in Chapter 8, focuses on laypeople users, who are faced with a legal issue and wish to receive information about their legal situation and next steps. However, there is nothing in the platform or methodology itself that limits it to this the layperson user. Sometimes, the exact same version focused on the layperson user could also be used by other users, with other purposes. Otherwise, changing the contents and/or pathways may make it possible for the JusticeBot tool to also work well for other target users. Let us explore which users could benefit from a JusticeBot tool.

9.2.2 The public

Previously, we assumed that the user of the JusticeBot would be in a conflict situation. Of course, this does not have to be the case. The public in general could also be interested in using the system in order to understand the law, something which may be difficult to accomplish by reading the law itself. By navigating through the JusticeBot, and exploring the different pathways, a curious citizen can understand how the law applies to them, even if they are not yet in a situation where this is relevant. This understanding can allow them to adjust their behavior in order to be in line with the requirements set out in the law. Since predictability is a key feature of the rule of law, being able to understand how the law applies to individuals is crucial in any society ruled by law. JusticeBot-based utilities could support individuals in understanding how laws may apply to them in hypothetical situations, and thus increase the rule of law in society and access to justice.¹²⁷⁶

9.2.3 Legal aid clinics

Another target user group that might significantly benefit from the introduction of a JusticeBot tool is legal aid clinics. As we saw above in 8.2.5, legal aid clinics can play an important role in providing people with information regarding their situation. Often manned by students, these clinics allow people to drop-in and present their situation. The

¹²⁷⁶ Compare 3.4.

aid clinic workers will then review the legal rules and explain the general legal situation to the individual. Legal clinics can be a very powerful tool to increase access to justice - a study conducted in Canada in certain provinces in Canada showed that 90% of the people who visited a legal clinic indicated that they gained a better understanding of their legal situation.¹²⁷⁷

Adapting a JusticeBot based tool in aid clinics could be a powerful method to increase access to justice. For example, it could be used by the workers of the legal aid clinic to support drop-in clients. The worker could act as a translator, by opening the JusticeBot site and asking the client the relevant questions. The legal aid clinic worker is, of course, not limited to asking a yes/no question, but can ask follow-up questions to correctly understand the situation. Once they have sufficient information, they can select the appropriate answer in the JusticeBot, and ask the next question. At the end of the pathway, they could explain the information shown by the JusticeBot to the user, targeting the level of complexity to the level of understanding of the user.

The advantage of using a JusticeBot tool in this way is that the client does not have to assess the legal criteria themselves – rather, the legal aid clinic worker can ask the individual for their pertinent facts and make the assessment of whether their situation fulfills a legal criterion in the JusticeBot, with the help of the explanations and case law in a JusticeBot decision support tool. They could also adapt the level of explanation to the level of understanding of the user. The legal aid worker would benefit from the comprehensive research that has gone into creating the JusticeBot, relieving them from having to do the legal research from scratch and guaranteeing consistency between different workers. As new situations are encountered at the legal clinic, these could be added to the JusticeBot using the JusticeCreator by the workers, in order to help their colleagues support future clients.

¹²⁷⁷ *Legal Clinics in Canada: Exploring Service Delivery and Legal Outcomes Among Vulnerable Populations in the Context of COVID-19*, J2-542/2022E-PDF (Department of Justice Canada, 2021) Last Modified: 2022-06-22.

9.2.4 Government employees

Likewise, government institutions could benefit from having an internally accessible version of the JusticeBot. Government institutions often have a duty to respond to phone calls and inform citizens of the laws that they are responsible for. This can be an enormous number of calls. In 2020-21, the tribunal administratif du logement received over 900k calls,¹²⁷⁸ and treated over 140k written communications.¹²⁷⁹ The “Commission des normes, de l’équité, de la santé et de la sécurité du travail” (CNESST), which applies the law regarding worker protections, responded to 1.1 million calls in 2021.¹²⁸⁰

Of course, it might be tricky for employees of the institution to rapidly assess the situation of the caller and provide relevant information. Having access to an internal version of the JusticeBot, that can guide the employee to verified and relevant information, could potentially increase the efficiency of responding to calls, and the correctness of the advice that is given. Further, any changes to the law or jurisprudence can immediately be reflected in the JusticeCreator and are thus immediately visible at the relevant position to all government employees, instead of requiring the teaching of the changes to all employees.

9.2.5 Lawyers

A JusticeBot tool could also be relevant for lawyers. When entering a new legal area, they could rely on a JusticeBot tool to quickly get an overview, or an initial assessment, of new situations. However, even experienced lawyers may forget important requirements. Using a JusticeBot tool could allow them to verify their procedure and provide information and case law that may serve as a useful initial place for research.

In “A Life-Saving Checklist”, Atul Gawande describes the power of a checklist to improve outcomes in hospitals. As championed by Peter Pronovost, such checklists can help doctors and nurses remember and make explicit the many steps required to treat

¹²⁷⁸ Simard, *supra* note 993 at 50.

¹²⁷⁹ *Ibid* at 3.

¹²⁸⁰ *Rapport annuel de gestion 2021* (Commission des normes, de l’équité, de la santé et de la sécurité du travail) at 13.

patients in the intensive care unit.¹²⁸¹ Perhaps, the JusticeBot could serve a similar function for lawyers (or even doctors), providing an adaptable and responsive list detailing the important steps to deal with certain legal matters.

9.2.6 Judges

Finally, judges may benefit from the implementation of a JusticeBot tool. By navigating through the JusticeBot, and selecting whether the legal criteria apply to the case that they are presiding over, they can easily get a quick insight into the legal situation, without being forced to rule in a certain way. This could be useful to help the consistency of decision making, and make sure that no criteria are unintentionally left out. Such a system could be even more useful in conjunction with a module that generates draft document templates, as described below.¹²⁸²

9.2.7 Conclusion

As we can see, the application of the Justicebot is not limited to parties that are immediately faced with a legal dispute. It could also have beneficial effects for other stakeholders in the judicial system. Some of these may be able to benefit from the same version that is targeted at layperson parties, while other stakeholders may benefit more from specifically created versions. Nonetheless, the JusticeBot framework should have the required flexibility to be able to support many different target users and use-cases.

9.3 Other legal and administrative areas

Let us examine some areas where it may be fruitful to build further tools using the JusticeBot methodology. This list is non-exhaustive and aims to serve as an exemplification of the kinds of areas that may benefit from the creation of a JusticeBot tool. I will use examples from Canada to motivate why these areas may be useful targets. Of course, such tools could potentially be created in any jurisdiction.

¹²⁸¹ Atul Gawande, “A Life-Saving Checklist”, (2 December 2007), online: *The New Yorker* <<https://www.newyorker.com/magazine/2007/12/10/the-checklist>>.

¹²⁸² See 9.4.4.1.2.

Further, this section will mostly discuss the JusticeBot methodology as-is. Below, in 9.4, I will discuss ways of expanding the JusticeBot methodology to include additional aspects, through features such as natural language interaction, document generation and integration in an ODR platform. These features may be useful to enhance the JusticeBot versions discussed in this section as well.

9.3.1 High-volume, low-intensity dispute areas

9.3.1.1 Introduction

In this section, I present a few legal areas that could be well suited for the creation of further JusticeBot decision support tools.

9.3.1.2 Consumer issues

One of the most prevalent areas of disputes is that of Consumer issues. A study in 2009 asked 6,665 respondents whether they had experienced any legal issues in the past three years. Of these, 22% responded that they had experienced a consumer issue. The author estimated that 5.7m individuals in the population of Canada could face such issues.¹²⁸³ The biggest problems were repairs or renovations being unsatisfactory, or the merchant not honoring a warranty following large purchases.¹²⁸⁴ A study conducted in 2021 (with a different methodology and categories) found that 18% of individuals faced a legal issue over the past few years, 15% of which related to a large purchase or service.¹²⁸⁵

In Quebec, consumer complaints are dealt with by the Office de la protection du consommateur (OPC). In the year 2020 to 2021, this office received 125k cases submitted by the consumer.¹²⁸⁶ The OPC uses the ODR platform PARLe (developed at the

¹²⁸³ Currie, *supra* note 297 at 12.

¹²⁸⁴ *Ibid* at 14.

¹²⁸⁵ Savage & McDonald, *supra* note 303 at 5.

¹²⁸⁶ *Rapport annuel de gestion 2020-2021*, by François Paradis, Zotero (Office de la protection du consommateur) at 6.

Cyberjustice Laboratory) to solve conflicts. Using this platform, 75% of cases were solved. Further, the average delay was only 21.9 days.¹²⁸⁷

This certainly seems like an area of high volume. Further, consumer disputes generally seem to be areas of low intensity, relating to goods and services that are purchased by the general public, and often involve relatively small amounts of money.¹²⁸⁸ Despite this, the problems may have a significant impact on the individual – in the study conducted in 2009, 80.9% of individuals saw their problems as important to resolve, and 43.2% believed it to cause difficulty for everyday life.¹²⁸⁹

However, not everyone is able to resolve their issues to a satisfactory extent. In the 2009 study, for consumer problems, 58.7% decided to handle their problem on their own, but 42.3% of these individuals believe the outcome could have been improved with assistance.¹²⁹⁰ The 2021 study showed that only around 24% of issues related to a large purchase or service in Canada in the last three years was resolved.¹²⁹¹ More than half of respondents in the study who experienced a serious problem said that they did not understand the legal implications of the problem, and not knowing or understanding their rights was also a primary reason to not take any action.¹²⁹²

These statistics seem to indicate that there is a significant need for information in the area of consumer disputes. Building a JusticeBot-based tool in this area would thus be a worthwhile endeavour. Such a system could inform the user e.g. of their rights respecting warranties of purchased goods and services, what to do if a performed service is inadequate, what they can expect from certain types of good and how to file their claim with the OPC or equivalent. Likewise, a JusticeBot tool could serve as an initial entry point to an ODR platform, to support the user in identifying the correct procedure. Or, in

¹²⁸⁷ *Ibid* at 8.

¹²⁸⁸ Currie, *supra* note 297 at 38–39.

¹²⁸⁹ *Ibid* at 35–36.

¹²⁹⁰ *Ibid* at 58–59.

¹²⁹¹ Savage & McDonald, *supra* note 303 at 27.

¹²⁹² *Ibid* at 10.

the course of an ODR procedure, the JusticeBot could inform the user of possible outcomes of their consumer dispute, thereby supporting them in settling their dispute.¹²⁹³

9.3.1.3 Employment issues

Another common type of issue is employment issues. According to a study performed in 2009, 17.8% of respondents had these types of issues, an estimated 4.6m individuals in Canada.¹²⁹⁴ These issues commonly included wages being owed, issues with health and safety, unfair disciplinary actions, harassment or unfair dismissal.¹²⁹⁵ People under 30 and individuals with disabilities were especially affected by these issues.¹²⁹⁶ Employment issues had a big impact on individuals, with 88.7% believing the problem to be important to resolve, and 68.8% claiming that their daily life is affected by the issues.¹²⁹⁷ 30.5% of people with employment issues decided to do something about the issue themselves. Out of these, 52.2% believed that their outcome would have been better with more assistance.¹²⁹⁸

In Quebec, many employees are protected by the “Act respecting labour standards”.¹²⁹⁹ It sets out rules regarding issues such as minimum wage, termination of employment and harassment.¹³⁰⁰ Employees that have issues with these situations can file their case with the “Commission des normes, de l’équité, de la santé et de la sécurité du travail” (CNESST), which applies the act and informs the public.¹³⁰¹ In 2021, this institution received 135k claims work workplace accidents or sicknesses, and the employees responded to 1.1 million calls.¹³⁰² Together with the Cyberjustice Laboratory, the

¹²⁹³ See below under 9.4.4.2 for an in-depth discussion of the use of JusticeBot in ODR.

¹²⁹⁴ Currie, *supra* note 297 at 12.

¹²⁹⁵ *Ibid* at 15.

¹²⁹⁶ *Ibid* at 27.

¹²⁹⁷ *Ibid* at 35–36.

¹²⁹⁸ *Ibid* at 58–59.

¹²⁹⁹ *Act respecting labour standards*, CQLR c N-11.

¹³⁰⁰ “Workplace Protections in Quebec”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/workplace-protections-in-quebec/>>.

¹³⁰¹ *Ibid*.

¹³⁰² note 1280 at 13.

CNESST has launched a version of the PARLe ODR platform, which promotes the mediation between employers and non-unionized employees in Quebec.¹³⁰³

Clearly, this is an area with a high volume of issues, with hundreds of thousands of complaints per year just in Quebec. There is also a significant demand for information. Building a JusticeBot based tool in this area could thus be very helpful. Such a system could inform the user e.g. whether they are owed wages and how to be repaid, what to do in case of safety issues at their workplace, or their rights in cases of unfair disciplinary actions or dismissals. Likewise, the JusticeBot system could be integrated with, and enhance, an ODR platform in the domain, to enhance and extend the platform.¹³⁰⁴

9.3.1.4 Debt issues

Around 20.4% of the respondents in the 2009 survey reported having had a debt issue in the last three years, corresponding to an estimated 5.2m Canadians.¹³⁰⁵ These included disputes over bills, collecting money owed, dealing with a collection agency or unfairly being refused credit or bankruptcy.¹³⁰⁶ Debt issues were seen as quite serious by the affected individuals, with 86.1% arguing that the problems are important to resolve, and 49.9% of the affected respondents stating that they caused issues in daily life.¹³⁰⁷ What is more, debt issues frequently appeared together with other issues – 78.5% of individuals with at least 6 problems had issues with debt.¹³⁰⁸ As we can see, debt issues are prevalent in society, and are thus a high-volume area.

Around half of the debt issues that were not related to bankruptcy included a relatively low amount that would likely not include engaging counsel at normal rates.¹³⁰⁹ This

¹³⁰³ note 475.

¹³⁰⁴ See 9.4.4.2

¹³⁰⁵ Currie, *supra* note 297 at 12.

¹³⁰⁶ *Ibid* at 15.

¹³⁰⁷ *Ibid* at 35–36.

¹³⁰⁸ *Ibid* at 47.

¹³⁰⁹ *Ibid* at 39.

shows that those issues can be seen as low-intensity. 59.4% of individuals with debt issues dealt with these issues themselves.¹³¹⁰

Debt issues are both high-volume and low-intensity, making them another good target for a JusticeBot, which could inform the users of how to deal with money that is owed to them, what their rights are when dealing with a debt-collection agency,¹³¹¹ and how to challenge decisions to refuse credit.

9.3.1.5 Neighborhood issues

Neighbourhood issues (such as vandalism and property damage) were not reported as a separate category in the 2009 study. However, they were the most commonly reported serious problem in the 2021 study. These issues are described as including vandalism and property damage.¹³¹²

Some of these occurrences may be criminal offences, warranting a call to the police, while other issues may be against city bylaws. In these instances, such as noisy neighbors, there may be a specific municipality that can be contacted to deal with the issue.¹³¹³ A JusticeBot decision support tool could help the user understand e.g. whether their neighbor infringes upon any of their rights, and how to respond to this, including who to contact.

9.3.1.6 Divorce

In 2019, 56,937 divorces were granted in Canada.¹³¹⁴ Divorcing can be a complex and stressful endeavour, requiring the understanding of multiple legal consequences, such as family patrimony, union, parenthood, child support and former spousal support. If the spouses agree on a divorce, they are able to present a joint application for divorce at the

¹³¹⁰ *Ibid* at 58.

¹³¹¹ Compare “Dealing with a debt collector”, (5 August 2022), online: *Financial Consumer Agency of Canada* <<https://www.canada.ca/en/financial-consumer-agency/services/debt/collection-agency.html>>.

¹³¹² Savage & McDonald, *supra* note 303 at 5.

¹³¹³ “Understand Neighbour Law”, (October 2017), online: *Dial-A-Law* <<https://dialalaw.peopleslawschool.ca/neighbour-law/>>.

¹³¹⁴ *A fifty-year look at divorces in Canada, 1970 to 2020*, The Daily 11-001-X (Statistics Canada, 2022) Last Modified: 2022-03-09.

superior court, based on a draft agreement that is turned into a judgment.¹³¹⁵ Otherwise, the divorce is conducted through a trial.¹³¹⁶

The Ministère de la Justice Quebec, which handles divorce in Quebec, explicitly asks couples to inform themselves before drafting an application.¹³¹⁷ It thus seems like legal information is an important part of planning and executing a successful divorce.

Building a JusticeBot in such an area could provide this information to couples, by allowing them to understand the criteria and possible consequences of their situation. An extended version of the JusticeBot (see 9.4.4.1.3) may even be able to draft an agreement to help the parties divorce, based on a mutually agreed upon solution.

9.3.1.7 Gender-based violence and sexual assault

Gender-based violence is violence individuals face due to their gender, gender expression, gender identity or perceived gender.¹³¹⁸ It can include “sexual, physical, mental and economic harm inflicted in public or private”.¹³¹⁹ Such violence can target anyone. However, certain groups, such as women, people with sexual orientation other than heterosexual, and women belonging to Indigenous groups are more likely to be affected.¹³²⁰

¹³¹⁵ “Amicable divorce”, online: *Ministère de la Justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/amicable-divorce/>>; “Divorce by Agreement”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/divorce-by-agreement/>>.

¹³¹⁶ “How to apply for divorce?”, online: *Ministère de la Justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/divorce/how-to-apply-for-divorce/>>.

¹³¹⁷ “Filing an application”, online: *Ministère de la justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/amicable-divorce/application-for-divorce-based-on-a-draft-agreement/filing-an-application/>>.

¹³¹⁸ “What is gender-based violence?”, (31 March 2021), online: *Government of Canada* <<https://women-gender-equality.canada.ca/en/gender-based-violence-knowledge-centre/about-gender-based-violence.html#women>> Last Modified: 2022-02-07; Adam Cotter & Laura Savage, “Gender-based violence and unwanted sexual behaviour in Canada, 2018: Initial findings from the Survey of Safety in Public and Private Spaces” (2019) *Juristat: Canadian Centre for Justice Statistics* 1–49 at 4.

¹³¹⁹ “Gender-based Violence”, online: *UNHCR* <<https://www.unhcr.org/gender-based-violence.html>>.

¹³²⁰ Cotter & Savage, “Gender-based violence and unwanted sexual behaviour in Canada, 2018”, *supra* note 1318 at 16–17; note 1318.

There are different types of gender-based violence, such as intimate partner violence. In Canada, a woman is killed by her intimate partner every 6 days.¹³²¹ Sexual harassment in the workplace is another significant issues. 52% of women report having been the subject to sexual harassment.¹³²² One in three women report having been subject to unwanted sexual behavior in public spaces.¹³²³

Another significant type of gender-based violence is sexual assault. In a study conducted in 2018, 30% of women (representing 4.7m women) in Canada aged 15 or older reported having been sexually assaulted at least once. 8% of men (around 1.2m) reported having been sexually assaulted.¹³²⁴ High proportions of the women (96%) and men (78%) affected reported that they were emotionally impacted by the assault.¹³²⁵ However, these situations are severely underreported – only about 5% of women said that police found out about their most serious incident of sexual assault.¹³²⁶ Reasons for this reportedly included the internalization of shame, guilt and stigma, a perception that the victim will be blamed or not believed, or that it was not worth taking the time to report the assault.¹³²⁷

Gender-based violence has been recognized as a significant issue by the government in many areas. In Quebec, a committee was tasked with coming up with solutions to the problem, leading to a list of 190 recommendations,¹³²⁸ and the creation of a tribunal specifically targeting cases of sexual aggression and domestic violence.¹³²⁹

¹³²¹ “Gender Based Violence in Canada”, (1 June 2022), online: *Canadian Women’s Foundation* <<https://canadianwomen.org/the-facts/gender-based-violence/>>.

¹³²² “Sexual Assault And Harassment in Canada”, (22 November 2022), online: *Canadian Women’s Foundation* <<https://canadianwomen.org/the-facts/sexual-assault-harassment/>>.

¹³²³ Cotter & Savage, “Gender-based violence and unwanted sexual behaviour in Canada, 2018”, *supra* note 1318 at 6.

¹³²⁴ *Ibid* at 15.

¹³²⁵ *Ibid* at 21.

¹³²⁶ *Ibid* at 20.

¹³²⁷ *Ibid* at 20–21.

¹³²⁸ *Contre la violence sexuelle, la violence conjugale et Rebâtir la confiance*, Stratégie Gouvernementale Intégrée 2022-2027, Stratégie Gouvernementale Intégrée 2022-2027 (Gouvernement du Québec).

¹³²⁹ *Loi visant la création d’un tribunal spécialisé en matière de violence sexuelle et de violence conjugale*, LQ 2021, c 32 2021; “À propos du tribunal spécialisé”, (14 September 2022), online: *Gouvernement du Québec* <<https://www.quebec.ca/justice-et-etat-civil/systeme-judiciaire/processus-judiciaire/tribunal-specialise-violence-sexuelle-violence-conjugale/a-propos>>.

There have been a number of projects aiming to use technology, such as chatbots, to support survivors of gender-based violence.¹³³⁰ Likewise, there have been research projects exploring how NLP can be used to, e.g., detect a type of gender-based violence from a description in natural language.¹³³¹

As we can see, the issue of gender-based violence is significant, deeply affecting millions of individuals per year. This means that designing a JusticeBot in this area would carry with it specific challenges, such as hitting the right tone with regards to the survivors and preventing the reliving of traumatic experiences. Likely, the system would be less focused on assessing the situation of a user, and more focused on providing a low-threshold point of interaction for survivors, supporting survivors in articulating their experiences, directing survivors towards the appropriate resources, and documenting an event for potential future reporting.

9.3.1.8 Conclusion

As we can see, there are plenty of legal areas that correspond to the requirements of having a high volume and low intensity of cases, while also having a significant social need for information. The study conducted by Ab Currie in 2009 found that for the people that believe that assistance would have led to a beneficial outcome, 67.6% of individuals believed that public information would have helped, whereas 30.4% of individuals believed that someone explaining the law and helping with filling out forms

¹³³⁰ Jessica Galang, “Botler.AI’s new chatbot analyzed 300,000 court documents to help sexual harassment and assault survivors”, (6 December 2017), online: *BetaKit* <<https://betakit.com/botler-ais-new-chatbot-analyzed-300000-court-documents-to-help-sexual-harassment-and-assault-survivors/>>; Julie Thomas, “Agressions sexuelles : Lila, le chatbot qui aide les victimes”, (16 July 2019), online: *Revis* <<https://revisherault.org/agressions-sexuelles-lila-le-chatbot-qui-aide-les-victimes/>>; “Spot A.I. tool for documenting harassment and discrimination receives major updates”, (20 October 2020), online: *Yukon Human Rights Commission* <<https://yukonhumanrights.ca/2020/10/spot-a-i-tool-for-documenting-harassment-and-discrimination-receives-major-updates/>>.

¹³³¹ Tobias Bauer et al, “# MeTooMaastricht: Building a chatbot to assist survivors of sexual harassment” (2020) *Machine Learning and Knowledge Discovery in Databases: International Workshops of ECML PKDD 2019*, Würzburg, Germany, September 16–20, 2019, Proceedings, Part I 503–521.

would have helped.¹³³² This is very promising for the application of the JusticeBot, since this is exactly the kind of information this tool can provide.

9.3.2 Public law and administrative law

9.3.2.1 Introduction

Previously, we have discussed building JusticeBot utilities in areas that involve a dispute between individuals, i.e. private law. However, there may be an even bigger opportunity to build JusticeBot tools in areas of public law, that govern the relationship between individuals and the state. Public law includes constitutional law, administrative law, tax law, and criminal law.¹³³³ While most individuals have to go through disputes very rarely, we all need to interact with the government when we file our taxes, apply for a driving license or enter a country. Further, we are always expected to follow government rules regarding e.g. restrictions during a pandemic.

The use of artificial intelligence in the sector of public administration has been a growing topic.¹³³⁴ Just like the previously discussed legal areas, these administrative procedures are governed by rules. The rules have certain criteria that can lead to rights and obligations when they are fulfilled. Instead of being decided by a court or tribunal, these criteria are often applied by government employees. Just like judges, these employees need to assess often vague criteria, in order to decide whether a certain situation falls under a certain rule or not, and the consequences of this. In case of disputes, the decisions can often be appealed to a court.

As you can see, the elements required to build JusticeBot-based systems are present in administrative decisions. In Chapter 4, I described the steps of legal decision making,

¹³³² Currie, *supra* note 297 at 59.

¹³³³ Jonathan Law & Elizabeth A Martin, “public law” in *A Dictionary of Law*, 7th ed (Oxford University Press) container-title: A Dictionary of Law.

¹³³⁴ Sonia LeBel & Éric Caire, “Stratégie d’intégration de l’intelligence artificielle dans l’administration publique 2021-2026” (2021) Gouvernement du Québec; “L’implantation et l’utilisation des outils d’intelligence artificielle dans l’administration publique”, online: *Laboratoire de cyberjustice* <<https://www.cyberjustice.ca/2022/06/21/limplantation-et-lutilisation-des-outils-dintelligence-artificielle-dans-ladministration-publique/>>.

including assessing evidence to determine the facts that underlie an application, deciding whether these facts underlie a certain legal criterion, identifying the logical consequence of legal criteria being fulfilled, and deciding on an outcome. I believe that these exact steps will also be carried out by government employees when deciding, for example, whether to grant a liquor license, a driving license or whether a police officer should fine an individual. Therefore, the JusticeBot methodology should be able to model these kinds of decisions as well.

One difference between administrative procedures and court procedures may be the lack of case decision materials. While court decisions are usually publicly available, individual administrative decisions may be harder to find. This could make it harder to identify decisions for inclusion in the JusticeBot pathways. However, in practice, this may not be a problem. Government decisions have to be based on rules. These rules can be encoded into JusticeBot pathways. Further, administrative decisions can usually be appealed, at which point a court will render a decision that will be followed by the agency. This decision can be integrated into the JusticeBot. Government agencies are also likely to employ some sort of guidelines or training materials that inform the workers how they should proceed with judging individual situations.¹³³⁵ These should be excellent sources of information to integrate into the JusticeBot, in order to support individuals in understanding how the laws apply to them.

Building JusticeBot tools in these areas could have an enormously positive impact on individuals. People are likely faced with administrative decisions significantly more frequently than with legal disputes. However, just like with legal disputes, individuals may struggle to understand which rules apply to them, how they apply and how to proceed with their situations, such as by filling out a certain form. A JusticeBot-based tool could allow these individuals to get a better idea of the rules that apply to their factual situation (such as whether they are allowed to enter a country), or how they might achieve their goals within the scope of the administrative rules (such as building a pool in

¹³³⁵ See e.g. note 888.

their garden). Informing the individuals of the laws that apply to them in an easy-to-understand way through a JusticeBot would also make it significantly easier for individuals to follow the law.

In this section, I will briefly describe some administrative areas where it could be useful to build a JusticeBot. This list is not intended to be exhaustive.

9.3.2.2 Travel and Immigration

International travel is incredibly common. In 2019,¹³³⁶ over 22 million tourists visited Canada.¹³³⁷ Of course, when travelling to another country, it is important to be aware of the rules that govern entry into a foreign country. The entry requirements may depend upon factors such as the source country, the entry point in the destination country, immigration status in the destination country, and more recently vaccination status. Further, there may be administrative hurdles such as filling out forms, obtaining visas and restrictions on what can be brought into the country. Building a JusticeBot in such an area could make it easier for people to understand and follow these rules to ensure a smooth travel experience, especially as travel requirements rapidly change due to developments regarding e.g. the COVID-19 pandemic. The JusticeCreator would allow the government agencies to rapidly adjust the pathway to reflect the latest rules.

The procedure may be even more complex when it comes to more permanent immigration. In 2019 in Canada, 404k temporary work permits were issued, 75k individuals transitioned from temporary to permanent residents, and 341k permanent residents were admitted to Canada.¹³³⁸ Immigrating, whether as a student, worker or permanent resident, can involve a significant number of steps and rules that need to be

¹³³⁶ Here chosen due to a marked decrease in tourists and immigrants from 2020 onwards due to the COVID-19 pandemic.

¹³³⁷ *Travel between Canada and other countries, December 2019*, The Daily (Statistics Canada, 2020) Last Modified: 2020-02-21.

¹³³⁸ *2020 Annual Report to Parliament on Immigration*, by The Hon Marco EL Mendicino (Immigration, Refugees and Citizenship Canada, 2020) at 3.

followed. A JusticeBot based tool could potentially inform individuals of which paths are available to them and the steps required to immigrate into another country.

9.3.2.3 Licenses

The government issues an enormous number of licenses to individuals and companies. These may grant the right to perform actions such as fishing, hunting, driving a car, providing medical or legal services, engaging in commerce and providing food and alcohol.¹³³⁹ Millions of such licenses are issued every year. Each of these underlie some requirements, that need to be fulfilled in order to be eligible for the license. Tools based on the JusticeBot methodology could help individuals and companies understand whether they need a license to perform a certain action, and what the requirements are to obtain such a license.

9.3.2.4 Social Aid

Government issues also provide a large amount of aid to individuals in need. This often includes program providing social aid. There can be many such programs – an abridged list regarding income assistance in Canada lists 43 different programs.¹³⁴⁰ However, the complexity of the system may prevent individuals from benefitting from the aids. For example, studies conducted in 2018 found that there are many individuals who are likely eligible for a disability tax credit, but never claim it, potentially due to the complexity and barriers in the process. Further, the appeals process can be very complex.¹³⁴¹

Building a JusticeBot that allows the citizens to easily understand the programs that they might qualify for, and where to apply, could support individuals in gaining access to these vital services.

¹³³⁹ Compare “Government Licenses” in *West’s Encyclopedia of American Law*, 2d ed (The Gale Group, 2008).

¹³⁴⁰ “Income Assistance”, (24 August 2019), online: *Employment and Social Development Canada* <<https://www.canada.ca/en/employment-social-development/services/benefits/income-assistance.html>> Last Modified: 2018-08-24.

¹³⁴¹ *Canada’s Tax System: What’s so Wrong and Why it Matters* (Chartered Professional Accountants of Canada, 2018) at 35.

Many countries also offer some sort of legal aid, that supports low-income individuals in legal conflicts. In Quebec, around 250k such applications are processed each year.¹³⁴² However, the affected individuals may not understand the requirements to obtain this legal aid, or which documents they need to prove their eligibility. A JusticeBot that is able to guide the individuals to these resources and allow them to understand the requirements that need to be fulfilled, could potentially allow a larger percentage of individuals to discover and use these resources.

9.3.2.5 Pandemic restrictions

During 2020, 2021 and early 2022, a number of pandemic restrictions were in force in many countries. These governed what individuals were allowed to do, in order to stop the spread of the COVID-19 pandemic. Rules included areas such as how many visitors one was able to invite to one's home or meet outside, what kind of sport activities were permitted and with how many people, where masks needed to be worn, and which establishments were allowed to be open. The rules further depended on vaccinal status, age, the type of the establishment and the time of day. Breaches of the rules could lead to significant fines.

In situations like this, a JusticeBot-based tool could be an excellent utility to help individuals and businesses understand the rules. Instead of trying to understand their situation from reading long static explanations of the text, the tool could narrow down the situation of the user and what they wanted to do, and then inform them of whether this was currently allowed and the restrictions that would need to be observed. The JusticeCreator could be used to update the frequently changing rules. The JusticeBot could thus serve as an easy-to-use, always up-to-date and user-adaptive explanation of the restrictions that are currently in force.

¹³⁴² note 365.

9.3.2.6 Tax

Many people have to file their taxes. In Canada in between February and September 2022, over 30 million individual tax returns were filed.¹³⁴³ The process can be time-consuming and frustrating. The process of filing taxes depends on understanding different legal criteria and applying them to your own situation. There are commercial entities that offer the filing of taxes using online expert systems, such as TurboTax.¹³⁴⁴ Using the JusticeBot, it would be possible to build a system that asks individuals for their situation and informs them of possible tax implications. In a future generation, it may even be possible to generate the forms required for tax declaration directly in the system.¹³⁴⁵

9.3.2.7 Police

Finally, JusticeBot could be useful in interactions with the police. Police are governed by rules, telling them what they are able to do in different situations and how they should behave.¹³⁴⁶ Informing citizens of these rules through a JusticeBot could allow individuals to better understand how police are able to intervene, ensuring a smoother interaction between individuals and police. It could also help them understanding when an officer may have treated them in a way that is not allowed, and how to file a complaint. Further, a system could be built for the police itself, that allows them to assess whether certain interventions are legal in certain situations, such as entering a house or searching a vehicle.

9.3.2.8 Conclusion

As we can see, there are a number of important public law and administrative areas where JusticeBot utilities could be helpful. Of course, the listed examples are just scratching the surface of the areas that could potentially be improved by the creation of JusticeBot

¹³⁴³ “Individual income tax return statistics for the 2022 tax-filing season”, (24 January 2023), online: *Canada Revenue Agency* <<https://www.canada.ca/en/revenue-agency/corporate/about-canada-revenue-agency-cra/individual-income-tax-return-statistics.html>> Last Modified: 2022-09-13.

¹³⁴⁴ “TurboTax® Official Site: File Taxes Online, Tax Filing Made Easy”, online: <<https://turbotax.intuit.com/>>.

¹³⁴⁵ See 9.4.4.1.3.

¹³⁴⁶ “Police Ethics”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/police-ethics/>>.

utilities. Overall, there are millions of interactions between citizens and government agencies each year. Building JusticeBot systems to assist citizens in understanding how the rules governing these interactions apply to them, we could potentially have an even bigger impact on access to justice than supporting individuals in disputes.

9.4 Future Research

9.4.1 Introduction

Beyond building JusticeBot tools for other areas, there are also a number of important improvements that can be developed to improve the methodology and platform, or even expand what is possible using the platform.

This section will explain some of the possible future research relating to the JusticeBot methodology. Some of this research will focus on improving the usability and efficiency of the JusticeBot frontend (9.4.2). Other research will instead target the JusticeCreator, using natural language processing methods to make the creation of pathways more efficient (9.4.3). Finally, it is possible to use schemas created in the JusticeCreator to enable completely new functionalities, such as generating documents, structuring evidence or being used to enhance Online Dispute Resolution platforms (9.4.4).

One interesting aspect of these enhancements is that they are not linked to a single version of the JusticeBot. Since the JusticeBot methodology is domain-independent, most of the research suggested in this section would benefit every present and future JusticeBot version.

9.4.2 Evaluating and improving the end-user interaction with the JusticeBot

First, let us examine some improvements that could be made to improve the end-user experience of the JusticeBot. This could be done through a more in-depth evaluation of the system, or the use of natural language processing to guide the user towards appropriate pathways in the system.

9.4.2.1 Further evaluation of the JusticeBot user experience

Understanding the impact of legal decisions support tools is crucial in order to determine how and where they should be deployed, and how they should be designed in order to maximize the positive impact.

Of course, what should be considered a positive impact is in itself an important question regarding legal decision support tools. According to G elinas *et al*, remodeling the justice system often requires the balancing of multiple competing values, including values relating to the satisfaction of the user, and the integrity of the justice system.¹³⁴⁷ Once the positive impact has been established, it is also important to decide how this impact should be measured. This can also be a tricky problem. Pasca describes best practices for evaluation, including seven steps that can be followed to evaluate tools.¹³⁴⁸ In the Autonomy Through Cyberjustice Technologies and Artificial Intelligence (ACT) project, subproject 8 focuses on evaluating legal decision support tools.¹³⁴⁹

For the JusticeBot, we collect statistics and feedback that can tell us how the user interacts with the system. As described in 7.7, these can be collected from a variety of sources, including analytics, ratings of pages, surveys and a page asking the user to describe their missing question. In 8.4, I analyzed these statistics with regards to the JusticeBot TAL.

However, while informing people is a worthwhile end-goal in itself, this data cannot tell us the impact that receiving the information has on the user. For example, are they able to leverage the information to settle their cases? How many of the users end up going to court compared to people who have not used the JusticeBot? Is the experience of going to court or the outcome in court impacted by the information?

¹³⁴⁷ Fabien G elinas et al, *Foundations of Civil Justice: Toward a Value-Based Framework for Reform* (Cham: Springer International Publishing, 2015); Alexandra Pasca, "Fabien G elinas et al, Foundations of Civil Justice: Toward a Value-Based Framework for Reform, Cham, Springer, 2015" (2017) 30:1 rqi 149–152.

¹³⁴⁸ Alexandra Pasca, "Promising Evaluation Practices Guide: A Few Basic Tips" 16.

¹³⁴⁹ "Subproject 8 - Tools for Self-Represented Litigants", online: *Projet AJC | ACT Project* <<https://www.ajcact.org/en/organisation/chantier-8-outils-dautorepresentation-pour-plaideurs-citoyens/>>.

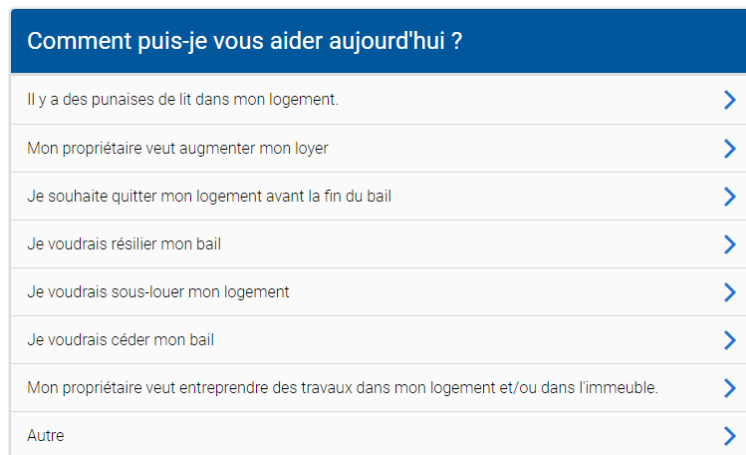
Collecting such information would be very beneficial to understand the true impact of the JusticeBot. However, it is also very complex, since people need to be followed over a longer time period and compared against people that have not used a JusticeBot. In future work, this would be a very interesting avenue to pursue, in order to understand the true impact of the JusticeBot on the users, and how the platform could be improved to be more helpful and beneficial.

9.4.2.2 Connecting lay language to pathways in the JusticeBot

Another promising avenue for future work is the inclusion of Natural Language Processing to link the lay description of a situation to legal pathways that are present in the JusticeBot.

9.4.2.2.1 The problem

When entering the JusticeBot tool, users may have difficulty identifying the legal pathway that applies to them. They are faced with a screen as can be seen in Figure 57, which shows them the list of possible paths that they can explore. The items on this page can be phrased in terms of either a situation (“There are bedbugs in my apartment”) or a goal (“I would like to terminate my lease”). It is further possible to link multiple items to the same pathway, in order to cover more situations.



Comment puis-je vous aider aujourd'hui ?	
Il y a des punaises de lit dans mon logement.	>
Mon propriétaire veut augmenter mon loyer	>
Je souhaite quitter mon logement avant la fin du bail	>
Je voudrais résilier mon bail	>
Je voudrais sous-louer mon logement	>
Je voudrais céder mon bail	>
Mon propriétaire veut entreprendre des travaux dans mon logement et/ou dans l'immeuble.	>
Autre	>

Figure 57 - A list of possible pathways for the user to explore

However, the user still has to decide which of the pathways applies to their situation. In 5.2.3, I described how users tend to think of their situation in terms of facts or goals.

Sometimes, however, their conceptualization of the situation may not correspond to the pathways. For example, an individual may face issues with heating, water leaks or noise. If these are intense enough, they may be sufficient to terminate their lease. Therefore, the user may want to click on the “I would like to terminate my lease” option, to obtain more information about this possibility.

However, the individual may not be aware of this consequence, since the specific issues are not on the list. Instead, they will click on “other”, and not receive the high-quality information that the JusticeBot can offer. In 8.4.1.3, we saw that a majority of users in the JusticeBot TAL end up on the “missing question” screen – users being unable to identify an existing pathway that applies to them may be a big part of this reason.

A possible solution to this issue is to extend the list with more possible situations, i.e. “I have issues with my heating” etc. However, this would make the list much longer, requiring the user to read and analyze many possible pathways to find the path that is right for them, which is not desirable.¹³⁵⁰

9.4.2.2.2 The solution

A possible solution to this issue is to present the user with a text prompt that allows them to enter their situation in plain language, upon which they would be presented with a list of suggested pathways.

¹³⁵⁰ Compare 5.2.6.

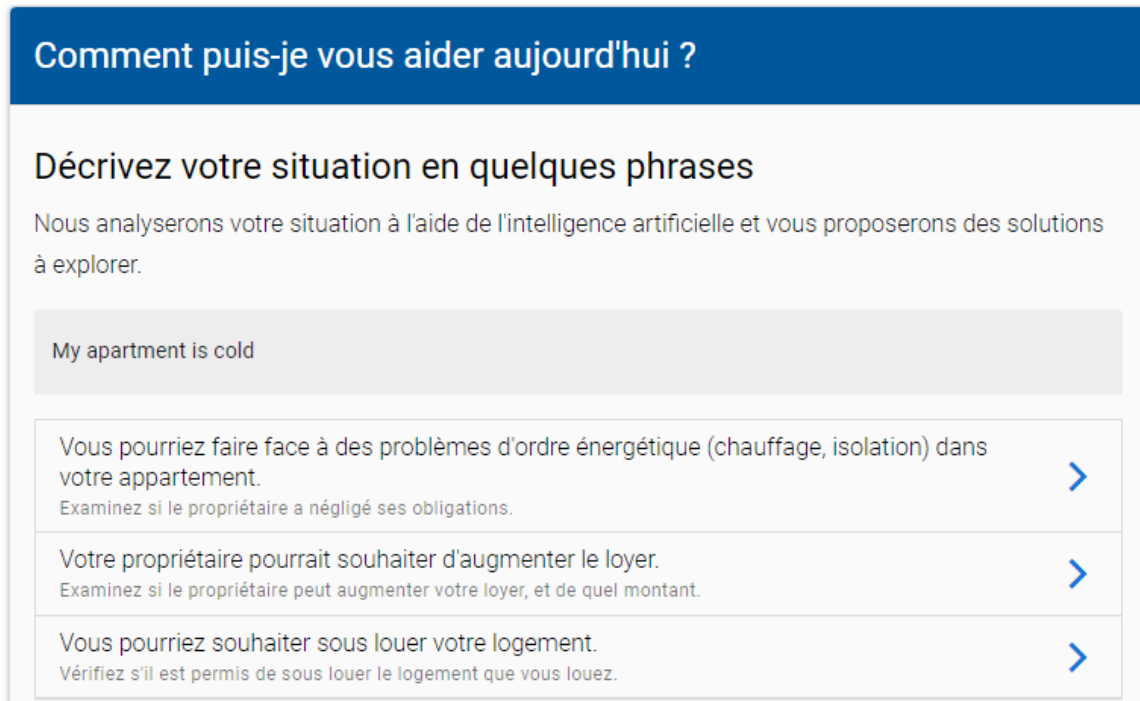


Figure 58 - Interface for lay language description of issue and path suggestion

Figure 58 shows how such an interface could look. The lay user is asked to describe their issue in their own words. After this, they are presented with a list of possible pathways they may be interested in exploring. In the example, the user has written: “My apartment is cold”. They are then given three suggestions, the first of which points them to a path exploring whether the landlord has failed in carrying out their obligations due to the heating issues.

Instead of the user having to determine which pathways fit their issue, the system here supports the user in taking this decision. At the same time, the pathways are just suggestions, and if a pathway seems irrelevant the user can instead pick from the list of pathways.

Developing such an interface could thus have the potential to guide users through the system more actively, by connecting their lay explanation of an issue to relevant pathways. Through this, it could increase the number of people that the JusticeBot is able to support and have a greater impact on access to justice. It is thus important research to

pursue. An initial exploration of the topic was presented at the ASAIL 2023 workshop.¹³⁵¹

9.4.2.3 Leveraging language models to support user interaction

Recently, the development of sophisticated large language models, such as ChatGPT and GPT-4,¹³⁵² have opened a new door when it comes to increasing access to justice. For example, one might consider directly asking legal questions to ChatGPT. However, one needs to be very careful in using this approach. Tan *et al* compared the information given by the JusticeBot to directly asking legal questions to ChatGPT, finding that the information given by ChatGPT is often inaccurate.¹³⁵³ These language models sometimes hallucinate, and e.g. make up facts or even previous cases.

One method of overcoming this issue is the use of “augmented” language models, where additional context (such as previous cases or legal information) is given to the model on each run. In this manner, the model can e.g. summarize or explain legal information that it has been given, which can reduce the risk of hallucinations, since the model already has direct access to correct legal information. This approach was shown to work well for e.g. explaining the meaning of statutory terms based on their use in previous cases.¹³⁵⁴

Combining the JusticeBot with these large language models thus seems like a promising approach. The JusticeBot contains a structured, verified representation of the legal questions that need to be answered to arrive at a legal decision. This structured representation could be combined with a language model to e.g. allow the user to interact with the system more naturally, while relying on the accurate content provided by the JusticeBot.

¹³⁵¹ Hannes Westermann et al, “Bridging the Gap: Mapping Layperson Narratives to Legal Issues with Language Models” (2023) 3441 Proceedings of the 6th Workshop on Automated Semantic Analysis of Information in Legal Text (CEUR Workshop Proceedings) 37–48.

¹³⁵² See 2.6.2.4.

¹³⁵³ Jinzhe Tan, Hannes Westermann & Karim Benyekhlef, “ChatGPT as an Artificial Lawyer?” (2023) 3435 Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice (CEUR Workshop Proceedings) , online: <<https://ceur-ws.org/Vol-3435/#short2>>.

¹³⁵⁴ Savelka et al, *supra* note 259.

9.4.3 Making the creation of JusticeBot pathways more efficient

In this section, I will look at ways to improve the efficiency of creating new pathways in the JusticeBot. The JusticeCreator offers legal experts an easy way to create and update pathways in the JusticeBot. Further, it integrates natural language processing to identify new relevant decisions more efficiently.¹³⁵⁵ However, understanding which pathways to build, and the act of building pathways can still take a long time and represent a significant bottleneck in the creation of new JusticeBot tools.

9.4.3.1 *Discovering the content of corpora of texts*

An important use of natural language processing techniques in the JusticeBot framework is to discover the content of corpora of texts. There are multiple ways that the methodology could benefit from having a quick and efficient way to discover the topics that appear in a text:

- One step in building JusticeBot tools is discovering frequent types of cases.¹³⁵⁶ This is important in order to build JusticeBot tools that are able to deal with the most common issues that face individuals in a certain legal area. For example, in the JusticeBot TAL, the most frequent type of cases was landlords aiming to evict tenants that stopped paying their rent.¹³⁵⁷ If the creator of the system has access to a corpus of documents, exploring this corpus could allow them to gather statistics on which issues or facts are frequent, and thus important to integrate into a pathway.
- Another important use for exploring the content of a corpus of texts is analyzing the feedback gathered in the JusticeBot tool. As we have seen in the JusticeBot TAL, users filled out thousands of responses to the question what their current

¹³⁵⁵ See 7.3.2.4.1.

¹³⁵⁶ See 7.3.3.3.

¹³⁵⁷ See 8.3.2.2.

issue was, if it was not covered by a pathway. This feedback can be a great way to discover which new issues should be added to a JusticeBot, or which pathway names should be clarified so that affected users can find them.¹³⁵⁸

As we can see, having a method to explore a corpus of texts in an efficient manner could allow us important insights that could better inform the creation of JusticeBot tools. Such a tool could work on the level of the entire decision, or on the level of individual sentences.

There are several commercial tools that could support such an analysis, including Nvivo¹³⁵⁹ and QDA miner.¹³⁶⁰ These tools are meant for the qualitative and quantitative analysis of texts, and as such offer methods to identify and classify components of texts. Applying them to the corpora described above could be a promising avenue in discovering the content of the corpora to enhance the JusticeBot.

Some of my own research also targets this area. In 2019, we presented a tool that allows a human to create search rules that are able to identify texts of a certain class. The tool showed the user statistical distributions of which words appear in texts. Assisted by this information, the user can decide that some of these words signify a certain class and should thus be included in a classifier. While the resulting classifiers were not as powerful as some traditional methods, they were significantly more compact and explainable.¹³⁶¹ In 2020, we presented a tool that supports the user in annotating corpora of documents, by offering a way to rapidly find similar sentences to a given sentence across documents.¹³⁶² Branting *et al* presented a method to annotate very few documents, and then project those across an entire corpus of documents, in order to perform scalable

¹³⁵⁸ See 8.4.4.

¹³⁵⁹ “NVIVO - Qualitative Data Analysis Software for Researchers”, online: <<https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software/home>>.

¹³⁶⁰ “Qualitative Data Analysis Software, Mixed Methods Research Tool”, online: *Provalis Research* <<https://provalisresearch.com/products/qualitative-data-analysis-software/>>.

¹³⁶¹ Westermann et al, *supra* note 194.

¹³⁶² Westermann et al, *supra* note 195.

and explainable legal predictions.¹³⁶³ Continuing these lines of research, one could imagine a tool that ingests a corpus of texts (whether cases or sentences) and allows the user to explore and classify, thereby learning about the content of the corpus.

There are also more automated ways that could be explored to build such a tool. For example, latent Dirichlet allocations are a method for automatically detecting a set of topics from textual documents, by exploring which words tend to co-occur in texts.¹³⁶⁴ Such methods have been used to model the topics of legislative texts,¹³⁶⁵ and to discover explainable factors for machine learning predictions of case outcomes.¹³⁶⁶

In 2022, Salaün *et al* used such topic models to explore the facts that appear in decisions where tenants sue landlords.¹³⁶⁷ The researchers used several clustering methods to discover factors that are relevant to such decisions, such as “water leakage”, “noise” and “bedbugs”, and explored whether the discovered topics matched the topics discovered in the course of the FactorBot research.¹³⁶⁸ While some of the topics ended up being nonsensical, many of the topics were relevant, and some were even discovered that were not discovered during the FactorBot research.¹³⁶⁹ Integrating this approach with the JusticeBot methodology could be very interesting for enhancing the productivity of the legal experts.

¹³⁶³ Branting *et al*, *supra* note 54.

¹³⁶⁴ Blei, Ng & Jordan, *supra* note 1097.

¹³⁶⁵ James O’Neill *et al*, “An Analysis of Topic Modelling for Legislative Texts” (2017) 2143 Proceedings of the Second Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 16th International Conference on Artificial Intelligence and Law (ICAIL 2017), London, UK, June 16, 2017 (CEUR Workshop Proceedings), online: <<http://ceur-ws.org/Vol-2143/paper7.pdf>>.

¹³⁶⁶ Aletras *et al*, “Predicting judicial decisions of the European Court of Human Rights”, *supra* note 56.

¹³⁶⁷ Salaün *et al*, “Why Do Tenants Sue Their Landlords?”, *supra* note 1098.

¹³⁶⁸ *Ibid* at 119.

¹³⁶⁹ *Ibid* at 120–121.

9.4.3.2 Linking legal criteria to paragraphs in legal decisions automatically

One key step in the JusticeBot methodology is reading legal decisions and integrating them into the pathway.¹³⁷⁰ These cases serve two purposes:

- They illustrate how judges tend to reason about certain legal criteria.
- They give information about the outcome of previous similar cases to the user.

Both of these purposes depend on being able to map the cases to the reasoning steps (i.e. applied criteria) by the judge. This can be a tedious process, requiring the reading and annotation of many individual cases by legal experts.

Therefore, a very interesting path of future work is attempting to automatically classify individual paragraphs in cases in terms of the specific legal criterion that is discussed in a paragraph, and whether the judge found the criterion to apply or not.

This could make the process of annotating case law more efficient. The legal expert could be shown suggestions for specific paragraphs that relate to a specific legal criterion, which would allow them to rapidly summarize the reasoning of the judge. In the absence of manual summaries, perhaps the automatically identified paragraphs could even be shown directly to the user, which could significantly increase the speed of building new JusticeBot tools.

Even more importantly, if the prediction is accurate enough, the identification of similar cases could be performed automatically. The system could classify the cases in terms of the legal criteria that the judge found to apply, and add an automatically extracted outcome to the pathway. Currently, this is done manually, but automating this step could increase the number of case outcomes that are shown to the user, giving the user a better statistical insight into the possible outcomes of their situation.

¹³⁷⁰ See 7.3.2.

In terms of research, bridging the gap between legal texts and symbolic models of legal reasoning has been seen as a potential revolution of the field of AI & Law, e.g. by Ashley.¹³⁷¹ This research may respond to this possibility, by automatically linking the paragraphs in the case text to the symbolic and explainable reasoning schema of the JusticeBot methodology. This research would also sit in a line of important prior work, as presented in 4.4.3.4.2.6. Exploring whether it is possible to link a JusticeBot schema to the text of decisions is thus a very promising avenue for future research.

9.4.3.3 Automatically summarizing decisions for inclusion in the JusticeBot

The JusticeBot relies on the summarization of case decisions. These summaries are targeted at specific legal criteria. The annotator tries to capture how the judge reasoned about a specific legal criterion, and the outcome they came to. The purpose of these summaries is to illustrate how judges tend to reason about certain criteria, so that they can understand how the criterion may be applied in their own case.¹³⁷²

Summarizing these decisions can be a major bottleneck for the creation of new JusticeBot systems. A promising avenue for future research is thus the investigation of whether this step can be automated. Even if not perfectly accurate, such a method could suggest a possible summary to the legal expert, allowing them to accept the summary or edit it to improve its usefulness. Hopefully, such a system could save the legal expert time, thereby making it more efficient to create JusticeBot systems.

Summarizations are a well-studied field in the domain of AI & Law.¹³⁷³ An investigation into this issue was presented by Salaün *et al.*¹³⁷⁴ The researchers used the data from the

¹³⁷¹ Ashley, *supra* note 44 at 3.

¹³⁷² See 7.3.2.2.4.

¹³⁷³ See e.g. M Saravanan & B Ravindran, “Identification of Rhetorical Roles for Segmentation and Summarization of a Legal Judgment” (2010) 18:1 *Artif Intell Law* 45–76; Linwu Zhong et al, “Automatic Summarization of Legal Decisions using Iterative Masking of Predictive Sentences” (2019) *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law* 163–172; Huihui Xu et al, “Using Argument Mining for Legal Text Summarization” (2020) *Legal Knowledge and Information Systems* 184–193; Huihui Xu, Jaromir Savelka & Kevin D Ashley, “Toward summarizing case decisions

JusticeBot TAL, and trained language models to perform the task of summarizing the decisions with regard to a specific legal criterion. While the resulting summaries were sometimes correct, a manual analysis revealed that they often contained factual errors.¹³⁷⁵ Extending this research is a promising avenue of future work in making legal experts more efficient in creating new JusticeBot systems.

9.4.4 Expanding the capabilities of the JusticeBot framework

The previous sections of future work focused on expanding the current functioning of the JusticeBot system, by improving the front-end user experience or the way JusticeBot based tools can be created.

In this section, I will instead look at ways that the capabilities of the framework can be expanded. This includes the utilization of the tool to generate documents, the integration of the tool into an Online Dispute Resolution platform, and the changing of the JusticeBot representation from a static to a flexible schema.

9.4.4.1 Using JusticeBot to generate documents

The JusticeBot methodology relies on a representation of the reasoning steps a judge performs to come to a decision on certain legal issues. By following the schema, and assessing whether the criteria apply in certain ways, it is possible to understand how the rules apply to a certain situation.

Currently, this capability is used to give the user information about their rights. An interesting avenue of further research would be to use the schema to generate documents, that could be used to formulate and structure the arguments and evidence by parties, or even serve as a template that can be used by the judge to more quickly produce decision documents. Let us explore how this could work.

via extracting argument issues, reasons, and conclusions” (2021) Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL ’21) 250–254.

¹³⁷⁴ Olivier Salaün et al, “Conditional Abstractive Summarization of Court Decisions for Laymen and Insights from Human Evaluation” (2022) Legal Knowledge and Information Systems 123–132.

¹³⁷⁵ *Ibid* at 127–131.

9.4.4.1.1 Generating letters or submissions for parties

In order to obtain relief from the court system, it is necessary to know which legal criteria apply to a situation, and which facts are relevant to establish the existence of these criteria. This information is required when submitting a claim to a court, in order to give the judge, the necessary information to render a decision. Further, it is required when writing a letter to the opposing party, in order to establish the existence of a claim. Such letters may lead to the settlement of a situation, as the other party realizes that the claim is likely to prevail in court. Further, such letters may have legal importance – generally, an opposing party needs to be notified and given the chance to respond before a case can proceed in court.¹³⁷⁶

However, writing such letters may be difficult for pro se litigants. The litigant may not be aware of the rights that they have,¹³⁷⁷ the criteria that need to be fulfilled,¹³⁷⁸ the facts that are relevant to establish the criteria,¹³⁷⁹ and the evidence that can support these facts.¹³⁸⁰ According to Branting *et al*, “pro se litigants seldom know what facts they need to establish or how to articulate and organize the facts in a manner that makes their claims amenable to evaluation.”¹³⁸¹ Susskind mentions the example of self-represented litigants arriving in court with a bag of un-indexed documents.¹³⁸² Landsman suggests that computer programs that help self-represented litigants preparing briefs and assembling proofs could be a useful way to provide help to such litigants.¹³⁸³

¹³⁷⁶ See e.g. 8.2.4.2.

¹³⁷⁷ See 4.2.

¹³⁷⁸ See 4.4.

¹³⁷⁹ See 4.3.

¹³⁸⁰ See 4.2.

¹³⁸¹ Branting *et al*, *supra* note 348 at 218.

¹³⁸² Susskind, *supra* note 22 at 122.

¹³⁸³ Landsman, *supra* note 338 at 456.

Frequent lateness of rent?

According to article 1971 of the Code Civil Quebec, terminating the lease requires the tenant to be frequently late with paying their rent. In this case, the tenant was frequently late.

Describe how your tenant was frequently late. How often were they late? How many weeks were they late?

Evidence

Payment receipts

Figure 59 - A mock-up of an interface used to structure evidence using the JusticeBot methodology

A system based on the JusticeBot could be very helpful in these instances. Each step in the JusticeBot schema is a legal criterion that has to be fulfilled or not in order to obtain a certain outcome. A special version of a JusticeBot tool could thus use the schema as a way to structure the argumentation of a user. At each step, the user would not just be asked whether a criterion applies, but also *why* they believe it does, and what *evidence* they have. The system could further give hints as to which facts and evidence might be relevant to each specific criterion. Figure 59 shows how such an interface could look.

The output of such a system could be a letter that presents a structured legal argument based upon the responses of the user. It would go through the relevant criteria, and describe the relevant facts and evidence. The letter would be both clear and legally correct. Depending on the output format, it could be used as a letter to the opposing party to inform them of the basis of the claim, leading to an efficient settlement, or a clear and structured court submission, that would make it easy for the judge to recognize the relevant facts and criteria. Previously, Walker *et al* and Satoh *et al* proposed systems that

could use expert systems to structure evidential reasoning.¹³⁸⁴ Here, we could use the JusticeBot schema to structure such reasoning.

9.4.4.1.2 Generating decision templates for judges

Another application of the generation of documents in the JusticeBot methodology is to support judges, by quickly generating a template for a decision based upon their own decision making.

After or during a hearing, judges would gain access to a JusticeBot version. For each legal criterion, they would be asked whether it applies or not in the case they are presiding over. The judge would answer these questions based on the case that they heard, which would take them further through the schema, until they arrive at an endpoint.¹³⁸⁵ At this point, they would be given a template for a possible decision, that is aligned with the answers they selected, and gives them space to explain their reasoning.

Let us take an example from the domain of landlord-tenant disputes to illustrate how this could work. We imagine that a judge decides that a tenant was frequently late with paying their rent, but that this did not cause the landlord a serious prejudice. After answering the questions in the JusticeBot, they could be given a document similar to the one shown in Table 10.

Table 10 - Example of decision template

<p>Facts</p> <p>[Describe facts]</p> <p>Frequent lateness of rent?</p> <p>According to article 1971 of the Code Civil Quebec, terminating the lease requires the tenant to be frequently late with paying their rent. In this case, the court finds that the tenant was frequently late in paying their rent, because [REASONING]</p>

¹³⁸⁴ See 4.3.3.2.

¹³⁸⁵ See 7.4.1.2.

Did the frequent lateness cause the landlord a serious prejudice?

Terminating the lease further requires this frequent lateness causing the landlord a serious prejudice. In this case, the court finds that this is not the case. [REASONING]

Outcome

For the aforementioned reasons, the court has decided to reject the claim of the landlord.

This provides the judge with a clear and pedagogical structure, that only requires them to fill in the reasoning sections. Such a template could make it more efficient to write decisions and serve as a clear and pedagogical start of a decision. At the same time, the power stays completely with the judge. They answer the questions in the JusticeBot that lead to the generation of the structure, and also fill in the reasoning sections to complete the decision. Further, the decision is only a guideline, and presented in an editable word document. If the judge wishes to interpret the legal criteria in a different manner, they are thus free to change the document.

Of course, it is important to assess the effects this would have on judges, to make sure that they are able to diverge from the template and are not unduly influenced. However, generating such documents is a very interesting avenue to explore for further research.

9.4.4.1.3 Generating contracts and other legal documents

While the JusticeBot methodology as described here is focused on emulating legal reasoning, the tooling around the framework could also be used to generate other types of documents. The JusticeCreator allows the encoding of a schema containing a number of questions (question blocks) and consequences to these questions (information blocks). Currently, the questions are designed to correspond to legal criteria, while the information blocks correspond to legal conclusions. However, one could also imagine a version of JusticeBot where the questions are more general questions, and the information blocks have other effects, such as creating a written contract or filling out a form.



Figure 60 - Mock-up of part of contract creator

Figure 60 shows a crude mock-up of how a pathway to create contracts might work. The first question asks what type of contract the user might want to create. After this, the questions try to understand the aims of the user. Each information block adds a certain clause to a document. At the end, the user can download and use the contract. Like this, individuals can create contracts that encode their needs in legal language encoded by experts, without needing to understand the law themselves.

A similar platform could be used beyond contracts to fill out forms, generate divorce agreements or prepare wills. The output is not even limited to written documents – potentially, the information blocks could correspond to lines of code, that are added to a smart contract¹³⁸⁶ that can then be deployed to the blockchain.¹³⁸⁷

Some of these functions are possible today, with the help of tools such as docassemble¹³⁸⁸ and A2J author¹³⁸⁹, that explicitly target filling out and assembling documents. Including this functionality in the JusticeCreator would leverage the easy-to-use interface to allow

¹³⁸⁶ “Introduction to smart contracts”, (23 January 2023), online: *ethereum.org* <<https://ethereum.org/en/smart-contracts/>>.

¹³⁸⁷ Hannes Westermann, “An Interactive Guide To The Blockchain, Part I-V”, (May 2020), online: *Cyberjustice Laboratory* <<https://www.cyberjustice.ca/mots-cles/blockchain/>>; compare Ilham Qasse, Shailesh Mishra & Mohammad Hamdaqa, “Chat2Code: Towards conversational concrete syntax for model specification and code generation, the case of smart contracts” (2021), online: <<http://arxiv.org/abs/2112.11101>> arXiv:2112.11101 [cs].

¹³⁸⁸ note 456.

¹³⁸⁹ note 455.

lawyers and institutions to create tools that could support individuals in writing legal documents, thereby increasing access to justice.

9.4.4.2 Integrating the JusticeBot with an ODR platform

Next, let us explore some ways to integrate a JusticeBot-based tool into a platform for online dispute resolution. As we have discussed in 3.5.2.3, online dispute resolution is a very effective way for individuals to resolve their disputes. ODR systems can contain multiple steps, including negotiation and mediation, that allow for the amicable settlement of a situation by allowing the parties to build a mutually acceptable solution.¹³⁹⁰ Even in cases where these options fail, the platform can allow the quick and efficient adjudication of an issue, without requiring the parties to wait for months for a court date, or physically be present at the court. Integrating AI into ODR has the potential to further increase their effectiveness. This has been referred to as ODRAI.¹³⁹¹

In this section, I will describe some ways the JusticeBot platform could be used to enhance an ODR platform.

9.4.4.2.1 Using JusticeBot as an entry point to an ODR platform

One way that the JusticeBot could serve to enhance an ODR platform is through serving as an entry point to such a platform.¹³⁹² In order to use an ODR platform, individuals may need to know where they should start. Which type of claim should they file, and what type of outcome can they ask for?

A JusticeBot based tool could serve as an initial orientation that allows the user to overcome these issues. Just like a normal JusticeBot, it would ask questions to understand their situation. If the user answers the questions in a way that would give rise to a claim, they could be given a button that says, “Open this claim in an ODR platform”. Clicking this button would transfer the user and their answers to an ODR platform. They would arrive at the page that allows them to file their specific issue. Further, the answers they

¹³⁹⁰ Benyekhlef et al, *supra* note 322 at 11.

¹³⁹¹ note 485 at 134.

¹³⁹² Compare Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 797.

selected in the JusticeBot would be transferred to the platform, allowing parts of the forms to be pre-filled and preventing the user from having to answer the same question twice. Previous such systems, such as the Civil Resolution Tribunal, were discussed above in 3.5.2.4.1.

This avenue of future work would thus use the JusticeBot as an entry point to an ODR platform, that combines the power of the JusticeBot in helping users understand their rights, with the power of an ODR platform to obtain effective, quick and hopefully amicable solutions to legal disputes.

9.4.4.2.2 Using JusticeBot as a tool to facilitate settlement in an ODR platform, by presenting a BATNA

Once already inside an ODR platform, the users are often given the option to settle their dispute. Settling at this stage is very positive. If the parties settle, it means that they have found a mutually satisfactory solution to their issue. Further, they were able to avoid the long and complicated court procedure. Since the settling of a dispute during the negotiation phase does not involve the support of a human mediator or adjudicator, it is further easier to scale such approaches

However, settling may be difficult if the parties have very different ideas of what is a fair outcome for their situation. For example, if the defendant is willing to offer 100\$, while the opposing party wants 10,000\$, it would be difficult to find a mutual agreement. Instead, the parties are likely to take their issue to court, which involves the issues of expenses, time and frustration.

Showing the outcome of previous cases to the users could be a way of making it easier for them to settle. If they can see that the previous cases in their area tend to lead to an outcome of between 300 and 400 CAD, it may be easier for the parties to find an agreement. This was previously discussed in 5.4.1.2.

A system based on the JusticeBot methodology could be used to provide this kind of information. JusticeBot based tools are able to give information about the previous outcomes of cases based on the criteria applied by a judge. This represents a BATNA¹³⁹³ to the situation of the user, since it shows the outcome of previous decisions that are similar to that of the user.¹³⁹⁴ Showing this information to the users of an ODR platform could thus be a powerful way of aligning their expectations and allowing them to settle their dispute.

This functionality would be even easier to implement if the user has entered the system via a JusticeBot-based entry point, as described in 9.4.4.2.1. In this case, the information regarding how the user believes that the criteria would be assessed are already captured in the system. The similar cases that are shown to the user in the JusticeBot methodology could then be shown to the user as a BATNA, without any additional input from the user.

9.4.4.2.3 Using the JusticeBot schema as the structure for an ODR platform

Finally, the JusticeBot-based schema could be used as the core of an ODR platform. In current ODR platforms, in order for individuals to settle their dispute during a negotiation stage, they may be given an interface that allows them to communicate via chat and make offers to settle with the other party. However, this chat is typically freeform. Since the lay users may not know which arguments are relevant, they may talk past each other and bring up unrelated issues, which could prevent the parties from settling. Further, if the case moves on to mediation or adjudication, the human in the loop may have to waste precious time to establish which relevant facts are agreed upon by both parties, and which issues are contentious.

Potentially, the JusticeBot schema could be used to provide a structure for the communication inside an ODR platform. This would start with the user entering their claim via a system such as described in 9.4.4.2.1. For each criterion, the system could ask the user to provide the factual basis and evidence and upload it into the system. Once the

¹³⁹³ Best alternative to a negotiated agreement, see 3.5.2.2.

¹³⁹⁴ Benyekhlef & Zhu, “Intelligence artificielle et justice”, *supra* note 305 at 796–797.

defendant accesses the system, they would see the structured arguments and evidence prepared by the plaintiff, and could accept or challenge individual criteria. Figure 61 shows how such an interface may look. This way, it would become very clear which of the criteria are disputed, and what the evidence is on either side.

Frequent lateness of rent?

Your landlords argues that you were frequently late.

Over the past 8 months, the tenant was late 7 times. The lateness varied between 1 and 3 months.

Evidence

8 x Payment receipts [See files](#)

[Admit](#) [Challenge](#) [Counter](#)

Figure 61 - A mock-up of an ODR interface following the structure of a JusticeBot schema

In some instances, either of the parties may decide that they agree with the other party and drop the claim or settle. In other cases, if the claim goes to adjudication inside the ODR platform, the adjudicator will receive a structured overview over the contentious issues, with the relevant evidence and discussion prepared by the parties. This would allow them to render a decision more efficiently.

Such a system would share similarities with the PROLEG system developed by Ken Satoh, which allows individuals to communicate by choosing desired outcomes and counterarguments, and also presents the judge with issues that need to be assessed for evidence.¹³⁹⁵ It would also be similar to a system presented by Tianyu Yuan, which allowed parties to challenge individual legal criteria in an online platform, to allow for

¹³⁹⁵ See 4.3.3.2.

the collaborative creation of decisions,¹³⁹⁶ and the Rechtwijzer 2.0 system, presented above in 3.5.2.4.2. Building such a system on the basis of the JusticeCreator platform would benefit from the ease of encoding new pathways. It is thus a promising avenue for further research.

9.4.4.3 A flexible schema?

Currently, the JusticeBot assumes that judges follow the same path of legal reasoning for solving cases in each situation. This often occurs in areas of high-volume, low-intensity disputes. However, it might fail in more complex situations, where the questions asked for reasoning are very dependent upon which arguments are raised by the parties, or the particular situation of a case. In this case, a static schema may not be sufficient to capture the individual situations treated in each case.

An interesting avenue of future research could involve switching from a static schema representation to a dynamic one. Instead of encoding a schema for each legal area, such a system would work by building an individual schema for each legal decision. This would allow the system to adapt to different paths of legal reasoning taken by the judge. All of these schemas could then be programmatically combined to create an overall schema, that is able to treat each unique situation, and handle all possible seen previous variants of an argumentation flow. Such a system would share similarities with PROLEG, which allows for the inclusion of “open” legal criteria, that are only included in the reasoning pathway when raised by either party.¹³⁹⁷ Likewise, the VJAP system is able to generate a graph structure containing all of the possible arguments with regards to a certain issue.¹³⁹⁸

A flexible schema could allow the capturing of more complex legal decisions, and more easily linking decision texts to the schema. It is thus a very promising area for future work.

¹³⁹⁶ Tianyu Yuan, *Structured relation – Increasing decision efficiency through collaborative decision creation* (Groningen, Netherlands, 2018).

¹³⁹⁷ Satoh et al, “PROLEG”, *supra* note 48.

¹³⁹⁸ Grabmair, *supra* note 52 at 46–68; Ashley, *supra* note 44 at 156–158.

9.4.5 Conclusion

In this section, I have presented a number of enhancements of the JusticeBot methodology and platform. These could make it easier to interact with the system, make it more efficient to build JusticeBot pathways, or even add completely new use-cases to the JusticeBot. Further, they could benefit any current or future JusticeBot, no matter the legal area.

Of course, before deploying these improvements, it is important to investigate the impact they might have on the user, and how they relate to the prohibition against giving legal advice for non-bar members. These questions will also be explored in future work.

9.5 Conclusion

As we can see, the future for the JusticeBot methodology is bright. In this section, I have described many areas that could benefit from the implementation of a JusticeBot decision support tool. The methodology can be applied not just to support a layperson party to a dispute, but also to support stakeholders such as the public, legal aid clinics, government employees, lawyers and judges. The JusticeBot could be useful in numerous areas of high-volume, low-intensity disputes (including consumer issues, employment issues and debt issues). Beyond this, it could be used in administrative areas, to support individuals in understanding laws and rules that affect them in interacting with government agencies. As we can see, there are plenty of JusticeBot versions to be built.

Beyond this, I also described a wide variety of future research to be performed on the platform itself, including the evaluation and enhancement of the user experience, the increase of the efficiency of the JusticeBot creation process, and the adding of new features to the platform. I am looking forward to investigating these avenues and integrating them into the JusticeBot platform.

In the next chapter, I will wrap up and conclude this thesis.

Chapter 10 Conclusion

Research Objective: Overall objective (1.2.1)

Research Topics:

- How can artificial intelligence be used to increase access to justice and access to legal information through the creation of a methodology for developing legal decision support tools?

10.1 Introduction

We have now arrived at the final chapter of this thesis. At the beginning, I set out to design a methodology that can be used to build legal decision support tools in order to increase access to justice and access to legal information. Here, I will summarize and discuss this methodology, and assess whether the objective has been achieved. First, let us take a very brief look at the different steps I took in answering the research question.

I started out by investigating the state of artificial intelligence. The capabilities of AI systems have evolved tremendously over the past few years, as the field went from expert systems to machine learning based approaches and deep learning. However, even today, machine learning models may lack some important prerequisites for legal reasoning, such as common sense and human values.¹³⁹⁹

Then, I explored the issues of access to legal information and access to justice. Overall, it seems like individuals have difficulties understanding their rights or resolving their disputes. This can cause issues for both the individual and society.¹⁴⁰⁰

Next, I explored legal reasoning, and previous research in automating legal reasoning. Legal decision makers carry out a number of steps in reaching a decision with regards to a situation, including finding a guiding legal norm, assessing evidence, facts and legal criteria, and eventually deciding on an outcome. There has been a significant amount of

¹³⁹⁹ See Chapter 2.

¹⁴⁰⁰ See Chapter 3.

research in performing such reasoning automatically. However, some of the steps may require AI that has common sense, which is beyond the current state of the art.¹⁴⁰¹

Then, I discussed a few design criteria that are important for designing a methodology to build legal decision tools that can increase access to legal information and justice. I decided to build a system that targets laypeople, focuses on areas of high-volume, low-intensity legal problems, gives specific and useful information, and is practical, i.e. can be used to build real-world tools.¹⁴⁰²

The first attempt at building such a methodology resulted in the FactorBot, which represented previous cases in terms of facts that occurred in them, and the outcome of the case. Thus, users could enter their facts, and receive information about possible outcomes. However, this methodology was not able to accurately predict the outcome of new cases, and the annotation of the cases was both difficult and time-consuming.¹⁴⁰³

Therefore, I designed the JusticeBot methodology. Here, legal decisions are represented by the reasoning path that a judge took in coming to a decision. A user is then asked to answer a number of questions regarding their situation, forming a sort of hypothesis about how judges would reason about their case, which can then be used to provide them with information and outcomes of previous similar cases. Instead of predicting the case of the user, this methodology supports the user in understanding their own situation. I also developed the JusticeCreator, a tool to design such legal decision support tools in a visual, drag-and-drop manner.¹⁴⁰⁴

This methodology was used to create the JusticeBot TAL, which focuses on rental disputes and was developed together with the Tribunal Administratif du Logement du Quebec. It was launched to the public in the summer of 2021 and has since then been

¹⁴⁰¹ See Chapter 4.

¹⁴⁰² See Chapter 5.

¹⁴⁰³ See Chapter 6.

¹⁴⁰⁴ See Chapter 7.

used by over 17k users. 86% of users who responded to a survey indicated that they would recommend the tool to a friend.¹⁴⁰⁵

However, the research I have done on the JusticeBot methodology is just a start. I discussed administrative and legal areas where further JusticeBot tools could be implemented, and ways to expand and improve the methodology, for example by integrating more artificial intelligence techniques and by connecting a JusticeBot to an online dispute resolution platform.¹⁴⁰⁶

Here, I will briefly recapitulate the main contribution of this thesis, namely the JusticeBot methodology (10.2) and the implemented toolchain (10.3). I will then discuss some aspects of the methodology, and what I believe makes it interesting (10.4), as well as some limitations (10.5). Finally, I will come back to the question of whether I have achieved the research objective of building a methodology that can be used to build tools that can increase access to legal information and access to justice (10.6), and the impact such tools may have on the legal profession (10.7).

10.2 The JusticeBot methodology

The JusticeBot methodology was presented and described in Chapter 7. It offers a novel method to encode legal information and cases in a hybrid system and expose this information to the layperson user to provide legal information. Since this methodology is one of the key contributions of this thesis, I will briefly summarize the methodology, and how it can be used to create legal decision support tools, here.

10.2.1 The encoding of legal rules and cases

The JusticeBot methodology relies on the encoding of rules and cases.

10.2.1.1 The legal reasoning schema

The rules that are encoded in the methodology aim to replicate the steps that judges, or other legal decision makers, in reality perform to come to a decision regarding specific

¹⁴⁰⁵ See Chapter 8.

¹⁴⁰⁶ See Chapter 9.

cases. The rules are thus extracted from reading the relevant laws and cases dealing with a specific issue. These rules are then encoded in the form of a rule-based reasoning schema, which contains two main types of blocks: Question blocks (which correspond to legal criteria applied by the judge), and information blocks (which correspond to legal conclusions made by judges). By encoding these rules in the system, a schema of the possible reasoning paths taken by a judge is created. Each question and information block has simplified, plain language explanations as to their significance.¹⁴⁰⁷

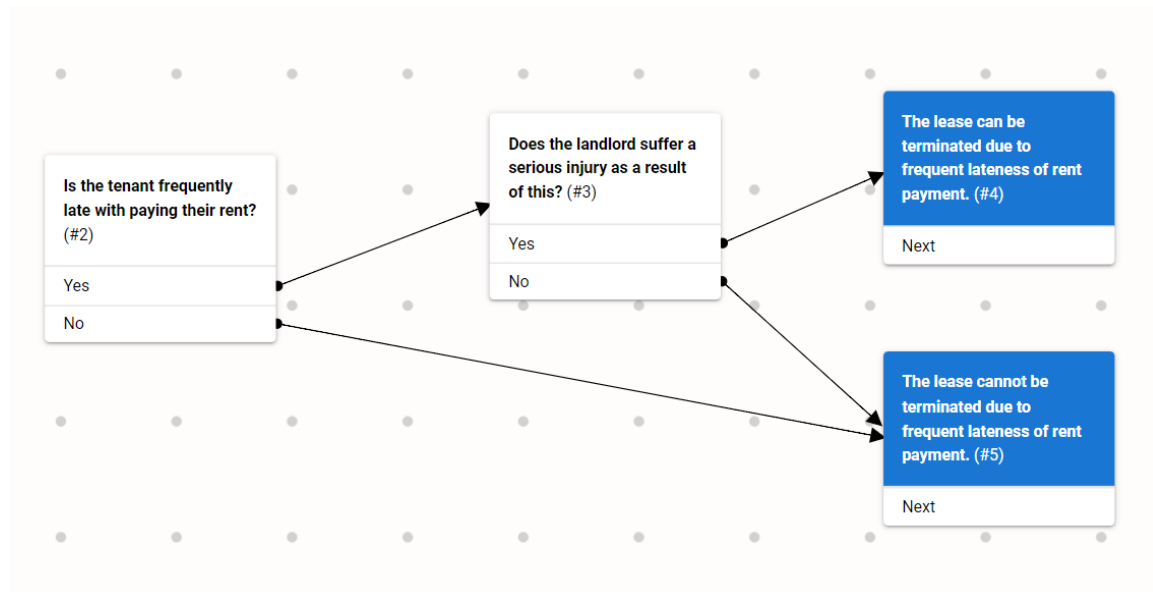


Figure 62 - reasoning schema for lease termination due to frequent lateness of rent payment

Figure 62 shows such an encoded reasoning schema for determining whether a lease can be terminated due to frequent lateness of rent. By reading cases and legislation, it was determined that judges tend to apply the cumulative requirements of “Is the tenant frequently late with paying their rent?” and “Does the landlord suffer a serious prejudice due to this lateness?”. These criteria and the logical connection between them were then encoded in the schema, using the white question blocks connected by arrows. Further, it

¹⁴⁰⁷ See 7.3.1.

was determined that judges typically arrive at legal conclusions (e.g. the lease can be terminated or not), based on whether these criteria apply or not. These were encoded in blue information blocks.

10.2.1.2 Cases to illustrate legal criteria

After the rules have been encoded, cases are introduced to the reasoning schema. In order to do so, cases are read in parallel with the schema. For each case, the creator identifies the legal criteria that are being discussed by a judge in the case. For each criterion, the creator then analyzes *why* the judge finds that a certain legal criterion applies or not. This reasoning is summarized and added as an illustration to the relevant question block.¹⁴⁰⁸

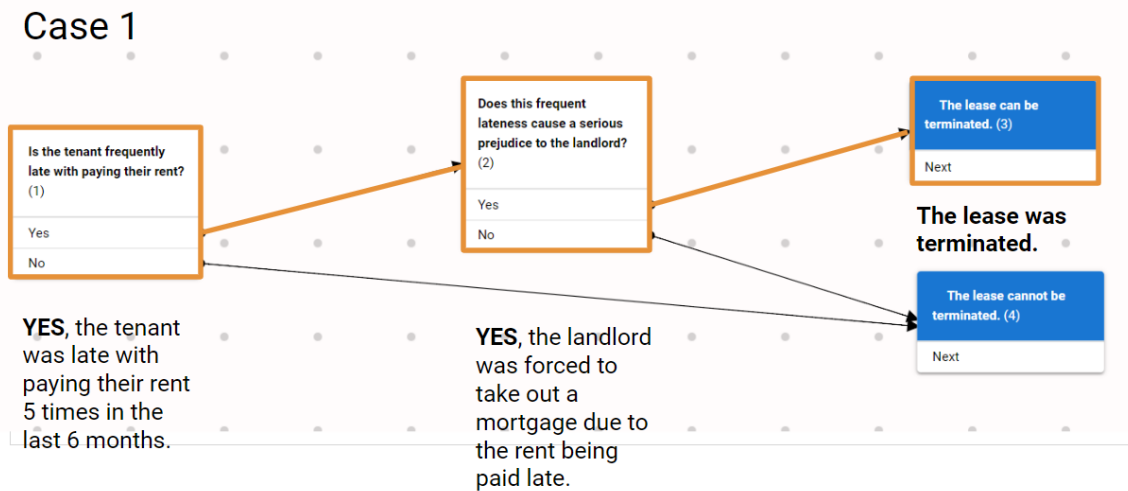


Figure 63 - A case is annotated and added to the reasoning schema

For example, to annotate a case in connection with the reasoning schema presented in Figure 62, the legal expert would read a case about frequent lateness of rent payment in conjunction with the schema. Figure 63 shows an example of how such a case could be introduced into the schema. The judge first assesses whether the tenant was frequently late with paying their rent. Since the tenant was late 5 times in the previous 6 months, the judge finds that the tenant was frequently late in paying their rent, i.e. the first criterion

¹⁴⁰⁸ See 7.3.2.2.

applies. Therefore, the legal expert adds a summary of this reasoning step, indicating that the judge found that the criterion applies, and summarizing the reasoning.

Since frequent lateness of payment did occur, the judge subsequently assesses whether this serious lateness caused the landlord a serious prejudice. This was also found to be the case. Once more, the legal expert annotates the fact that a serious prejudice was found, and why.

10.2.1.3 Cases to illustrate outcomes

By following the pathway in line of the reasoning of the judge, the creator of the decision support tool will eventually arrive at an information block, that corresponds to a legal conclusion taken by the judge. To this block, the creator attaches a summary of the outcome ordained by the judge in connection with that conclusion, e.g. that damages should be paid, and how much. In our example in Figure 63, the judge decided to terminate the lease. Therefore, the legal expert would add “the lease was terminated” to the corresponding information block.¹⁴⁰⁹

10.2.1.4 Cases to validate the schema

In reading the decisions in conjunction with the schema, the legal expert may notice that the schema does not correspond to the reasoning performed by judges in decisions. In this case, they may need to update the schema to reflect the real reasoning path of the judges, for example by adding new criteria or altering the logical structure of the schema. Thus, each further decision that is encoded empirically validates the correctness of the logical flow of the schema. Once the schema arrives at a state where most decisions correspond to the logical flow of the schema, the system likely accurately captures the different reasoning paths a judge may take to arrive at a conclusion regarding decisions in a certain legal area.¹⁴¹⁰

¹⁴⁰⁹ Compare 7.3.2.3.

¹⁴¹⁰ See 7.3.2.1.

10.2.2 The capturing of a new user case

Once the legal information has been encoded into the system, it can be used to capture the information of a user. The system can be accessed by a user who wants to know their rights in a certain situation, or who wants to achieve a certain goal, and wants to know whether this is possible. At first, the user is presented with a menu of different pathways that the JusticeBot tool can treat.

After the user has selected the pathway that is relevant to them, they are guided through the schema created in the previous step. For each legal criterion, the user is asked whether they believe that the criterion would apply in their case. In order to support them in making this assessment, the user is given the summaries that explain how judges reasoned about that specific criterion in prior cases. Based on these summaries, the user can hypothesize that the judge will find a certain criterion to apply or not. Depending on their answer, they are guided to the next question, until they reach the end of a pathway. By traversing the schema encoded by the system, the user thus creates a hypothesis of how their case may be assessed by legal decision makers.¹⁴¹¹

¹⁴¹¹ See 7.4.

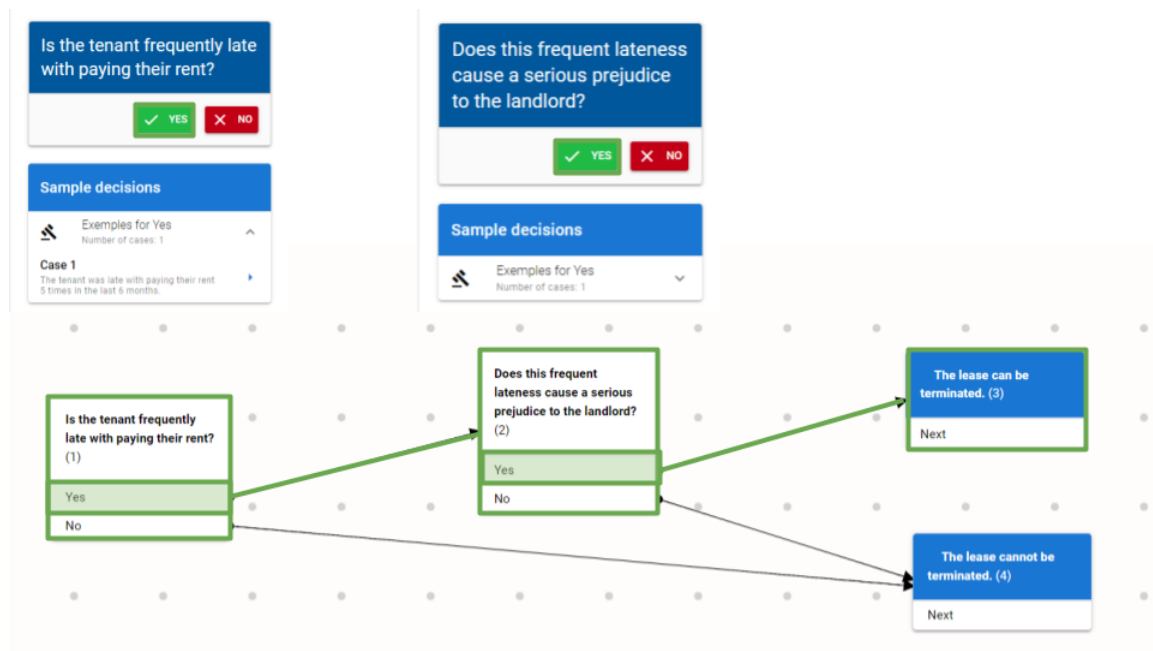


Figure 64 - A user navigates the JusticeBot interface to create a hypothesis for how their case will be treated

Figure 64 shows how a user would navigate our example schema, with the annotated cases. First, a user is asked whether in their case, the tenant can be seen to be frequently late with paying their rent. To assist them in this determination, they are given the previously frequently annotated case, showing that in “Case 1”, the judge found that the criterion “frequent lateness of rent” applies when the tenant was late in paying their rent 5 times in the 6 previous months. Based on this case, the user can estimate what a judge might find in their situation (in reality, there would be several cases, for “Yes” and “No”). In the example in Figure 64, they pick the “Yes” option. Next, they are asked whether this frequent lateness causes the landlord a serious prejudice, and are again shown the relevant case criterion summary. Again, they indicate that the answer is yes.

The system has thus captured the hypothesis of the user, indicating how they expect judges to reason about their situation. In Figure 64, this hypothesis is shown as a green line through the reasoning schema.

10.2.3 The analysis of the case of the user

The next step in the methodology is to analyze the hypothesis of the user, in order to find cases that are similar, and to provide them with legal information. The relevant information and cases are selected using the information blocks (corresponding to legal conclusions) that the user passed through in entering their case to select information from the schema.¹⁴¹²

As previously mentioned, each information block contains a simplified explanation as to the possible consequence of the judge arriving at a certain conclusion. This information can be shown to the user who passes through this information block in their hypothesis.

Further, each information block has a number of case outcome summaries attached to it. By selecting cases where the judge found the same legal criteria to apply, we can show the user of the system the outcomes that judges tend to award in cases like theirs.



Figure 65 - An example of a user hypothesis matching a previous case

Figure 65 shows an example of a case being retrieved based on a user hypothesis. The image corresponds to the example discussed above. The user has introduced the hypothesis that both criteria apply in their case. The logical consequence of this is that they arrive at the information block “The lease can be terminated”. We thus select the explanation from this block. Further, since the judge arrived at the same legal conclusion in the example case above in Figure 63, we can show the outcome summary to the user, to inform them of the outcomes that judges previously awarded in similar cases.

¹⁴¹² See 7.5.

10.2.4 Showing information to the user

Once the information and cases have been selected, they are shown to the user. The user is informed of the possible legal consequences that a judge may decide upon, based on the answers the user gave. Thus, the user can understand the potential rights they have, stemming from their situation.

The user is also provided with summaries of the outcomes ordered by the judge in previous cases, where the application of the legal criteria matches the hypothesis of the user. Thus, the user can get an insight into the real-world outcomes that cases such as theirs tend to result in, which can help them in their decision making.¹⁴¹³

Finally, the user is shown a list of possible next steps, which can help them understand the options that are available to them in resolving their situation.

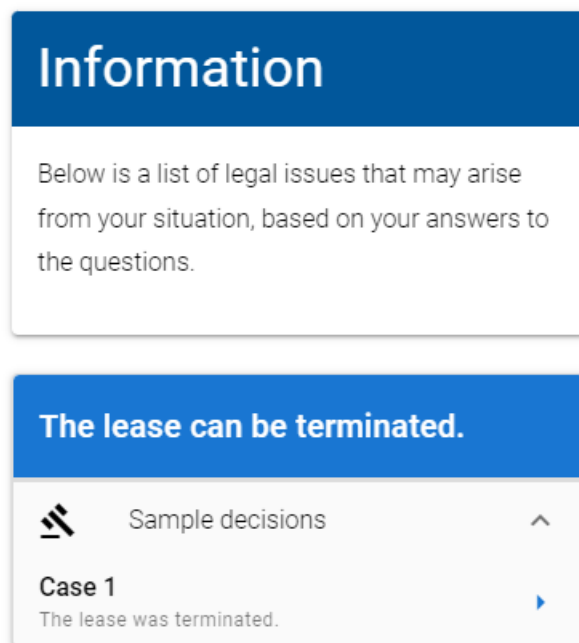


Figure 66 - The information a user is shown at the end of their circuit, including information and case law.

¹⁴¹³ See 7.6.

Figure 66 shows the information that a user selecting the answers above in Figure 64 would be provided in our example. The user is given the information that the lease can be terminated, should their hypothesis be correct. Further, they are shown a summary of a case example, where the lease was, in fact, terminated. The user thus gets an empirical understanding of the actual outcomes that were awarded by judges in situations similar to theirs.

Now that we have seen the core of the JusticeBot methodology, I will give an overview over the JusticeBot toolchain, which can be used to build such legal decision support tools.

10.3 The JusticeBot toolchain

My research goes beyond the conceptualization of a methodology to build legal decision support tools. I also practically implemented such a methodology in a software toolchain that allows the creation of JusticeBot tools in many legal domains.

The implemented system, here referred to as the JusticeBot toolchain, has two components: The JusticeCreator and the JusticeBot frontend. The JusticeCreator is used to encode the legal rules and cases, as described above. The resulting schema is exported in a JSON-file, which can then be consumed by the JusticeBot frontend. The JusticeBot frontend allows the user to interact with the system, by traversing the legal rules and receiving the resulting information and case law summaries.

10.3.1 The JusticeCreator

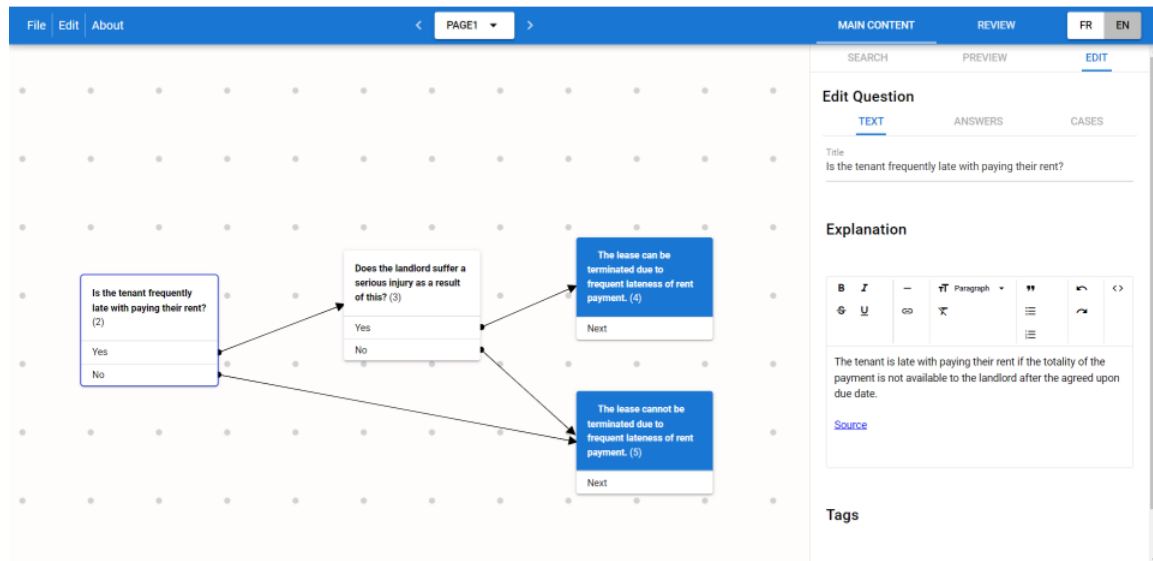


Figure 67 - A screenshot of the JusticeCreator interface

The JusticeCreator allows the creation of decision support tools. Figure 67 shows a screenshot of the JusticeCreator. The system is designed as a web-app, and thus allows the user to access it from their web browser. It presents the creator of such systems with a WYSIWYG (What You See Is What You Get) interface that allows the encoding of rules, by creating blocks, arranging them in visual patterns, and connecting them logically via arrows. The content of the blocks can also be edited very easily, by using a rich text editor interface.¹⁴¹⁴

Further, cases can be encoded in the same interface. For each question block, case summaries can be added to illustrate the application of criteria. For each information block, summaries of the outcomes of previous cases can be added.¹⁴¹⁵

The system has a number of features to make it easier to navigate and structure the pathways. For example, it offers the capability to split the pathway into multiple pages, to

¹⁴¹⁴ See 7.3.1.5.

¹⁴¹⁵ See 7.3.2.4.

preview the JusticeBot legal decision support tool and to search for questions in the pathway.¹⁴¹⁶ Further, it implements a machine learning feature that can search a database of previous case law to identify cases that are likely to correspond to the path the creator is currently working on, allowing them to rapidly identify cases suitable for annotation.¹⁴¹⁷

The JusticeCreator system is thus both efficient and friendly to users without a technical background. In fact, it is currently being used by a number of legal experts at the Cyberjustice Laboratory to create JusticeBot tools in a number of legal domains. After a brief introduction, users are quickly very comfortable in using the JusticeCreator to create, update and verify JusticeBot schemas.

Once the creation of a tool is done, it can be exported into a special JSON-file by the JusticeCreator. This file can then be integrated into the JusticeBot frontend, to be accessible to laypeople via the internet.

¹⁴¹⁶ See 7.3.1.5.1.

¹⁴¹⁷ See 7.3.2.4.1.

10.3.2 The JusticeBot frontend

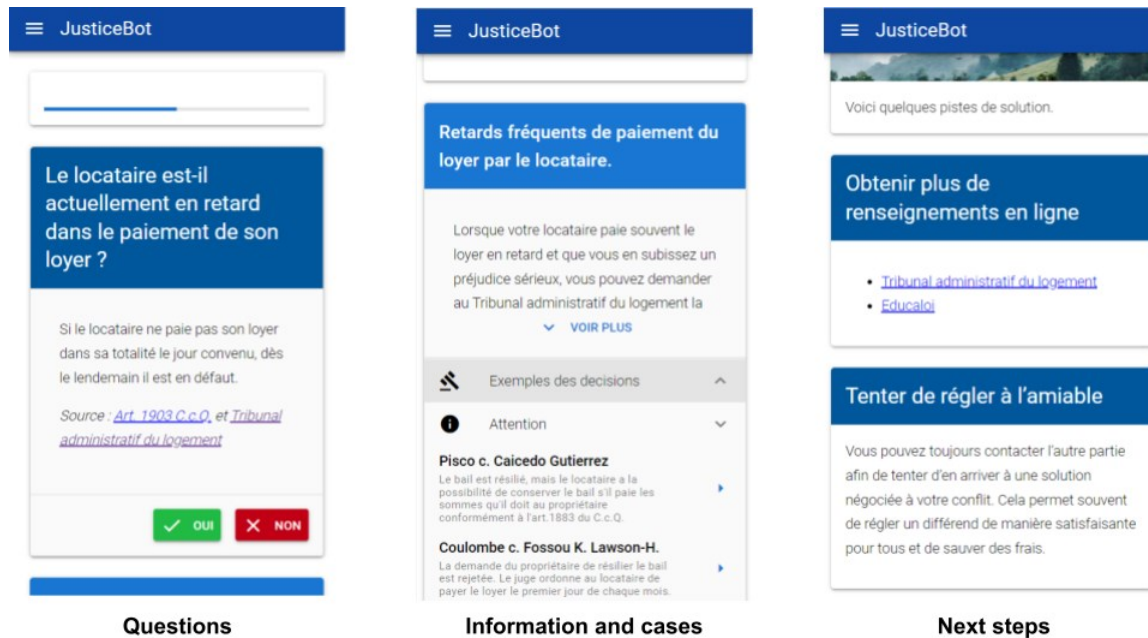


Figure 68 - Screenshots of the JusticeBot frontend

The JusticeBot frontend is a system that can read the schema of legal pathways and encoded cases exported by the JusticeCreator and expose it to the user. It can be easily made accessible on the internet and allows the user to access it via a smartphone or PC.¹⁴¹⁸ Figure 68 shows a few screenshots from the JusticeBot TAL, the first implemented version of the JusticeBot.

The user interacts with the system by first answering a few general questions, allowing the system to triage the issue of the user. The user is presented with a list of possible legal areas that the system can treat, and can choose one that corresponds to their situation, either because it is related to a goal that the user wishes to achieve (such as terminating their lease), or to a factual situation that the user has experienced (such as an infestation of bedbugs).¹⁴¹⁹

¹⁴¹⁸ See 7.4.1.1.

¹⁴¹⁹ See 7.2.2.

Based on their selection, the user answers a number of in-depth questions asked by the system, corresponding to the legal reasoning schema encoded using the JusticeCreator. For some questions, the user is shown summaries of how that specific question was answered by judges in previous situations, to help them make their decision. The simplified information and case law examples make it possible for layperson users to interact with the system.¹⁴²⁰

At the end, the system analyzes the provided information, and displays information regarding the situation of the user, and examples of results from previous cases. Further, they are given information about possible next steps that they can undertake.¹⁴²¹

At each stage, the user is able to return to a previous question, in order to change an answer and see the effects. The user is further able to provide feedback through a number of mechanisms, including star ratings on each page, a survey at the final page and a form that they can fill out if their question is not yet treated by the system.

The JusticeBot methodology and toolchain thus allows legal experts to build fully functional legal decision support tools. Next, I will discuss some important aspects of this methodology.

10.4 Discussion

Now that I have presented the JusticeBot methodology, let us delve into some of the particularities of the methodology, and discuss some interesting aspects of the approach.

As I see it, the most important motivating factor for the JusticeBot methodology is the desire to create a practical methodology, that can be used to implement real-world tools. Susskind discusses the differentiation between purists and pragmatists.¹⁴²² I believe the JusticeBot methodology to be very much on the pragmatic side, focused on building tools that can increase access to justice in the real world.

¹⁴²⁰ See 7.2.2.

¹⁴²¹ See 7.2.3.

¹⁴²² See 5.5.

This practical focus has shaped the development of the methodology, and been instrumental in many of the choices made. Part of this discussion will focus on the choices made to enable the building of practical legal decision support tools, within the current technological and legal constraints.

10.4.1 Augmented intelligence instead of artificial intelligence

One of the important insights of the FactorBot research was that predicting the outcomes of cases in a way that can be exposed to laypeople is difficult. In fact, I speculated that perfect legal prediction may be an impossible task – even lawyers are not able to indicate to an individual if their case will succeed or not with 100% certainty. Lawyers are, however, able to understand the unique situation of an individual, including how any new factors or policy considerations may affect a particular case. As discussed in 4.4.3, this kind of reasoning may be AI-complete, requiring general AI systems with common sense to fully carry out, which is beyond the current state of the art.

Based on these insights, the JusticeBot methodology takes a different approach. Tools built using this methodology do not seek to predict the situation of the user. Instead, the system seeks to *augment* the intelligence of the user, by providing them with relevant contextual information from previous cases.

In doing so, the system asks the user to perform tasks that it is not able to do. Since predicting whether certain criteria apply in a case based on input from laypeople may be beyond the scope of AI systems, the user is asked to carry out this task. However, the system supports the user in performing this task with the support of relevant case criterion summaries.

However, as we have seen, expert systems are good at logically traversing complex sets of logical rules. This is also an important part of legal reasoning, as rules are traversed to arrive at outcomes.¹⁴²³ In the JusticeBot, this task is performed by the computer system,

¹⁴²³ See 4.5.

which decides which questions to ask the user. The user thus does not have to think about the structure of the law, but merely about answering questions.

Finally, the system provides the user with summaries of the outcomes that previous similar cases resulted in. The system does not attempt to predict the outcome of a case, but rather gives the user an overview of the previous outcomes that courts have ordered in cases similar to theirs. This augments the intelligence of the user, by providing relevant, specific information regarding previous court cases. The user can use this information as a factor in their decision making.

I see this approach as greatly contributing to the practicality of the JusticeBot methodology. By tacitly acknowledging that some tasks (such as assessing the applicability of legal criteria and predicting the outcomes of cases) are beyond the scope of the system, it is possible to build a system that plays to the strengths both of the human user and the computer system. This approach is also very much in line with the seminal Human-Computer Interaction research by Bush and Licklider, both of whom suggested that the computer should symbiotically interact with the user, by helping them navigate enormous amounts of information. However, the human fills the gaps of the computer systems.¹⁴²⁴ The JusticeBot supports the user by selecting relevant cases from an enormous amount of case law, and showing these to the user. However, the user fills the gap in the system requiring common sense understanding.

Instead of trying to predict the outcomes of cases based on information provided by a layperson, the JusticeBot system focuses on providing the user with useful information, using a practical approach. Thus, it can be seen as a system that indexes previous cases, rather than predicting the individual user case. Let us explore this perspective a bit more.

10.4.2 Indexation instead of prediction

An important feature of the JusticeBot methodology is the focus on indexing previous cases rather than aiming to predict the outcome of new cases.

¹⁴²⁴ See 1.3.5.2.

Many AI systems aim to build a model of how certain inputs map to certain outputs. The FactorBot is an example of such a system, that aims to take a list of facts, and predict the outcome of the case.

However, as discussed in the previous section, such prediction can be tricky to perform correctly, due to the many complex steps involved in legal reasoning. Further, it is an open question whether predicting the outcomes of cases is even the right approach in the legal field. Starr discussed this point in the context of evidence-based sentencing tools. Here, the tools aim to predict the recidivism of an individual. Starr argues that what the tools are actually doing is to “predict the average recidivism rate of individuals who share with the defendant whichever characteristics are included as variables in the model”.¹⁴²⁵ The same criticism could be seen to apply to legal prediction models – in the end, while the model may seem to predict the outcome of an individual case, it actually provides information about how previous cases that share characteristics with the current case were decided. This does not have to have any direct bearing on the current case, since judges may decide to assess it in a completely different manner, based on the particularities of the case. Atkinson *et al* make a similar point, arguing that giving a prediction that a case is 80% likely to go a certain way would mean that 1 in 5 cases are not decided correctly, which is not justice.¹⁴²⁶ In prior work, many of the systems have instead focused on generating arguments for the user, or explaining the prediction in a way that can help the user understand the case, rather than the prediction being the main feature.

The JusticeBot methodology makes this underlying assumption explicit. The tools built using this methodology do not claim to be able to predict the outcome of new cases. Rather, the JusticeBot tool seeks to provide the user with information about how cases that are similar to theirs were *previously* treated at a court. Of course, there is a strong possibility that the case of the user will be treated in the same way as previous similar

¹⁴²⁵ Starr, *supra* note 537 at 842.

¹⁴²⁶ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 3.

cases, due to the concepts of local and personal stare decisis – we expect similar cases to be treated equally under the law.¹⁴²⁷ The user can use this knowledge to make better decisions, without being potentially misled by believing that the system is able to fully predict the outcome of their case.

I believe that this approach contributes substantially to the practical nature of the system. The JusticeBot methodology does not claim to solve the difficult or even impossible task of predicting new cases from information provided by laypeople. Rather, it claims to be able to index legal information in a way that it can retrieve cases that are useful to the user, based on information provided by them. In a way, the system can be seen as an intelligent search engine of previous cases, that uses a novel indexation scheme to be able to surface cases that could be relevant to the user. This task is more feasible than predicting new cases, meaning that such tools can be built and deployed to the public today.

Of course, this framing of the system puts significant importance on how cases are indexed by the system. Which cases are seen as similar? Next, I will explore how the methodology relies on a legal realism view to encode case law.

10.4.3 A novel way of encoding rules and cases, based on legal realism

The JusticeBot methodology introduces a novel way of dealing with case law, inspired by legal realism.

The JusticeBot methodology is focused on commonplace cases, where judges treat everyday decisions, rather than landmark cases, that may be useful in making an argument.¹⁴²⁸ The purpose of the JusticeBot is to use these commonplace decisions to provide contextual information to the user, based on how their case might be treated should it go to court. The way cases are treated is in line with the concept of legal realism¹⁴²⁹, which focuses on discovering how cases are *actually* dealt with at court,

¹⁴²⁷ See 4.9.3.

¹⁴²⁸ Compare 4.9.3.

¹⁴²⁹ See 1.3.2.1.

rather than learning what the legislation says about a certain legal area, i.e. how their case *should* be treated by courts. Since the purpose of JusticeBot tool is to inform the user of how their case may be treated at court, this realistic view may be more interesting to them.

Cases are thus used to discover how courts have previously treated certain cases. Each case is seen as the judge traversing a path of legal criteria that stem from a statute or precedent, applying these criteria to the case that they are dealing with, and deciding on an outcome. Thus, previous cases have three separate uses:

- Cases are used to discover the legal reasoning schema that judges tend to apply to deal with cases in a certain area, and to validate this schema.
- Cases are used to illustrate the reasoning that judges have used to decide whether a legal criterion applies or not.
- Cases are used to illustrate the outcomes that judges have awarded in previous cases, if certain legal criteria are fulfilled.

Let us briefly explore these uses of previous cases.

10.4.3.1 Using cases to discover a schema for legal reasoning

First, cases are used to discover the pathway of legal criteria that are applied by judges to solve certain legal cases. By reading statutes and cases, the legal expert building a JusticeBot system will get an overview over the real-world legal reasoning pathway traversed by judges in dealing with the cases. This is unlikely to be easy – each case just describes a single pathway through the schema, depending on what is found by the judge in a case, and thus only contains a small puzzle piece of the overall schema of possible reasoning paths. However, by reading many cases in conjunction, the legal expert can recreate the legal reasoning schema that judges use in dealing with certain cases, including the important criteria and legal conclusions. This can be seen as related to the

grounded theory method,¹⁴³⁰ as the schema, which is used to index the cases, is discovered from the cases themselves. As more and more cases are read in conjunction with the schema, the correspondence of the schema with the way judges actually reason about the cases is validated.

By focusing on areas with a lot of cases, this approach can be used to overcome the syntactic ambiguity of the legislation, by discovering how laws are applied in practice.

10.4.3.2 Using case summaries to illustrate how individual criteria are applied

Of course, the legal structure itself is not sufficient to learn the outcome of individual cases. For example, the structure does not inform the user of whether a legal criterion (such as “reasonable”, or “frequently late”) applies in their case.

Therefore, the second use of cases is to provide an exemplification of how judges tend to apply legal criteria contained in the schema. The legal expert creates case-criterion summaries, that summarize the facts that lead to a certain legal criterion being fulfilled or not. These summaries are then added to the schema, and can be read by the user of the system, to understand how the criterion was previously applied, and better understand how the criterion may apply in their case.

This method of encoding case law is very flexible, since the legal expert is able to summarize cases in free language, rather than having to rely on a rigid encoding. Further, fewer cases are needed, since the cases are not used to build a model of how a criterion is applied, but are instead used as an example of how the criterion was previously applied.

10.4.3.3 Using case law summaries to illustrate the outcome of cases

Finally, cases are used to illustrate the outcomes of previous cases. Once the user has entered a hypothesis about their situation, the hypothesis will be used to select a number of previous cases, and show the user the outcome of those cases. Thus, the user can understand the possible outcomes that they may obtain in court.

¹⁴³⁰ See 1.3.3.

This way of treating case law is a crucial part of making a system that is practical. As discussed, the system explicitly does not attempt to predict the case of the user, which as we have seen may be an AI complete problem, since there can always be exceptional cases. Instead, it focuses on providing a useful service to the user, by allowing them to easily enter a hypothesis about their situation, supported by case law summaries, and obtain information from cases regarding previous outcomes of similar cases. This information can be very helpful for the user in understanding how a court might treat their case, and thus support them in their decision making.

The focus on providing references to case law rather than predictions of the user case is also important for another reason: In many jurisdictions, predicting specific cases counts as giving legal advice, which is an activity exclusive to members of bar associations. Next, let us take a look at why the JusticeBot methodology provides legal information instead of legal advice.

10.4.4 Legal information instead of legal advice

An important consideration with regards to legal decision support tools is whether they constitute the unauthorized practice of law. While the rules may become more and more amenable to self-help tools, the uncertainty may still make it challenging to deploy tools that could be seen to give legal advice instead of legal information. The dividing line in many jurisdictions seems to be whether the tool gives the user opinionated advice regarding their specific case, or merely informs them of how cases tend to be treated in general.¹⁴³¹

In the JusticeBot methodology, the legal decision support tool never advises the user what they should do, or how their case is likely to be treated in court. Rather, it supports the user in empirically exploring previous cases. At all stages where forming an opinion is necessary, the system asks the user what they believe. For example, the system asks the user whether they believe that a certain legal criterion applies to their case, and supports

¹⁴³¹ See 3.5.3.

them by providing general information about how such criteria have previously been assessed. The system also does not tell the user the outcome they can expect should the case go to court – rather, it informs them of how previous cases where the judge found that certain criteria apply were decided.

Thus, the system can be seen more like a search engine than a lawyer giving legal advice. The system allows the user to enter a query, by specifying whether they believe that their situation fulfills any of the legal criteria inquired about by the system. After arriving at the end of a pathway, the full query is used to select legal information and legal cases from the database. This information is likely useful to the user, since they can use it to make a better decision. However, it consists of references to previous cases, and general legal information, rather than a specific analysis of the case of the user. In fact, all of the tasks that require an opinion on a specific situation are performed by the user. Just like websites that provide the functionality of searching for previous cases, such as CanLII, or a book that contains an explanation of the rules in a legal area, the JusticeBot should not be seen as giving legal advice. Therefore, it is possible to practically deploy JusticeBot-based systems, which is required to have a real-world impact on access to justice.

Even if a system does not provide legal advice, it should not be deployed if it can cause harm, for example by providing misleading or biased information. The JusticeBot framework therefore depends on a deterministic approach, where the creator of a system is always in full control over the information the user sees.

10.4.5 Determinism and predictability instead of probability

The alignment problem and lack of common sense have been identified as important shortcomings of AI systems. Since AI systems lack an understanding of human values, they can give information or advice that is harmful or discriminatory.¹⁴³² This is

¹⁴³² See 2.6.3.3 and 2.6.3.5.

especially problematic if the creator of the system has no way of verifying how the system will react in certain circumstances, or why it acts in a certain way.¹⁴³³

The JusticeBot approach is very cognizant of this issue. Therefore, the JusticeBot methodology uses the symbolic approach to provide information to the user. The system works by traversing question blocks and information blocks in a very intuitive way, akin to a flowchart. Understanding and reasoning about when certain questions are asked, and which information will be provided is thus trivial. The JusticeCreator further provides an intuitive way to create, edit and view the pathways of a system and previewing the logic of the resulting system.

Using the JusticeCreator and the JusticeBot methodology, the creator of the system fully defines how a JusticeBot tool will reason, and which information the user will see. Thus, there is no risk of the user being provided information that the creator did not approve. Systems can be fully vetted and understood for bias and inaccuracies. This feature is crucial in building tools that can be exposed to the public.

Due to the deterministic nature of the system, the decisions taken by the system are also fully transparent and explainable, both to the creator and the user of the system. There is no black box that uses complex mathematics to decide on an output. Rather, the legal expert is in all instances able to understand why a certain piece of information is shown, based on the flowchart as displayed by the JusticeCreator. Further, this transparency extends to the user. Since the legal sources, such as laws and case law, are included in the information given by the system, the user is able to verify why they are being given certain information, and where their rights stem from. The user is also always able to see the responses they gave and change any of them to understand how different responses may affect the information they are shown.

This does not mean that machine learning does not play a role in the JusticeBot methodology. The JusticeCreator integrates machine learning as a way to select cases for

¹⁴³³ See 2.6.3.4.

annotation, by searching for cases that contain phrases similar to a legal criterion. This can represent a significant gain in efficiency, as legal experts are able to find cases for annotation quickly, without having to think of possible synonyms and manually devising search terms.¹⁴³⁴ However, the legal expert decides whether a case warrants inclusion in the pathway, and how it should be summarized. Thus, the system acts to augment the intelligence of the legal expert (just as it augments the intelligence of the user), while still leaving them in complete control over the output the user will see.

Of course, to be able to practically deploy JusticeBot tools, an important pre-requisite is to create the tooling and interfaces required to build and publish such a tool. I have done this as part of my research. Next, I will discuss the JusticeCreator and JusticeBot interface.

10.4.6 An implemented, easy-to-use methodology

An important aspect of the JusticeBot methodology is that it is fully implemented, in a production-ready end-to-end toolchain. This toolchain has further been used to create a JusticeBot tool that is publicly accessible and has been accessed by thousands of users. Here, I will examine some important aspects of this concrete implementation.

10.4.6.1 The JusticeCreator

The crown jewel of the implementation is undoubtedly the JusticeCreator. Built as a web application, this tool can be used to create legal decision support tools based on the JusticeBot methodology.

The JusticeCreator follows is a no-code tool, i.e. it does not require any programming. This allows anyone to use it to create legal decision support tools, even without technical knowledge. All of the actions required to build a JusticeBot tool, including the structuring of the legal reasoning pathway, the writing and formatting of content for the pathway and the summarization of case law, can be done using this interface. Each of these functions can be done easily, after a few minutes of instruction.

¹⁴³⁴ See 7.3.2.4.1.

The JusticeBot methodology and the JusticeCreator deliberately use very simple building blocks. These make it very easy to build new systems, but also to understand why a question is asked, and which information will be provided at different points. Further, as we have seen, the conceptual underpinnings of the JusticeCreator map closely to the way judges and other legal decision makers reason, i.e. by applying legal criteria in sequence and eventually coming to legal conclusions, making it possible to use the methodology to build useful decision support tools.

The JusticeCreator was used in creating the first deployed JusticeBot version, the JusticeBot TAL. The process of building this decision support tool validated the choices taken in building the interface, as a team of legal experts without any programming knowledge was able to use it to build a working legal decision support tool. Further, the simplicity of the representation chosen allowed us to export the entire content of the JusticeBot TAL into a word document, that maintained the logical flow of the system, and pass this document to the Tribunal Administratif du Logement for verification.¹⁴³⁵ Likewise, legal teams are currently building additional JusticeBot versions using the JusticeCreator tool.

10.4.6.2 The JusticeBot frontend

The JusticeBot frontend is the other crucial component of the methodology. By allowing the ingestion of the schema exported from the JusticeCreator, it allows the creation of legal decision support tools based on the JusticeBot methodology. Since the frontend takes its content from the JSON-file exported by the JusticeCreator, it is fully domain independent, and can be implemented off-the-bat in any legal area. Any JusticeBot tool could thus get ease-of-use, mobile-friendliness and integrated analytics and feedback methods for free. Further, the JusticeBot is built like a static page web app, making it very cheap to host, even for millions of potential users.

¹⁴³⁵ See 8.3.3.

Just like the JusticeCreator, the JusticeBot frontend has been validated through building the JusticeBot TAL. The system has been running without issue since summer 2021, and been accessed by over 17k users. As we saw, the feedback has been very positive, with 86% of survey respondents answering that they would recommend the system to their friends.

As discussed in 7.8.2.4.1, the methodology has been used to create one tool that is publicly accessible. Several others are currently under development, and are showing promising initial indications. This is a positive sign for the ability of the toolchain to generalize to other legal domains.

In prior work, we saw multiple authors point to the lack of a toolset and methodology as important issues constraining the number of legal decision support tools.¹⁴³⁶ I hope that the methodology introduced in this thesis, coupled with the JusticeCreator and JusticeBot tools, can provide such a toolset and methodology, allowing lawyers to build legal decision support tools, thereby opening the door to many such tools being created, leading to an increase in access to justice and access to information and a positive impact on society.

We have seen many of the advantages of the JusticeBot methodology. Next, let us take a look at some possible limitations of the approach.

10.5 Limitations

Let us now turn to some of the limitations of the JusticeBot approach. Here, I will examine some of the limitations and trade-offs that are represented by the approach, and how they may be overcome.

10.5.1 User effort required

One limitation of the system is that the user effort required is relatively high. While in the FactorBot, the user merely had to provide a list of facts, in the JusticeBot, the user has to

¹⁴³⁶ See 5.5.1 and 5.5.4.

apply a legal criterion to their own facts. The user is supported in this assessment, through provided information and previous cases. However, if the user makes a mistaken assumption (e.g. indicating that their tenant should be seen as being “frequently late”, while a judge would not agree), the information they obtain at the end may not be accurate. In a sense, the system thus relies on the layperson user to perform a step of legal reasoning.

Requiring the user to perform this assessment is a deliberate trade-off of the JusticeBot methodology. As described above, the user carrying out the assessment themselves means that the system does not constitute the unauthorized practice of law, which is crucial for the real-world deployment of the tools. Further, as we have seen, the FactorBot was not able to accurately predict whether legal criteria apply or not. I argue that such a prediction may be impossible using current AI systems. While the JusticeBot approach may require more work from the user, it allows the building of useful legal decision support tools, without overcoming this limitation.

Further, even systems that rely purely on the “facts” of a case, such as the FactorBot, make a number of hidden assumptions about the user case. The first such assumption is that the user is able to prove the facts of their case. The second assumption is that the user understands what a fact means – laypeople may think of a situation very differently than a trained legal professional.¹⁴³⁷ If either of these assumptions turn out to be wrong, so will be the information given by the system. The JusticeBot approach makes these assumptions explicit, informing the user that they need to prove the facts of their situation, and providing summaries of and references to previous cases to allow the user to understand the judicial reasoning with regard to situations such as theirs.

There may also be situations where a JusticeBot tool does not require the assessment of legal criteria. First, this kind of assessment is only required with regard to legal criteria that are open-textured. Many questions, such as “Is the tenant more than three weeks late

¹⁴³⁷ See 4.4.3.4.3.1.

with paying their rent?” are clear enough to not require case law summaries to answer. In areas of high-volume, low-intensity disputes, or administrative procedures, such clear concepts may be the only assessments required to arrive at a conclusion.

There are also functions of the JusticeBot that are not impacted by an inability to determine whether a legal criterion applies or not. Even if the user is unable to perform this assessment correctly, they can understand that their problem has a legal solution, and find the right forum. They may also be able to use this information to settle their case. Further, upon seeing the possible outcomes of their case, such as low monetary values that are usually rewarded even if a fact can be proven, a user may decide that it is not worth the effort to pursue their case. The usefulness of these types of information are not conditional upon a user being able to determine how judges would see their case.

10.5.2 User questions not covered

Another potential limitation of the JusticeBot approach is that not every question of a user can be answered. At its core, the JusticeBot methodology consists of an expert system, meaning that the system cannot generalize to situations that have not been encoded.¹⁴³⁸ As we saw in the JusticeBot TAL, sometimes less than half of individuals are able to receive information regarding their question – the rest fall outside of the system, and end up on the “missing question” screen, that informs them that their issue is not yet covered.¹⁴³⁹

This is a limitation, but one that may be less serious than it seems. First, trying the JusticeBot is very low stakes for the user. We saw that users who end up on the missing question screen spend an average of 30 seconds on the platform. Therefore, they lose very little by trying the JusticeBot.

Further, this limitation may not be possible to overcome. We have seen that AI is not capable of common-sense reasoning. Thus, providing answers to questions of the user

¹⁴³⁸ See 2.5.2.3.

¹⁴³⁹ See 8.4.1.3.

without any form of training or encoding may be beyond the scope of current AI systems. Since the legal issues that people face in reality may be very varied, this necessarily means that some people with specific situations will always fall outside of the system. The only solution to this is to spend time in order to encode situations that users tend to face into the system.

The JusticeBot methodology provides a feedback option that allows individuals to describe their situation if it is not yet covered. This feedback allows the creator to identify the most important situations, and encode these into the system, to gradually increase the coverage of the system. In the meantime, even helping only a few users is better than helping no users.

10.5.3 Interpretation by the legal expert

Another possible limitation is that subjectivity may be required by the legal expert in order to build the system. If the legal expert needs to interpret a legal area to build the system, the user of the system will, in the end, interact with the interpretation of this legal expert.¹⁴⁴⁰ If the legal expert thus has made an incorrect assumption, the system may provide incorrect information.

The JusticeBot methodology has been specifically designed to reduce the amount of subjectivity that a legal expert introduces in the system. Each time a case is annotated, the legal decision support schema is validated. The interpretation of how the law itself should be interpreted is thus performed by the judge, and the legal expert integrates the interpreted version into the schema.

Despite this, there is necessarily some amount of interpretation that needs to be made by the legal expert. Sometimes, it may not be exactly clear which criteria a judge applies, and in which order. In these cases, the legal expert has to infer the reasoning path a judge took in annotating the case and creating the schema. Hopefully, in areas where there are

¹⁴⁴⁰ Compare 4.5.3.3.4.

many cases, reading more cases will clarify the correct schema of legal reasoning, applied by judges.

The legal expert also has discretion in which cases to summarize, and how they should be summarized. Currently, the JusticeBot platform requires the identification of a few (in our case, five) cases on each side of whether a legal criterion applies or not. Due to the limited number of cases that the user is able to read, the legal expert needs to make an important choice regarding which cases to include. I explain a few criteria influencing which cases to choose above in 7.3.2.2.3. Further, the legal expert needs to decide how the cases should be summarized in a few sentences (see 7.3.2.2.4). Of course, there is not a single correct way to summarize such a case, introducing another element of subjectivity.

While I have tried to minimize these elements of subjectivity, by providing guidelines on how to make the choices such as creating the schema, and choosing which cases to include, it is important to note that there are still decisions that the legal experts needs to take. I think this is inevitable, as legal reasoning is not a hard science. Therefore, the best we can do is to clearly specify the considerations that the legal expert should take and be honest with the user about what exactly the system is doing.

Beyond requiring interpretation, there may also be areas where it is essentially impossible to identify a single coherent schema. Let us examine what happens in these cases.

10.5.4 Complex cases and disagreement among judges

The JusticeBot works best where the interpretation of a judge occurs regarding specific criteria (i.e. “Do these facts mean that the rent is paid frequently late?”) as compared to on the syntactic level (i.e. “Does the rent being frequently late warrant a lease termination?”). Let us take a look at what happens when the syntactic reasoning in an area is not consistent.

In some legal areas, there may not be a single coherent legal schema that can be created, making it difficult to build a JusticeBot tool. This could be the case if there are very few legal cases, since it becomes difficult to infer the overall schema. It could also be the case

where the cases tend to involve a lot of varying arguments by the parties. Encoding these could be difficult, and lead to users having to answer a lot of very specific questions that are not directly related to them. Here, each case can be seen as having its own legal reasoning schema, with different criteria being applied depending in each case.

This is a limitation of the system. However, the system is targeted at areas of high-volume, low-intensity cases, where there are often fixed schemas that are applied to an individual situation by a judge or other legal decision maker, such as administrative officers. When determining whether a lease can be terminated due to late rent payment, or whether an individual qualifies for a certain form of social aid, the structure of the reasoning is likely very similar for each case. I discussed a possible approach to handle more complex cases above in 9.4.4.3.

Another situation where the JusticeBot might face challenges could occur if the case of the user is very complex. For example, if the user is late with paying their rent because their landlord stole their money, the user may answer all of the questions in the system, and be told that they may not have the right to terminate their lease, since no question captures the criterion of theft. As discussed, AI systems in general have issues dealing with the effect of completely new factors. Once more, however, the JusticeBot does not claim to predict the unique situation of the user, but rather to give them previous similar cases. The user is told that their case may have unique features that makes it different from the previous cases.

It may also be difficult to create a JusticeBot tool where judges disagree on the criteria that are relevant to solve certain cases. The JusticeBot is able to handle differing interpretations of the same criterion (i.e. different judges seeing “frequently late” as meaning different things) by including summaries from the different perspectives and showing them to the user. However, if the logical connections between the criteria are inconsistent, it is not possible to represent these in the JusticeCreator.

This is also a limitation. However, it is not clear whether there is a correct solution in such situations. If there is no consistent way of deciding certain cases, any legal decision support tool would have to choose one of the competing interpretations. One possible

solution could be presenting different schemas as arguments (compare 4.4.3.4.2.7), however it is not clear how useful this would be for laypeople.

This brings us to the next possible critique of the system – why does it focus on providing information instead of helping the user argue their case?

10.5.5 Information instead of arguments

In prior work, a lot of research focuses on generating arguments for either side of an issue.¹⁴⁴¹ An argument can be used for educational purposes, or to help legal professionals make their case more efficiently.¹⁴⁴² In some instances, it may be more useful to obtain an indication of how a case should be argued in order to be won, rather than receiving a context-less prediction of whether a case will be won or lost – as we have seen, the decision of a judge can depend on many factors. One might argue that the JusticeBot would be more useful if it also focused more on generating arguments for the parties.

However, the JusticeBot targets laypeople, and is intended to be used right after a situation has arisen. At this stage, the user is not necessarily aware that there is a legal solution to their problem. Thus, it cannot be assumed that the user is interested in arguments that may be useful in court – they may not even know if they have a claim at all, or even wish to go to court. At this stage, informing the user of their potential rights, and possible outcomes to their case, might therefore be more useful in helping them understand their situation and how they might want to proceed.

Further, I believe that legal argumentation may be less important in the high-volume, low-intensity cases targeted by the JusticeBot methodology. As we have seen, such cases are often governed by relatively rigid reasoning schemas. The legal decision maker (whether they are an administrative officer or a judge) seem more focused on assessing whether the legal criteria are fulfilled based on a factual situation, and deciding on the

¹⁴⁴¹ Atkinson, Bench-Capon & Bollegala, “Explanation in AI and law”, *supra* note 46 at 1.

¹⁴⁴² See e.g. 4.4.3.4.

consequences of this, leaving less room for legal argumentation by the parties. Such argumentation may be more relevant at higher court instances, which are beyond the scope of the JusticeBot.

Thus, I believe that laypeople facing everyday problems can be helped significantly by informing them of their potential legal rights, which criteria and facts are relevant, and how to practically proceed with their situation. That said, I also explored how the JusticeBot methodology could be used to structure party submissions in 9.4.4.1.1.

The type of reasoning described above, which relies on assessing criteria defined by statutes, may be seen as being characteristic of civil law systems. Let us explore whether the JusticeBot methodology could be useful both in common law and civil law jurisdictions.

10.5.6 Civil law or common law?

Once created, the schema of the JusticeBot remains static. While the user is able to reason about the individual legal criteria, the system does not support reasoning about different syntactic structures. Let us explore whether this makes the system more useful in civil law jurisdictions than in common law jurisdictions.

In the civil law, statutory law is the main source of law. Courts are seen as merely interpreting the law, to apply it to new cases.¹⁴⁴³ This style of reasoning is a perfect fit for the JusticeBot, since the law serves like an outer constraint that encompasses the reasoning of legal decision makers. Especially at the first instance, which is likely to be the most interesting for laypeople, and in areas of high-volume, low-intensity cases, it would seem likely that most of the interpretation of the legal decision maker lies in whether a factual situation fulfills certain legal criteria, mirroring the structure of the JusticeBot reasoning schemas.

In the common law, previous cases are incorporated in sophisticated reasoning steps. These cases are treated as a collection of material facts, that are tied to an outcome. By

¹⁴⁴³ See 4.9.2.

analogizing to certain cases, and distinguishing other cases, judges come to a decision.¹⁴⁴⁴ In this system, the courts seem to have a much larger freedom to affect the syntactic structure of legal reasoning, by drawing analogies to previous cases and reinterpreting how cases should be understood in the light of previous facts. This kind of reasoning may be more difficult to capture in a schema, as is being done in the JusticeBot, which relies on there being a rigid, legislation-like structure that is used to determine the outcomes of cases. Systems targeting such areas, such as HYPO and CATO, seem to place more emphasis on generating arguments and predicting cases by comparing cases across different layers of abstractions.¹⁴⁴⁵

However, there may be instances where there is a rigid reasoning schema, even in common law systems. When administrative decision makers determine whether social aid should be granted, or judges determine whether a lease should be terminated at the first instance, it seems unlikely that every case relies on common law style reasoning, as described above. Since lay people are unlikely to have the resources to appeal a decision or seek clarification on a nuanced point of law, they may often be interested in commonplace decisions, where local and personal *stare decisis* is more important than traditional *stare decisis*. Therefore, even in common law systems, I believe that there are many legal areas where useful JusticeBot tools could be created.

10.5.7 Rule skepticism?

The JusticeBot methodology relies on the assumption that legal rules are an important component in legal decision making. As we have seen, some researchers argue that the laws are irrelevant for judicial decision making, as judges are influenced more by what they had for breakfast, or other extraneous factors, than the actual legal rules in making decisions.¹⁴⁴⁶

¹⁴⁴⁴ See 4.9.1.

¹⁴⁴⁵ See 4.4.3.4.2.2.

¹⁴⁴⁶ See 1.3.2.1.

First of all, I question the assumption of rules not being relevant for legal decision making. When it comes to administrative decisions, for example, it is often very clear which criteria need to be fulfilled for a certain outcome. In these cases, it is obvious that the rules are important. Further, internal guidelines often constrain the decision makers in how they can exercise their discretion. Likewise, in high-volume, low-intensity cases, it is very clear which criteria need to be fulfilled to achieve a certain outcome.

Of course, this does not make it impossible that rules only account for some of the outcome of a case. For example, in assessing an open-textured legal concept or making a discretionary determination of an outcome, a judge may hypothetically let extraneous factors influence their decision, including potential biases.

Even if this is true, it is not clear how such information would be helpful to the user of a JusticeBot legal decision support tool. The tool is intended to be used right after a situation or need has arisen, before the user has even decided what to do about their situation. At this point in time, it is impossible to determine which judge will preside over a case, or which extraneous factors may affect the judge at the time of the hearing. Thus, in the JusticeBot, I focus on what we can deal with at this stage, namely by supporting the user in understanding their potential rights and potential outcomes, based upon which legal criteria a judge may find to apply to their situation.

10.5.8 Perpetuating bias

In 3.6.2.4, I explored the risk of machine learning systems perpetuating biases inherent in society. Let us explore the extent of this risk when it comes to the JusticeBot methodology.

As discussed in 10.4.5, the JusticeBot takes the approach of relying on a predictable system, created and verified by legal experts. This means that the user never directly interacts with a machine learning-based system, that may have observed discriminatory trends from data, and would e.g. give increased chances of success to people of certain ethnicities or genders. Further, the system seeks to empirically model legal rules as applied by the legal system, limiting the potential influence and bias stemming from the creator of the system.

Of course, this does not mean that no bias can enter into the system. This may be the case, for example, if previous case decisions regarding certain legal issues are biased. Such biases, that stem from the distribution of the training data itself, are referred to as “historical biases” by Suresh and Guttag.¹⁴⁴⁷ For example, judges in certain types of cases may be biased against parties belonging to certain groups, e.g. awarding lower damages to people that are poor.

Such biases would not affect the output of the JusticeBot system directly. The system never asks the user for their socio-economic status, their gender, or their ethnicity. It only asks the user whether they believe that certain legal criteria are fulfilled in their case. Thus, the output seen by people with different characteristics will be the same, provided that they select the same options.

But bias is tricky. While the JusticeBot does not provide any direct discrimination, the discrimination could still be included in an indirect manner. People with different characteristics are likely to be affected by different kinds of issues. For example, as previously discussed, poor individuals are more likely to be affected by bedbug infestations.¹⁴⁴⁸ Women, people that are not heterosexual, and women belonging to Indigenous groups are more likely to be affected by gender-based violence.¹⁴⁴⁹

If people belonging to these groups were discriminated against in certain jurisdictions, this bias would be reflected in the outcome of the cases of these types. The JusticeBot, which relies on such cases to give information, would thus contain these biases. This could result in users with certain cases being shown similar cases that resulted in lower damages than they should have in an ideal world.

There may also be other, more overarching, biases that are connected to the JusticeBot system. For example, people belonging to certain groups may have better access to computers and smartphones, which are needed to access the JusticeBot, than others.

¹⁴⁴⁷ Suresh & Guttag, *supra* note 166 at 4–5.

¹⁴⁴⁸ See 8.2.2.

¹⁴⁴⁹ See 9.3.1.7.

Further, as discussed in 3.6.2.6, the main data source of the JusticeBot is case law, which may introduce its own sets of issues, as many cases are settled before ever reaching court, and are thus not part of the JusticeBot database.

Understanding, and deciding how to respond to these potential sources of bias, is undoubtedly an important step in future research. However, it is also important to note the purpose of the JusticeBot. Throughout the system, it is very clear that the system does not aim to answer the question of how a case *should* be dealt with, or what the user should do. Instead, it exclusively aims to help the user understand their case and how other cases such as theirs were *previously* treated by legal decision makers. Even cases that are affected by bias may still be a part of accurately informing the user of how cases in a certain area have historically been decided.

We have now discussed some of the limitations of the JusticeBot methodology approach. Next, let us discuss whether decision support tools created using the methodology may increase access to legal information and access to justice, which is the main research question of this thesis.

10.6 Can the JusticeBot methodology increase access to justice and legal information?

To wrap up, I wanted to reconnect to the main research question of this thesis:

How can artificial intelligence be used to increase access to justice and access to legal information through the creation of a methodology for developing legal decision support tools?

Previously in this chapter, I summarized and discussed the JusticeBot methodology, which is the main result of my research. As we have seen, the methodology was successfully used to implement the JusticeBot TAL, a decision support tools accessed by thousands of users. It thus seems like the methodology that I developed can be used to create decision support tools using artificial intelligence.

However, in order to fully answer the question, I will also need to explore whether the methodology can increase access to legal information and access to justice. I will do so in this section.

It is important to answer the question with humility. Access to justice and access to legal information are immensely complex issues, encompassing many socio-economic and legal aspects. As I explored in 9.4.2.1, in order to fully understand the impact of JusticeBot tools on the user, a user study would have to be conducted, which is beyond the scope of this research and thus left for future work.

However, I believe the initial results to be very promising, both in theory and in practice.

We saw that laypeople have issues with regards to legal information and access to justice. Users think of their situation in terms of a factual occurrence or goals, rather than in terms of the legal or administrative procedures that correspond to their situation. This means that they are not aware of their rights and obligations, or how to enforce them in a specific forum. Laypeople struggle, for example, with identifying which form to use to file a claim, and often have a poor court experience.¹⁴⁵⁰

JusticeBot legal decision support tools, such as the JusticeBot TAL, should be able to overcome this issue. They allow the user to select a situation or goal, and then explore the legal rules linked to this situation or goal. In doing so, the user is not required to understand the law. Rather, the system guides them through the relevant rules, and gives them the tools to understand how judges may decide on their situation. At the end of the pathway, the user is given information about the legal rights they may have, previous outcomes, and possible next steps.

The information provided by such tools can be seen to increase *access to legal information*.¹⁴⁵¹ The tool makes the law accessible, by allowing the user to explore their rights. However, it goes beyond just providing the user with access to the law, as it can

¹⁴⁵⁰ See Chapter 3.

¹⁴⁵¹ Compare 3.4.

also be used to understand the law. The simplified explanations can help the user understand what the legal terms mean. Further, the tool allows the user a realist insight into how courts interpret the laws, and how their situation may thus be assessed by legal decision makers. This understanding is a crucial aspect of the rule of law, as it can help the user arrange their affairs, and take well-informed decisions regarding their situation. Understanding the rules that affect an individual may also increase their sense of belonging and inclusion in society.

Likewise, the tool has the potential to increase *access to justice*.¹⁴⁵² The information provided by the system can give the user a much better standing in deciding how to proceed with their situation. With the help of the information, they can understand the avenues of solution to their issues, such as realizing that their situation does, in fact, have a legal solution. The information can further give laypeople an enormous leg up in settling their case. Since the user understands their legal rights and potential outcomes, they can rely on this information in negotiations with the other party, to create an amicable solution that is beneficial to both parties, allowing the continuation of a peaceful co-existence after the settlement. Finally, if the user is unable to settle their case, the information provided by the JusticeBot could be helpful in deciding whether to hire a lawyer or whether they want to take their case to court, and how to do so. Finding practical ways to resolve their issues can lead to an increase in welfare on a societal level, and an increase of trust in legal institutions.

In practice, the feedback received through the JusticeBot TAL shows that users are generally happy with the system. 57% of survey respondent indicated that they received the information necessary to understand their situation from the JusticeBot, while 53.6% of the users responded that they gained a good idea of the next steps from using the JusticeBot.¹⁴⁵³ These numbers may seem relatively low, however I believe that they represent a big success. Over half of the users were able to understand their legal

¹⁴⁵² Compare 3.3.

¹⁴⁵³ See 8.4.3.6.

situation and how they should proceed through the use of the tool. Since, as indicated by Chapter 3, the alternative may be that very few users understand their legal situation or how to proceed, this could represent a significant gain with regards to access to legal information and access to justice.

Further, 86% of individual indicated that they would recommend the JusticeBot to their friends.¹⁴⁵⁴ This indicates that even users that were not able to gain a better understanding of their case through use of the JusticeBot saw the potential of the tool and would share it with others. I consider this to be a significant endorsement, and a validation that real-world users seem to think that the tool is useful.

In order to have the biggest possible impact on access to justice, many JusticeBot tools should be created, covering multiple domains in many jurisdictions. Therefore, a key aspect of my research has been to design a methodology that is repeatable and can be used by any legal expert to create legal decision support tools. Through the description of the methodology, and the creation of easy-to-use tools like the JusticeCreator and the JusticeBot interface, I hope that many more such tools can be created, helping individuals in diverse situations and areas. Currently, several further JusticeBot tools are under development at the Cyberjustice Laboratory.

It is important to acknowledge that the JusticeBot is unlikely to be enough to fully overcome the aforementioned issues with access to justice and access to legal information. As mentioned, these are complex, inter-disciplinary issues, requiring broad coordinated action and institutional buy-in to fully tackle. However, it is my hope that the JusticeBot can represent a meaningful contribution to the field, and a step in the right direction. I hope that it can lay at the basis for many legal decision tools that can increase access to justice and legal information, and serve as a starting point for further research that will explore how it could be expanded and integrated with, for example, online

¹⁴⁵⁴ See 8.4.3.6.

dispute resolution, to have an even greater impact on the well-being of individuals across society.

10.7 How will tools like the JusticeBot impact the legal profession?

Finally, let us take a more general look at how tools such as the JusticeBot could affect the legal profession.

As discussed, lawyers may be concerned that AI tools may replace them in the future. But as evidenced by the comparatively limited scope of supporting the user in understanding their legal situation and deciding how to proceed, JusticeBot tools clearly do not try to replace lawyers. Lawyers are able to legally analyze the unique situation of the client, recommend next steps, draft documents and even represent them in court, which are all beyond the scope of the JusticeBot.

However, in cases where consulting a lawyer is not an option for the user, due to the possible cost exceeding their financial means, AI-based tools such as the JusticeBot can provide a very low-threshold option for individuals who wish to understand their situation and be supported in resolving their legal issue.

In other instances, AI-based legal information tools can be the first step for users to decide whether to pursue their case. A user might not have consulted a lawyer at all because they assumed that they would lose in court, or were not aware that their situation had a legal solution at all. Thanks to the information provided by an AI tool, they may realize that they have a right, or that their chances are better than expected and therefore decide to hire a lawyer. In such cases, the system functions as a necessary precondition for engaging a lawyer.

Thus, tools such as the JusticeBot should not be seen as threats to lawyers. Rather, they can extend the scope of the legal system, to serve populations that previously had few effective ways of understanding how to resolve their conflicts. JusticeBot tools could also be a step toward a society where people are able to find amicable solutions to their legal problems before a lawsuit or outside of a lawsuit, which may also require the involvement of a lawyer for further negotiation and settlement.

At the same time, it seems unlikely that legal work will be unaffected by the recent advances in AI systems. I discussed ways that the JusticeBot methodology could be used to support lawyers, judges and government employees. However, again, none of these approaches aim to replace these positions – rather, they aim to support the person in working more efficiently, while still relying on the human to perform complicated cognitive tasks that are beyond the scope of the current AI systems.

In this sense, I believe it is currently more useful to see the AI systems as tools that support humans, rather than machines that are autonomously able to perform complicated tasks in the legal field and replace human legal professionals, just like the hammer did not replace the carpenter. I believe that AI is poised to make the legal field more accessible, effective and diligent, and I sincerely hope that the JusticeBot methodology can represent a valuable contribution towards this goal.

Bibliography

Statutes & Cases

Act respecting labour standards, CQLR c N-11.

Act respecting the Administrative Housing Tribunal, CQLR c T-1501.

Act respecting the Barreau du Québec, CQLR, c B-1.

By-law respecting the allocation of dwellings in low rental housing, S-8, r 1.

Canada Evidence Act, RSC, 1985, c C-5 2022.

Charter of human rights and freedoms, C-12.

Civil Code of Québec, CQLR c CCQ-1991.

Code of Civil Procedure, CQLR c C-2501.

Council Regulation (EU) No 216/2013 of 7 March 2013 on the electronic publication of the Official Journal of the European Union, OJ L 2013.

European Convention for the Protection of Human Rights and Fundamental Freedoms, as amended by Protocols Nos. 11 and 14, (Council of Europe) 1950.

Loi visant la création d'un tribunal spécialisé en matière de violence sexuelle et de violence conjugale, LQ 2021, c 32 2021.

Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS, COM/2021/206 final.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), OJ L119/1 2016.

Regulation (EU) No 524/2013 of the European Parliament and of the Council of 21 May 2013 on online dispute resolution for consumer disputes and amending Regulation (EC) No 2006/2004 and Directive 2009/22/EC (Regulation on consumer ODR), OJ L165/1 2013.

Regulation respecting mandatory lease forms and the particulars of a notice to a new lessee, CQLR c T-1501, r3.

Rules of procedure of the Administrative Housing Tribunal, CQLR c T-1501, r5.

Hryniak v Mauldin, 2014 Supreme Court of Canada, 2014 SCC 7 (CanLII), 1 SCR 87.

Reference re Code of Civil Procedure (Que), art 35, 2021 Supreme Court of Canada.

Books

American Bar Association & Temple University, eds, *Legal needs and civil justice: a survey of Americans: major findings of the comprehensive legal needs study* (Chicago, Ill: Consortium on Legal Services and the Public, American Bar Association, 1994).

Ashley, Kevin D, *Artificial intelligence and legal analytics: new tools for law practice in the digital age* (Cambridge New York Melbourne Delhi Singapore: Cambridge Univ Press, 2017).

———, *Modeling Legal Argument: Reasoning with Cases and Hypotheticals* (Cambridge, MA.: The MIT Press / Bradford Books, 1990).

———, *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*, PhD Thesis (Amherst, MA, USA: University of Massachusetts, 1988).

Benyekhlef, Karim et al, eds, *eAccess to justice, Law, technology and media* (Ottawa: University of Ottawa Press, 2016).

Canada, Law Society of Upper, *Access to Justice for a New Century: The Way Forward* (Law Society of Upper Canada, 2005).

Cane, Peter & Herbert M Kritzer, eds, *The Oxford handbook of empirical legal research*, 1st ed, Oxford handbooks in law (Oxford ; New York: Oxford University Press, 2012).

Chapman, Francis, *Principles of the Law of Evidence with Illustrative Cases* (Philadelphia: Cyrus M. Dixon, 1930).

Davis, Stephen F & William Buskist, *21st Century Psychology: A Reference Handbook* (SAGE, 2008).

Ekbia, H R, *Artificial Dreams: The Quest for Non-Biological Intelligence* (Cambridge: Cambridge University Press, 2008).

Farrar, John H & Anthony M Dugdale, *Introduction to Legal Method*, 3rd revised ed (London: Sweet & Maxwell, 1990).

Fick, Göran & Ralph H Sprague, *Decision Support Systems: Issues and Challenges: Proceedings of an International Task Force Meeting June 23-25, 1980* (Elsevier, 2013).

Fisher, Roger, William Ury & Bruce Patton, *Getting to Yes: Negotiating Agreement Without Giving In*, revised ed (New York: Penguin Books, 2011).

Gagnon, Pierre & Isabelle Jodoin, *Louer un logement*, 2e ed (Cowansville, Québec: Éditions Y. Blais, 2012).

Gélinas, Fabien et al, *Foundations of Civil Justice: Toward a Value-Based Framework for Reform* (Cham: Springer International Publishing, 2015).

Genn, Hazel G, *Paths to justice: what people do and think about going to law* (Oxford, England ; Portland, Or.: Hart Pub, 1999).

Glenn, Patrick, *Legal Traditions of the World: Sustainable Diversity in Law*, 5th ed (Oxford, New York: Oxford University Press, 2014).

Goodfellow, Ian, Yoshua Bengio & Aaron Courville, *Deep learning*, Adaptive computation and machine learning (Cambridge, Massachusetts: The MIT Press, 2016).

Grabmair, Matthias, *Modeling purposive legal argumentation and case outcome prediction using argument schemes in the value judgment formalism*, PhD Thesis (University of Pittsburgh, 2016).

Haugeland, John, *Artificial intelligence: the very idea* (Cambridge, Mass: MIT Press, 1985).

Heuschling, Luc, *État de droit, Rechtsstaat, Rule of law*, Nouvelle bibliothèque de thèses 16 (Paris: Dalloz, 2002).

Holland, James & Julian Webb, *Learning Legal Rules: A Students' Guide to Legal Method and Reasoning*, 10th ed (Oxford, United Kingdom ; New York, NY: Oxford University Press, 2019).

Lawless, Robert M, Jennifer K Robbennolt & Thomas S Ulen, *Empirical Methods in Law*, 2nd ed (New York: Wolters Kluwer Law & Business, 2016).

Leeuw, Franciscus L & Hans Schmeets, *Empirical legal research: a guidance book for lawyers, legislators and regulators* (Cheltenham, UK Northampton, MA: Edward Elgar Publishing, 2016).

MacCormick, Neil, *Legal Reasoning and Legal Theory* (Clarendon Press, 1994).

Macfarlane, Julie, *The National Self-Represented Litigants Project: Identifying and Meeting the Needs of Self-represented Litigants : Final Report* (National Self-Represented Litigants Project, 2013).

Mackaay, Ejan & Pierre Robillard, *Predicting judicial decisions: The nearest neighbour rule and visual representation of case patterns* (De Gruyter, 1974).

MacKenzie, I Scott, *Human-computer interaction: an empirical research perspective*, 1st ed (Amsterdam: Morgan Kaufmann is an imprint of Elsevier, 2013).

McIntyre, Joe, *The Judicial Function: Fundamental Principles of Contemporary Judging* (Singapore: Springer Singapore, 2019).

Mitchell, Tom M, *Machine Learning*, 1st ed (New York: McGraw-Hill Education, 1997).

Popple, James David, *Legal Expert Systems: The Inadequacy of a Rule-based Approach* (Oxford: Blackwell, 1987).

———, *SHYSTER: a pragmatic legal expert system* (Canberra: Dept. of Computer Science, Australian National University, 1993).

Rainey, Daniel, Ethan Katsh & Mohamed S Abdel Wahab, eds, *Online Dispute Resolution - Theory and Practice*, 2d ed (Eleven International, 2021).

Raz, Joseph, *The authority of law: essays on law and morality* (Oxford University Press on Demand, 2009).

Rhode, Deborah L, *Access to justice* (Oxford University Press, 2004).

Russell, Stuart J & Peter Norvig, *Artificial intelligence: a modern approach*, fourth ed, Pearson series in artificial intelligence (Hoboken: Pearson, 2021).

Russell, Stuart J, Peter Norvig & Ernest Davis, *Artificial intelligence: a modern approach*, 3rd ed, Prentice Hall series in artificial intelligence (Upper Saddle River: Prentice Hall, 2010).

Sandefur, Rebecca L, *Access to Justice: Classical Approaches and New Directions*, 1st ed, Sociology of Crime, Law and Deviance volume 12 (Bingley: Emerald JAI, 2009).

Sayn, Isabelle, *Le droit mis en barèmes ?* (Dalloz, 2014).

Susskind, Richard E, *Online courts and the future of justice* (Oxford: Oxford University Press, 2019).

Thulasiraman, K & M N S Swamy, *Graphs: theory and algorithms* (New York: Wiley, 1992).

A Tale of Cyberjustice: A Modern Approach to Technology in the Canadian Justice System (Montréal, Québec, Canada: Cyberjustice Laboratory, 2019).

Al-Abdulkarim, Latifa et al, “Angelic environment: demonstration” (2017) Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law (ICAAIL '17) 267–268.

———, “Angelic Environment: Support for the Construction of Legal KBS” (2015) Proceedings of the 15th International Conference on Artificial Intelligence and Law 3–12.

———, “Noise induced hearing loss: Building an application using the ANGELIC methodology” (2019) 10:1 *Argument & Computation* 5–22.

Articles

Al-Abdulkarim, Latifa, Katie Atkinson & Trevor Bench-Capon, “A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks” (2016) 24:1 *Artif Intell Law* 1–49.

———, “Accommodating change” (2016) 24:4 *Artif Intell Law* 409–427.

———, “ANGELIC Secrets: Bridging from Factors to Facts in US Trade Secrets” (2016) *Legal Knowledge and Information Systems* 113–118.

Alarie, Benjamin, Anthony Niblett & Albert H Yoon, “Using machine learning to predict outcomes in tax law” (2016) 58 *Can Bus LJ* 231.

Aletras, Nikolaos et al, “Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective” (2016) 2 *PeerJ Comput Sci* e93.

Aleven, Vincent, “Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment” (2003) 150:1 *Artificial Intelligence (AI and Law)* 183–237.

Allen, Layman E & C Rudy Engholm, “Normalized Legal Drafting and the Query Method” (1978) 29:4 *Journal of Legal Education* 380–412.

Allen, Ronald J & Alex Stein, “Evidence, Probability, and the Burden of Proof” (2013) 55 *Arizona Law Review* 557–602.

Alom, Md Zahangir et al, “The history began from alexnet: A comprehensive survey on deep learning approaches” (2018) arXiv preprint arXiv:180301164.

Ashley, Kevin D, “Case-Based Models of Legal Reasoning in a Civil Law Context. Invited paper” (2004) *International Congress of Comparative Cultures and Legal Systems of the Instituto de Investigaciones Juridicas, Universidad Nacional Autonoma de México.*

———, “Reasoning with cases and hypotheticals in HYPO” (1991) 34:6 *International Journal of Man-Machine Studies* 753–796.

Ashley, Kevin D & Stefanie Brüninghaus, “Automatically Classifying Case Texts and Predicting Outcomes” (2009) 17:2 *Artif Intell Law* 125–165.

———, “Automatically classifying case texts and predicting outcomes” (2009) 17:2 *Artif Intell Law* 125–165.

Ashley, Kevin D & Edwina L Rissland, “Law, learning and representation” (2003) 150:1 *Artificial Intelligence (AI and Law)* 17–58.

Ashley, Kevin D & Vern R Walker, “Toward constructing evidence-based legal arguments using legal decision documents and machine learning” (2013) *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law - ICAIL '13* 176.

Atkinson, Katie, Trevor Bench-Capon & Danushka Bollegala, “Explanation in AI and law: Past, present and future” (2020) 289 *Artificial Intelligence* 103387.

Barendrecht, Maurits, “Legal aid, accessible courts or legal information? Three access to justice strategies compared” (2011) 11:1 *Global Jurist*.

Barkai, John & Elizabeth Kent, “Let’s stop spreading rumors about settlement and litigation: A comparative study of settlement and litigation in Hawaii courts” (2014) 29 *Ohio St J on Disp Resol* 85.

Barredo Arrieta, Alejandro et al, “Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI” (2019) arXiv e-prints, online: <<http://adsabs.harvard.edu/abs/2019arXiv191010045B>>.

Bauer, Tobias et al, “# MeTooMaastricht: Building a chatbot to assist survivors of sexual harassment” (2020) *Machine Learning and Knowledge Discovery in Databases: International Workshops of ECML PKDD 2019, Würzburg, Germany, September 16–20, 2019, Proceedings, Part I* 503–521.

Becker, Daniel & Isabela Ferrari, “VICTOR, the Brazilian Supreme Court’s Artificial Intelligence: a beauty or a beast?” (2020) *SIFoCC Standing International Forum of Commercial Courts*, online: <<https://sifocc.org/app/uploads/2020/06/Victor-Beauty-or-the-Beast.pdf>>.

Beemer, Brandon A & Dawn G Gregg, “Advisory Systems to Support Decision Making” in Frada Burstein & Clyde W Holsapple, eds, *Handbook on Decision Support Systems 1: Basic Themes* International Handbooks Information System (Berlin, Heidelberg: Springer, 2008) 511.

Bellucci, Emilia, Arno R Lodder & John Zeleznikow, “Integrating artificial intelligence, argumentation and game theory to develop an online dispute resolution environment” (2004) 16th *IEEE International Conference on Tools with Artificial Intelligence* 749–754.

Bellucci, Emilia & John Zeleznikow, “Developing Negotiation Decision Support Systems that support mediators: a case study of the Family_Winner system” (2005) 13:2 *Artificial Intelligence and Law* 233–271.

Bench-Capon, T J M, “HYPO’S legacy: introduction to the virtual special issue” (2017) 25:2 *Artif Intell Law* 205–250.

Bench-Capon, T J M & M J Sergot, “Toward a Rule-Based Representation of Open Texture in Law” in Charles Walter, ed, *Computer Power and Legal Language: The Use of Computational Linguistics, Artificial Intelligence, and Expert Systems in the Law* (New York: Praeger, 1988) 39.

Bench-Capon, Trevor et al, “A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law” (2012) 20:3 *Artificial Intelligence and Law* 215–319.

Benyekhlef, Karim et al, “ICT-Driven Strategies for Reforming Access to Justice Mechanisms in Developing Countries” (2015) 325 *The World Bank Legal Review*.

———, “Online Consumer Dispute Resolution: a narrative around (and an example of) postmodern law*” (2016) *Lex Electronica* 32.

Benyekhlef, Karim & Fabien Gélinas, “Online Dispute Resolution” (2005) 10:2 *Lex Electronica*, online: <<https://www.lex-electronica.org/en/articles/vol10/num2/online-dispute-resolution/>>.

Benyekhlef, Karim & Jie Zhu, “Intelligence artificielle et justice: justice prédictive, conflits de basse intensité et données massives” (2018) 30 *Cahiers de propriété intellectuelle* 789–826.

Berman, Donald H & Carole D Hafner, “The potential of artificial intelligence to help solve the crisis in our legal system” (1989) 32:8 *Commun ACM* 928–938.

Bickel, Esmée A, M Van Dijk & Ellen Giebels, “Online legal advice and conflict support: a Dutch experience” (2015) Report, University of Twente.

Bilson, Beth, Brea Lowenberger & Graham Sharp, “Reducing the ‘Justice Gap’ Through Access to Legal Information: Establishing Access to Justice Entry Points at Public Libraries” (2017) 34:2 *Windsor Yearbook of Access to Justice/Recueil annuel de Windsor d’accès à la justice* 99–128.

Blei, David M, Andrew Y Ng & Michael I Jordan, “Latent dirichlet allocation” (2003) 3:Jan *Journal of machine Learning research* 993–1022.

Bommarito II, Michael & Daniel Martin Katz, “GPT Takes the Bar Exam” (2022) arXiv, online: <<http://arxiv.org/abs/2212.14402>>.

Branting, Karl et al, “A computational model of facilitation in online dispute resolution” (2022) *Artif Intell Law*, online: <<https://doi.org/10.1007/s10506-022-09318-7>>.

———, “Judges Are from Mars, Pro Se Litigants Are from Venus: Predicting Decisions from Lay Text” in Serena Villata, Jakub Harašta & Petr Křemen, eds, *Frontiers in Artificial Intelligence and Applications* (IOS Press, 2020).

Branting, Karl L et al, “Semi-Supervised Methods for Explainable Legal Prediction” (2019) *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law (ICAIL '19)* 22–31.

Branting, L Karl, “Advisory systems for pro se litigants” (2001) *Proceedings of the 8th international conference on Artificial intelligence and law* 139–146.

———, “Building explanations from rules and structured cases” (1991) 34:6 *International Journal of Man-Machine Studies (AI and Legal Reasoning. Part 1)* 797–837.

———, “Scalable and explainable legal prediction” (2021) 29:2 *Artif Intell Law* 213–238.

Breyer, Stephen, “The Federal Sentencing Guidelines and the Key Compromises Upon Which They Rest” (1988) 17 *Hofstra L Rev* 1–50.

Brown, Tom et al, “Language models are few-shot learners” (2020) 33 *Advances in neural information processing systems* 1877–1901.

Buchanan, Bruce G, “A (Very) Brief History of Artificial Intelligence” (2005) 26:4 *AI Magazine* 53–53.

Buchanan, Bruce G & Reid G Smith, “Fundamentals of Expert Systems” (1988) 3:1 *Annual Review of Computer Science* 23–58.

Buolamwini, Joy & Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification” (2018) *Conference on Fairness, Accountability and Transparency* 77–91.

Bush, Vannevar, “As we may think” (1945) 176:1 *The atlantic monthly* 101–108.

Bushway, Shawn D & Anne Morrison Piehl, “Judging Judicial Discretion: Legal Factors and Racial Discrimination in Sentencing” (2001) 35:4 *Law & Society Review* 733.

Cabral, James E et al, “Using Technology to Enhance Access to Justice” (2012) 26:1 *Harv JL & Tech* 241.

Carneiro, Davide et al, “Online dispute resolution: an artificial intelligence perspective” (2014) 41:2 *Artif Intell Rev* 211–240.

———, “Using case-based reasoning and principled negotiation to provide decision support for dispute resolution” (2013) 36:3 *Knowledge and Information Systems* 789–826.

Chang, Yun-chien & Daniel M Klerman, “Settlement Around the World: Settlement Rates in the Largest Economies” (2021) 21–8 USC CLASS Research Paper Series No CLASS21-8, USC Legal Studies Research Paper Series.

Chen, Daniel L & Jess Eigel, “Can machine learning help predict the outcome of asylum adjudications?” (2017) Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law 237–240.

Chen, Mark et al, “Evaluating Large Language Models Trained on Code” (2021), online: <<http://arxiv.org/abs/2107.03374>>.

Colarusso, David, “Machine-Assisted Issue Spotting for Self-Represented Litigant Portals” (2022).

Cotter, Adam & Laura Savage, “Gender-based violence and unwanted sexual behaviour in Canada, 2018: Initial findings from the Survey of Safety in Public and Private Spaces” (2019) Juristat: Canadian Centre for Justice Statistics 1–49.

Currie, Ab, “The legal problems of everyday life” in Rebecca L Sandefur, ed, *Sociology of Crime, Law and Deviance* (Emerald Group Publishing Limited, 2009) 1.

Dahan, Samuel et al, “Predicting Employment Notice Period with Machine Learning: Promises and Limitations” (2020) 65:4 *McGill Law Journal/Revue de droit de McGill* 711–753.

Dahlman, Christian & Anne Ruth Mackor, “Coherence and probability in legal evidence” (2019) 18:4 *Law, Probability and Risk* 275–294.

Dainow, Joseph, “The Civil Law and the Common Law: Some Points of Comparison” (1966) 15:3 *The American Journal of Comparative Law* 419–435.

Del Duca, Louis F, Colin Rule & Kathryn Rimpfel, “eBay’s De Facto Low Value High Volume Resolution Process: Lessons and Best Practices for ODR Systems Designers” (2014) 6:1 *Arbitration Law Review* 204–219.

Deng, Jia et al, “Imagenet: A large-scale hierarchical image database” (2009) 2009 IEEE conference on computer vision and pattern recognition 248–255.

Devlin, Jacob et al, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding” (2019) arXiv:181004805 [cs], online: <<http://arxiv.org/abs/1810.04805>>.

Domingos, Pedro, “A few useful things to know about machine learning” (2012) 55:10 Commun ACM 78–87.

Doshi-Velez, Finale et al, “Accountability of AI Under the Law: The Role of Explanation” (2019) arXiv:171101134 [cs, stat], online: <<http://arxiv.org/abs/1711.01134>>.

Dosovitskiy, Alexey et al, “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale” (2021) arXiv:201011929 [cs], online: <<http://arxiv.org/abs/2010.11929>>.

Duda, R & E Shortliffe, “Expert Systems Research” (1983) 220:4594 Science 261–268.

Dumouchel, Paul, “Intelligence, Artificial and Otherwise” (2019) 24:2 Forum Philosophicum 241–258.

Eisenberg, Theodore & Charlotte Lanvers, “What is the Settlement Rate and Why Should We Care?” (2009) 6:1 Journal of Empirical Legal Studies 111–146.

Epstein, Lee & Gary King, “The Rules of Inference” (2002) 69 University of Chicago Law Review 1–209.

Epstein, Lee & Andrew D Martin, “Quantitative Approaches to Empirical Legal Research” (2010) The Oxford Handbook of Empirical Legal Research, online: <<http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199542475.001.0001/oxfordhb-9780199542475-e-38>>.

Falakmasir, Mohammad Hassan & Kevin D Ashley, “Utilizing Vector Space Models for Identifying Legal Factors from Text” (2017) Proceedings of the International Conference on Legal Knowledge and Information Systems (JURIX).

Farrow, Trevor, “What is Access to Justice?” (2014) 51:3 Osgoode Hall Law Journal 957–988.

Ferro, Lisa et al, “Scalable Methods for Annotating Legal-Decision Corpora” (2019) Proceedings of the Natural Legal Language Processing Workshop 2019 12–20.

Frey, Carl Benedikt & Michael A Osborne, “The future of employment: How susceptible are jobs to computerisation?” (2017) 114 Technological Forecasting and Social Change 254–280.

Fuller, Lon L, “Positivism and Fidelity to Law--A Reply to Professor Hart” (1957) 71:4 Harv L Rev 630–672.

Garth, Bryant G & Mauro Cappelletti, “Access to Justice: The Newest Wave in the Worldwide Movement to Make Rights Effective” (1978) BUFFALO LAW REVIEW 113.

Gaudreau, Louis, Guillaume Hébert & Julia Posca, “Analyse du marché de l’immobilier et de la rentabilité du logement locatif” (2020) Institut de recherche et d’informations socioéconomiques 20.

Gill, T Grandon, “Early expert systems: Where are they now?” (1995) *MIS quarterly* 51–81.

Glenn, H Patrick, “Doin’ the Transsystemic: Legal Systems and Legal Traditions” (2005) 50 *McGill LJ* 863–898.

Goertzel, Ben, Matt Iklé & Jared Wigmore, “The Architecture of Human-Like General Intelligence” in Pei Wang & Ben Goertzel, eds, *Theoretical Foundations of Artificial General Intelligence* (Paris: Atlantis Press, 2012) 123.

Goldschmidt, Jona, “How are courts handling pro se litigants” (1998) 82 *Judicature* 13.

Green, Leslie, “Positivism, Realism, and Sources of Law” (2019), online: <<https://papers.ssrn.com/abstract=3443743>>.

Greenleaf, Graham, Andrew Mowbray & Philip Chung, “The Meaning of Free Access to Legal Information: A Twenty Year Evolution” (2013) 1:1 *J Open Access L* 1–68.

Gross, Samuel R & Kent D Syverud, “Getting to No: A Study of Settlement Negotiations and the Selection of Cases for Trial” (1991) 90:2 *Michigan Law Review* 319–393.

Gunning, David et al, “XAI—Explainable artificial intelligence” (2019) 4:37 *Science Robotics*.

Hannah-Moffat, Kelly, “Actuarial Sentencing: An ‘Unsettled’ Proposition” (2012) 30 *Justice Quarterly* 1–27.

Hart, H L A, “Positivism and the Separation of Law and Morals” (1957) 71 *Harv L Rev* 593–629.

He, Kaiming et al, “Deep residual learning for image recognition” (2016) *Proceedings of the IEEE conference on computer vision and pattern recognition* 770–778.

Holt, Edward M, “How to treat fools: Exploring the duties owed to pro se litigants in civil cases” (2001) 25 *J Legal Prof* 167.

Horgan, Terence & John Tienson, “Representations without Rules” (1989) 17:1 *Philosophical Topics* 147–174.

Hossain, MD Zakir et al, “A Comprehensive Survey of Deep Learning for Image Captioning” (2019) 51:6 *ACM Comput Surv* 118:1-118:36.

- Jobin, Anna, Marcello Ienca & Effy Vayena, “The global landscape of AI ethics guidelines” (2019) 1:9 *Nature Machine Intelligence* 389–399.
- Jumper, John et al, “Highly accurate protein structure prediction with AlphaFold” (2021) 596:7873 *Nature* 583–589.
- Kaelbling, L P, M L Littman & A W Moore, “Reinforcement Learning: A Survey” (1996) 4 *Journal of Artificial Intelligence Research* 237–285.
- Karray, Fakhreddine et al, “Human-Computer Interaction: Overview on State of the Art” (2008) 1:1 *International Journal on Smart Sensing and Intelligent Systems* 137–159.
- Katz, Daniel Martin et al, “Gpt-4 passes the bar exam” (2023) Available at SSRN 4389233.
- Katz, Daniel Martin, Michael J Bommarito II & Josh Blackman, “A general approach for predicting the behavior of the Supreme Court of the United States” (2017) 12:4 *PLOS ONE* e0174698.
- Kehl, Danielle Leah & Samuel Ari Kessler, “Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in Sentencing” (2017) Responsive Communities Initiative, Berkman Klein Center for Internet & Society, Harvard Law School, online: <<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041>>.
- Kistemaker, Laura, “Rechtwijzer and Uitelkaar.nl. Dutch Experiences with ODR for Divorce” (2021) 59:2 *Family Court Review* 232–243.
- Komárek, Jan, “Reasoning with Previous Decisions: Beyond the Doctrine of Precedent” (2013) 61:1 *American Journal of Comparative Law* 149–171.
- Krippendorff, Klaus, “A Dictionary of Cybernetics” (1986).
- Kritzer, Herbert M, “Empirical legal studies before 1940: a bibliographic essay” (2009) 6:4 *Journal of Empirical Legal Studies* 925–968.
- Krizhevsky, Alex, Ilya Sutskever & Geoffrey E Hinton, “ImageNet Classification with Deep Convolutional Neural Networks” in F Pereira et al, eds, *Advances in Neural Information Processing Systems 25* (Curran Associates, Inc., 2012) 1097.
- Lage-Freitas, André et al, “Predicting Brazilian Court Decisions” (2022) 8 *PeerJ Comput Sci* e904.
- Landsman, Stephan, “The growing challenge of pro se litigation” (2009) 13 *Lewis & Clark L Rev* 439.
- Lawlor, Reed C, “What Computers Can Do: Analysis and Prediction of Judicial Decisions” (1963) 49:4 *American Bar Association Journal* 337–344.

LeBel, Sonia & Éric Caire, “Stratégie d’intégration de l’intelligence artificielle dans l’administration publique 2021-2026” (2021) Gouvernement du Québec.

LeCun, Yann, Yoshua Bengio & Geoffrey Hinton, “Deep learning” (2015) 521:7553 Nature 436–444.

Leiter, Brian, “Legal Realism and Legal Positivism Reconsidered” (2001) 111:2 Ethics 278–301.

Levi, Edward H, “An Introduction to Legal Reasoning” (1948) 15:3 The University of Chicago Law Review 501–574.

Levy, Mitchell, “Empirical Patterns of Pro Se Litigation in Federal District Courts” The University of Chicago Law Review 49.

Licklider, J C R, “Man-Computer Symbiosis” (1960) HFE-1:1 IRE Transactions on Human Factors in Electronics 4–11.

Lindquist, Stefanie A, “7. Stare Decisis as Reciprocity Norm” in Charles Gardner Geyh, ed, *What’s Law Got to Do With It?* (Stanford University Press, 2020) 173.

Liu, Vivian & Lydia B Chilton, “Design Guidelines for Prompt Engineering Text-to-Image Generative Models” (2022) Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI ’22) 1–23.

Lodder, Arno R & Ernest M Thiessen, “The role of Artificial Intelligence in Online Dispute Resolution” (2004) Proceedings from UN forum on ODR, online: <<https://research.vu.nl/en/publications/the-role-of-artificial-intelligence-in-online-dispute-resolution>>.

Lord Justice Bingham, “The Discretion of the Judge” (1990) 5 Denning LJ 27–44.

Malini, N & M Pushpa, “Analysis on credit card fraud identification techniques based on KNN and outlier detection” (2017) 2017 Third International Conference on Advances in Electrical, Electronics, Information, Communication and Bio-Informatics (AEEICB) 255–258.

Marcus, Gary, “Deep learning: A critical appraisal” (2018) arXiv preprint arXiv:180100631.

Marmor, Andrei & Alexander Sarch, “The Nature of Law” in Edward N Zalta, ed, *The Stanford Encyclopedia of Philosophy*, fall 2019 ed (Metaphysics Research Lab, Stanford University, 2019).

McBaine, J P, “Burden of Proof: Degrees of Belief” (1944) 32:3 Calif L Rev 242–268.

McCarthy, John, “Programs with common sense” (1960) RLE and MIT computation center Cambridge, MA, USA.

———, “Some expert systems need common sense” (1984).

McCarty, L T & N S Sridharan, “The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations” (1981) 1 Proceedings of the 7th international joint conference on Artificial intelligence-Volume 246–253.

McCarty, L Thorne, “An implementation of Eisner v. Macomber” (1995) Proceedings of the fifth international conference on Artificial intelligence and law - ICAIL '95 276–286.

McCormack, Hon Bridget Mary, “Access to Justice Requires Plain Language” (2021) 100:2 Michigan Bar Journal 44–46.

Medvedeva, Masha, Michel Vols & Martijn Wieling, “Using machine learning to predict decisions of the European Court of Human Rights” (2020) 28:2 Artif Intell Law 237–266.

Miller, Tim, “Explanation in artificial intelligence: Insights from the social sciences” (2019) 267 Artificial Intelligence 1–38.

Mitee, Leesi Ebenezer, “The Right of Public Access to Legal Information: A Proposal for its Universal Recognition as a Human Right” (2017) 18:6 German Law Journal 1429–1496.

Morison, John & Adam Harkens, “Re-engineering justice? Robot judges, computerised courts and (semi) automated legal decision-making” (2019) 39:4 Legal Studies 618–635.

Mumford, Jack, Katie Atkinson & Trevor Bench-Capon, “Reasoning with Legal Cases: A Hybrid ADF-ML Approach” (2022) 362 Legal Knowledge and Information Systems (Frontiers in Artificial Intelligence and Applications) 93–102.

Nelson Ford, F, “Decision support systems and expert systems: A comparison” (1985) 8:1 Information & Management 21–26.

Newell, Allen & Herbert A Simon, “Computer science as empirical inquiry: symbols and search” (1976) 19:3 Commun ACM 113–126.

Noreau, Pierre, “Accès à la justice et démocratie en panne: constats, analyses et projections” (2010) Révolutionner la justice, Montréal, Thémis 13.

Northcutt, Curtis G, Anish Athalye & Jonas Mueller, “Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks” (2021) arXiv:210314749 [cs, stat], online: <<http://arxiv.org/abs/2103.14749>>.

Olson, Gary M & Judith S Olson, “Human-Computer Interaction: Psychological Aspects of the Human Use of Computing” (2003) 54:1 Annual Review of Psychology 491–516.

- O'Neill, James et al, "An Analysis of Topic Modelling for Legislative Texts" (2017) 2143 Proceedings of the Second Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 16th International Conference on Artificial Intelligence and Law (ICAIL 2017), London, UK, June 16, 2017 (CEUR Workshop Proceedings) , online: <<http://ceur-ws.org/Vol-2143/paper7.pdf>>.
- Oriola, Taiwo A, "The use of legal software by non-lawyers and the perils of unauthorised practice of law charges in the United States: a review of Jayson Reynoso decision" (2010) 18:3 Artif Intell Law 285–309.
- Ouyang, Long et al, "Training language models to follow instructions with human feedback" (2022) arXiv, online: <<http://arxiv.org/abs/2203.02155>>.
- Pan, Sinno Jialin & Qiang Yang, "A Survey on Transfer Learning" (2010) 22:10 IEEE Trans Knowl Data Eng 1345–1359.
- Paquin, Louis-Claude, François Blanchard & Claude Thomasset, "Loge–expert: from a legal expert system to an information system for non-lawyers" (1991) Proceedings of the 3rd international conference on Artificial intelligence and law (ICAIL '91) 254–259.
- Pasca, Alexandra, "Fabien Gélinas et al, Foundations of Civil Justice: Toward a Value-Based Framework for Reform, Cham, Springer, 2015" (2017) 30:1 rqi 149–152.
- , "Promising Evaluation Practices Guide: A Few Basic Tips" 16.
- Posner, Richard A, "Legal formalism, legal realism, and the interpretation of statutes and the constitution" (1986) 37 Case W Res L Rev 179.
- Qasse, Ilham, Shailesh Mishra & Mohammad Hamdaqa, "Chat2Code: Towards conversational concrete syntax for model specification and code generation, the case of smart contracts" (2021), online: <<http://arxiv.org/abs/2112.11101>>.
- Radford, Alec et al, "Robust speech recognition via large-scale weak supervision" (2022) arXiv preprint arXiv:221204356.
- Ramesh, Aditya et al, "Hierarchical Text-Conditional Image Generation with CLIP Latents" (2022), online: <<http://arxiv.org/abs/2204.06125>>.
- Rao, Qing & Jelena Frtunikj, "Deep learning for self-driving cars: chances and challenges" (2018) Proceedings of the 1st International Workshop on Software Engineering for AI in Autonomous Systems (SEFAIS '18) 35–38.
- Redmon, Joseph & Ali Farhadi, "Yolov3: An incremental improvement" (2018) arXiv preprint arXiv:180402767.

Ribeiro, Marco Tulio, Sameer Singh & Carlos Guestrin, ““ Why should i trust you?’ Explaining the predictions of any classifier” (2016) Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining 1135–1144.

Rissland, Edwina L, “Examples in Legal Reasoning: Legal Hypotheticals.” (1983) IJCAI 90–93.

Rissland, Edwina L & Kevin D Ashley, “HYPO: A Precedent-Based Legal Reasoner” (1987) Defense Technical Information Center, online:
<<http://www.dtic.mil/docs/citations/ADA249335>>.

Rissland, Edwina L & David B Skalak, “CABARET: rule interpretation in a hybrid architecture” (1991) 34:6 International Journal of Man-Machine Studies (AI and Legal Reasoning. Part 1) 839–887.

Salaün, Olivier et al, “Analysis and Multilabel Classification of Quebec Court Decisions in the Domain of Housing Law” (2020) Natural Language Processing and Information Systems (Lecture Notes in Computer Science) 135–143.

———, “Conditional Abstractive Summarization of Court Decisions for Laymen and Insights from Human Evaluation” (2022) Legal Knowledge and Information Systems 123–132.

———, “Why Do Tenants Sue Their Landlords? Answers from a Topic Model” (2022) Legal Knowledge and Information Systems 113–122.

Salter, Shannon, “Online dispute resolution and justice system integration: British Columbia’s Civil Resolution Tribunal” (2017) 34:1 Windsor Yearbook of Access to Justice 112–129.

Sambasivan, Nithya et al, “‘Everyone wants to do the model work, not the data work’: Data Cascades in High-Stakes AI” (2021) proceedings of the 2021 CHI Conference on Human Factors in Computing Systems 1–15.

Saravanan, M & B Ravindran, “Identification of Rhetorical Roles for Segmentation and Summarization of a Legal Judgment” (2010) 18:1 Artif Intell Law 45–76.

Satoh, Ken et al, “PROLEG: An Implementation of the Presupposed Ultimate Fact Theory of Japanese Civil Code by PROLOG Technology” (2011) New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science) 153–164.

Satoh, Ken, Laura Giordano & Matteo Baldoni, “Implementation of Choice of Jurisdiction and Law in Private International Law by PROLEG Meta-interpreter” (2021) Logic and Argumentation (Lecture Notes in Computer Science) 60–75.

- Satoh, Ken, Kazuko Takahashi & Tatsuki Kawasaki, “Interactive system for arranging issues based on PROLEG in civil litigation” in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law* (New York, NY, USA: Association for Computing Machinery, 2021) 273.
- Savage, Laura & Susan McDonald, “Experiences of serious problems or disputes in the Canadian provinces, 2021” (2022) *Juristat: Canadian Centre for Justice Statistics* 1–28.
- Savelka, Jaromir et al, “Can GPT-4 Support Analysis of Textual Data in Tasks Requiring Highly Specialized Domain Expertise?” (2023) 3441 *Proceedings of the 6th Workshop on Automated Semantic Analysis of Information in Legal Text (CEUR Workshop Proceedings)* 1–12.
- , “Explaining Legal Concepts with Augmented Large Language Models (GPT-4)” (2023), online: <<http://arxiv.org/abs/2306.09525>>.
- , “Lex Rosetta: transfer of predictive models across languages, jurisdictions, and legal domains” (2021) *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL '21)* 129–138.
- Šavelka, Jaromir & Kevin D Ashley, “Segmenting U.S. Court Decisions into Functional and Issue Specific Parts” (2018) 313 *Legal Knowledge and Information Systems (Frontiers in Artificial Intelligence and Applications)* 111.
- Selbst, Andrew D et al, “Fairness and abstraction in sociotechnical systems” (2019) *Proceedings of the conference on fairness, accountability, and transparency* 59–68.
- Selbst, Andrew D & Julia Powles, “Meaningful information and the right to explanation” (2017) 7:4 *International Data Privacy Law* 233–242.
- Semple, Noel, “The cost of seeking civil justice in Canada” (2015) 93 *Can B Rev* 639.
- Sergot, M J et al, “The British Nationality Act as a logic program” (1986) 29:5 *Commun ACM* 370–386.
- Shaikh, Rafe Athar, Tirath Prasad Sahu & Veena Anand, “Predicting Outcomes of Legal Cases based on Legal Factors using Classifiers” (2020) 167 *Procedia Computer Science (International Conference on Computational Intelligence and Data Science)* 2393–2402.
- Shelton, Dinah, “Hierarchy of norms and human rights: Of trumps and winners” (2002) 65 *Sask L Rev* 301.
- Shin, Donghee & Yong Jin Park, “Role of fairness, accountability, and transparency in algorithmic affordance” (2019) 98 *Computers in Human Behavior* 277–284.

- Shortliffe, Edward H et al, “Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system” (1975) 8:4 *Computers and Biomedical Research* 303–320.
- Siems, Mathias M, “The Taxonomy of Interdisciplinary Legal Research: Finding the Way out of the Desert” (2009) 7:1 *Journal of Commonwealth Law and Legal Education* 5–17.
- Smith, Linda F & Barry Stratford, “DIY in Family Law: A Case Study of a Brief Advice Clinic for Pro Se Litigants” (2012) 14:2 *JL & Fam Stud* 167–222.
- Smolensky, Paul, “Connectionist AI, symbolic AI, and the brain” (1987) 1:2 *Artificial Intelligence Review* 95–109.
- Spahn, Thomas E, “Is Your Artificial Intelligence Guilty of the Unauthorized Practice of Law” (2017) 4 *Rich JL & Tech* 1–47.
- Sprague, Ralph H, “A Framework for the Development of Decision Support Systems” (1980) 4:4 *MIS Quarterly* 1–26.
- Starr, Sonja B, “Evidence-Based Sentencing and the Scientific Rationalization of Discrimination” (2014) 66 *Stanford Law Review* 71.
- Stein, Peter G, “Roman Law, Common Law, and Civil Law Symposium: Relationships Among Roman Law, Common Law, and Civil Law” (1991) 66 *Tul L Rev* 1591–1604.
- Stranieri, Andrew et al, “A hybrid rule–neural approach for the automation of legal reasoning in the discretionary domain of family law in Australia” (1999) 7:2 *Artificial intelligence and Law* 153–183.
- Stranieri, Andrew & John Zeleznikow, “The Role of Open Texture and Stare Decisis in Data Mining Discretion” (1998) *JURIX* 1998 11.
- Șulea, Octavia-Maria et al, “Predicting the Law Area and Decisions of French Supreme Court Cases” (2017) *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017* 716–722.
- Sunstein, Cass R, “On Analogical Reasoning” (1993) 106:3 *Harvard Law Review* 741–791.
- Surdeanu, Mihai et al, “Risk analysis for intellectual property litigation” (2011) *Proceedings of the 13th International Conference on Artificial Intelligence and Law (ICAIL '11)* 116–120.
- Surden, Harry, “The ethics of artificial intelligence in law: Basic questions” (2020) Forthcoming chapter in *Oxford Handbook of Ethics of AI* 19–29.

- Suresh, Harini & John V Guttag, “A Framework for Understanding Unintended Consequences of Machine Learning” (2020) arXiv:190110002 [cs, stat], online: <<http://arxiv.org/abs/1901.10002>>.
- Susskind, R E, “Expert systems in law: out of the research laboratory and into the marketplace” (1987) Proceedings of the 1st international conference on Artificial intelligence and law (ICAIL ’87) 1–8.
- Susskind, Richard, “Pragmatism and purism in artificial intelligence and legal reasoning” (1989) 3:1 AI & Soc 28–38.
- Susskind, Richard E, “Expert Systems in Law: A Jurisprudential Approach to Artificial Intelligence and Legal Reasoning” (1986) 49:2 The Modern Law Review 168–194.
- Sviokla, John J, “An Examination of the Impact of Expert Systems on the Firm: The Case of XCON” (1990) 14:2 MIS Quarterly 127–140.
- Szita, István & András Lőrincz, “Learning Tetris Using the Noisy Cross-Entropy Method” 18 Neural Computation 2006.
- Tan, Jinzhe, Hannes Westermann & Karim Benyekhlef, “ChatGPT as an Artificial Lawyer?” (2023) 3435 Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice (CEUR Workshop Proceedings) , online: <<https://ceur-ws.org/Vol-3435/#short2>>.
- Tay Hui Yuan, Denise, “Access to Legal Information and Counsel: Prevention Is as Important as Cure” (2017) 35 Sing L Rev 47–65.
- Taylor, Veronica et al, “Introduction: Nature of the Japanese Legal System” (2008) 1 Business law in Japan 3–8.
- Temme, Merle, “Algorithms and Transparency in View of the New General Data Protection Regulation” (2017) 3:4 European Data Protection Law Review 473–485.
- Thomasset, Claude, François Blanchard & Louis-Claude Paquin, “Loge-expert: an illustration of different phases of the development of an expert system in Law” Université du Québec à Montréal.
- , “Loge-expert: Strategies to integrate legal knowledge modelization, non-expert user interface, and textual data base into the development of an expert system in law” (1992) 4:4 Expert Systems with Applications (Special Issue: Expert Systems and Law) 379–395.
- Thomasset, Claude & Louis-Claude Paquin, “Expert Systems in Law and the Representation of Legal Knowledge: Can we Isolate it from the Why and the Who?”

(1989) III International Congress: Logica, Informatica, Diritto: Expert Systems in Law 751–772.

Thompson, Darin, “Creating New Pathways to Justice Using Simple Artificial Intelligence and Online Dispute Resolution” (2015) 2 IJODR 4–53.

Thoppilan, Romal et al, “LaMDA: Language Models for Dialog Applications” (2022), online: <<http://arxiv.org/abs/2201.08239>>.

Vaswani, Ashish et al, “Attention is all you need” (2017) 30 Advances in neural information processing systems.

Vellino, Andre, “Artificial intelligence: The very idea: J. Haugeland, (MIT Press, Cambridge, MA, 1985); 287 pp.” (1986) 29 Artificial Intelligence 349–353.

Vermeys, Nicolas, “Privacy v. Transparency: How Remote Access to Court Records Forces Us to Re-examine Our Fundamental Values” in Karim Benyekhlef et al, eds, *eAccess to Justice* (University of Ottawa Press, 2016) 123.

Vermeys, Nicolas W & Karim Benyekhlef, “Best Practices in the Field of Cyberjustice” (2011) Seminar on Recent Trends and Good Practices in the Application of Electronic Technology to Judicial Processes (E-Justice).

Villata, Serena et al, “Thirty years of artificial intelligence and law: the third decade” (2022) *Artif Intell Law*, online: <<https://doi.org/10.1007/s10506-022-09327-6>>.

Vinyals, Oriol et al, “Grandmaster level in StarCraft II using multi-agent reinforcement learning” (2019) 575:7782 *Nature* 350–354.

Waldron, Jeremy, “The Rule of Law” in Edward N Zalta, ed, *The Stanford Encyclopedia of Philosophy*, spring 2020 ed (Metaphysics Research Lab, Stanford University, 2020).

Walker, Vern R, “A default-logic paradigm for legal fact-finding” (2007) 47:2 *Jurimetrics* 193–243.

———, “Automatic Classification of Rhetorical Roles for Sentences: Comparing Rule-Based Scripts with Machine Learning” (2019) Proceedings of the 3rd Workshop on Automated Semantic Analysis of Information in Legal Text.

———, “Evidence Types, Credibility Factors, and Patterns or Soft Rules for Weighing Conflicting Evidence: Argument Mining in the Context of Legal Rules Governing Evidence Assessment” (2018) Proceedings of the 5th Workshop on Argument Mining 68–78.

———, “Representing the Logic of Statutory Rules in the United States” in Michał Araszkiewicz & Krzysztof Pleszka, eds, *Logic in the Theory and Practice of Lawmaking* (Cham: Springer International Publishing, 2015) 357.

———, “Representing the use of rule-based presumptions in legal decision documents” (2014) 13:3–4 *Law, Probability and Risk* 259–275.

Walker, Vern R et al, “Semantic Types for Decomposing Evidence Assessment in Decisions on Veterans’ Disability Claims for PTSD” (2017) *Proceedings of the 2nd Workshop on Automated Semantic Analysis of Information in Legal Text* 10.

Waltl, Bernhard et al, “Classifying Legal Norms with Active Machine Learning.” (2017) *Proceedings of the International Conference on Legal Knowledge and Information Systems (JURIX)* 11–20.

Wang, Pei, “On defining artificial intelligence” (2019) 10:2 *Journal of Artificial General Intelligence* 1–37.

Waterman, D A & M A Peterson, “Models of Legal Decisionmaking - Research Design and Methods” (1981) RAND Corporation.

———, “Rule-Based Models of Legal Expertise” (1980) 1 *AAAI* 272–275.

Webley, Lisa, “Qualitative approaches to empirical legal research” (2010) *The Oxford handbook of empirical legal research* 926–950.

Westermann, Hannes, “Automating Reasoning with Previous Judicial Decisions” in Karim Benyekhlef, ed, *AI and Law: a Critical Overview* (2020) 189.

———, “Bridging the Gap: Mapping Layperson Narratives to Legal Issues with Language Models” (2023) 3441 *Proceedings of the 6th Workshop on Automated Semantic Analysis of Information in Legal Text (CEUR Workshop Proceedings)* 37–48.

———, “Computer-Assisted Creation of Boolean Search Rules for Text Classification in the Legal Domain” (2019) *Frontiers in Artificial Intelligence and Applications* 123–132.

———, “Data-Centric Machine Learning in the Legal Domain” (2022) arXiv:220106653 [cs], online: <<http://arxiv.org/abs/2201.06653>>.

———, “Evidence-Based Sentencing : Risks and Opportunities” (2020) 25:3 *Lex Electronica* 71–93.

———, “Sentence Embeddings and High-Speed Similarity Search for Fast Computer Assisted Annotation of Legal Documents” (2020) *Legal Knowledge and Information Systems* 164–173.

———, “Toward an Intelligent Tutoring System for Argument Mining in Legal Texts” (2022) *Legal Knowledge and Information Systems* 133–142.

———, “Using Factors to Predict and Analyze Landlord-Tenant Decisions to Increase Access to Justice” (2019) Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law (ICAIL '19) 133–142.

Westermann, Hannes, Jaromir Savelka & Karim Benyekhlef, “LLMediator: GPT-4 Assisted Online Dispute Resolution” (2023) 3435 Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice (CEUR Workshop Proceedings), online: <<https://ceur-ws.org/Vol-3435/#paper1>>.

———, “Paragraph Similarity Scoring and Fine-Tuned BERT for Legal Information Retrieval and Entailment” (2021) New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science) 269–285.

Xu, Huihui et al, “Using Argument Mining for Legal Text Summarization” (2020) Legal Knowledge and Information Systems 184–193.

Xu, Huihui, Jaromir Savelka & Kevin D Ashley, “Toward summarizing case decisions via extracting argument issues, reasons, and conclusions” (2021) Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL '21) 250–254.

Yablon, Charles M, “The Indeterminacy of the Law: Critical Legal Studies and the Problem of Legal Explanation Responsive Scholarship from outside the Movement” (1984) 6:4 Cardozo L Rev 917–946.

Yampolskiy, Roman, “AI-Complete, AI-Hard, or AI-Easy: Classification of Problems in Artificial Intelligence” (2012) The 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Cincinnati, OH, USA.

Yang, Yinfei et al, “Multilingual Universal Sentence Encoder for Semantic Retrieval” (2019) arXiv, online: <<http://arxiv.org/abs/1907.04307>>.

Yin, Yifei, Farhana Zulkernine & Samuel Dahan, “Determining Worker Type from Legal Text Data using Machine Learning” (2020) 2020 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCoM/CyberSciTech) 444–450.

Yudkowsky, Eliezer, “The AI alignment problem: why it is hard, and where to start” (2016) Symbolic Systems Distinguished Speaker.

Zadeh, L A, “The role of fuzzy logic in the management of uncertainty in expert systems” (1983) 11:1 Fuzzy Sets and Systems 199–227.

Zelezniak, John, “Building Decision Support Systems in Discretionary Legal Domains” (2000) 3 Int'l Rev L Computers & Tech 341–356.

———, “Can Artificial Intelligence and Online Dispute Resolution Enhance Efficiency and Effectiveness in Courts” (2016) 8:2 *International Journal for Court Administration* 30–45.

———, “Using Web-Based Legal Decision Support Systems to Improve Access to Justice” (2002) 11:1 *Info & Comm Tech L* 15–34.

Zhong, Haoxi et al, “Legal Judgment Prediction via Topological Learning” (2018) *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* 3540–3549.

Zhong, Linwu et al, “Automatic Summarization of Legal Decisions using Iterative Masking of Predictive Sentences” (2019) *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law* 163–172.

Reports & Online Sources

A3nm, *Tred-Gprime.svg* (2013).

Angwin, Julia et al, “Machine Bias”, (23 May 2016), online: *ProPublica* <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>.

Ataei, Mehdi et al, *Understanding Dataset Shift and Potential Remedies*, Technical Report, by Mehdi Ataei et al, Technical Report (The Vector Institute, 2021).

Beaver, H Dennis, “Could ChatGPT and AI Change Delivery of Legal Services?”, (27 January 2023), online: *Kiplinger.com* <<https://www.kiplinger.com/personal-finance/chatgpt-artificial-intelligence-and-legal-services>>.

Bennet, Drake, “ChatGPT Is an OK Law Student. Can It Be an OK Lawyer?”, (27 January 2023), online: *Bloomberg.com* <<https://www.bloomberg.com/news/newsletters/2023-01-27/chatgpt-can-help-with-test-exams-it-may-even-offer-legal-advice>>.

Benyekhlef, Karim & Nicolas Vermeys, “‘Low-Value, High-Volume’ Disputes: Defining the Indefinable”, (29 January 2014), online: *Slaw* <<https://www.slaw.ca/2014/01/29/low-value-high-volume-disputes-defining-the-indefinable/>>.

———, “The ‘Success’ of Online Dispute Resolution in Europe”, (18 June 2018), online: *Slaw* <<http://www.slaw.ca/2018/06/18/the-success-of-online-dispute-resolution-in-europe/>>.

Bonaque, Audrey, “JusticeBot : un nouvel outil utile et intelligent”, (26 July 2021), online: *Droit-inc* <<https://www.droit-inc.com/article29094-JusticeBot-un-nouvel-outil-utile-et-intelligent>>.

Branwen, Gwern, “GPT-3 Creative Fiction”, (19 June 2020), online: <https://www.gwern.net/GPT-3>.

Breuel, Cristiano, “ML Ops: Machine Learning as an Engineering Discipline”, (3 January 2020), online: *Medium* <https://towardsdatascience.com/ml-ops-machine-learning-as-an-engineering-discipline-b86ca4874a3f>.

Callipel, Valentin, “New platform for online resolution of pay equity disputes in workplaces ruled by federal reglementation!”, (13 January 2021), online: *Laboratoire de cyberjustice* <https://www.cyberjustice.ca/en/2021/01/13/nouvelle-plateforme-pour-regler-en-ligne-les-differends-relatifs-a-lequite-salariale-dans-les-milieux-de-travail-sous-reglementation-federale/>.

Carabin, François, “JusticeBot, un nouvel outil web pour aider les locataires et les propriétaires à connaître et à faire respecter leurs droits”, (20 July 2021), online: *Le Journal de Montréal* <https://www.24heures.ca/2021/07/20/nouvel-outil-en-droit-du-logement-la-jurisprudence-au-bout-des-doigts>.

Carpenter, Julia, “Google’s algorithm shows prestigious job ads to men, but not to women. Here’s why that should worry you.”, (6 July 2015), online: *Washington Post* <https://www.washingtonpost.com/news/the-intersect/wp/2015/07/06/googles-algorithm-shows-prestigious-job-ads-to-men-but-not-to-women-heres-why-that-should-worry-you/>.

Chalkidis, Ilias, Ion Androustopoulos & Nikolaos Aletras, *Neural Legal Judgment Prediction in English* (2019).

Collins, Eli & Zoubin Ghahramani, “LaMDA: our breakthrough conversation technology”, (18 May 2021), online: *Google* <https://blog.google/technology/ai/lamda/>.

Corbett-Davies, Sam et al, “A computer program used for bail and sentencing decisions was labeled biased against blacks. It’s actually not that clear.”, (17 October 2016), online: *Washington Post* <https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-publicas/>.

Curtis, Christopher, “Renters are getting burned in Quebec’s red-hot housing market”, (27 May 2021), online: *Ricochet* <https://ricochet.media/en/3664>.

Daly, Ciarán, “‘I’m Not A Robot’: Google’s Anti-Robot reCAPTCHA Trains Their Robots To See”, (25 October 2017), online: *AI Business* <https://aibusiness.com/recaptcha-trains-google-robots/>.

Dastin, Jeffrey, “Amazon scraps secret AI recruiting tool that showed bias against women”, (10 October 2018), online: *Reuters* <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>.

Deitz, Amanda, “Artificial Intelligence | Towards A Law Assisted By Algorithms”, (8 May 2022), online: *AI Magazine* <<https://ai-magazine.com/problems/artificial-intelligence-towards-a-law-assisted-by-algorithms/>>.

Diop, Mouhamadou-Lamine, “Explainable AI: The data scientists’ new challenge”, (14 June 2018), online: *Towards Data Science* <<https://towardsdatascience.com/explainable-ai-the-data-scientists-new-challenge-f7cac935a5b4>>.

Edwards, Benj, “With Stable Diffusion, you may never believe what you see online again”, (6 September 2022), online: *Ars Technica* <<https://arstechnica.com/information-technology/2022/09/with-stable-diffusion-you-may-never-believe-what-you-see-online-again/>>.

Farrow, Trevor, *Everyday Legal Problems and the Cost of Justice in Canada*, by Trevor Farrow, Cost of Justice Project (Canadian Forum on Civil Justice, 2014).

François Paradis, *Rapport annuel de gestion 2020-2021*, by François Paradis, Zotero (Office de la protection du consommateur).

Future of privacy forum, *The Privacy Expert’s Guide To Artificial Intelligence and Machine Learning*, by Future of privacy forum (2018).

Gagnon, Katia, “Plaintes à la Régie pour des logements moisiss: des années d’attente”, (12 March 2019), online: *La Presse* <<https://www.lapresse.ca/actualites/sante/201903/12/01-5217878-plaintes-a-la-regie-pour-des-logements-moisiss-des-annees-dattente.php>>.

Gagnon, Sophie, *Rapport Annuel 2019-2020*, by Sophie Gagnon (JuriPop).

Galang, Jessica, “Botler.AI’s new chatbot analyzed 300,000 court documents to help sexual harassment and assault survivors”, (6 December 2017), online: *BetaKit* <<https://betakit.com/botler-ais-new-chatbot-analyzed-300000-court-documents-to-help-sexual-harassment-and-assault-survivors/>>.

Gawande, Atul, “A Life-Saving Checklist”, (2 December 2007), online: *The New Yorker* <<https://www.newyorker.com/magazine/2007/12/10/the-checklist>>.

Government of Canada, Canadian Radio-television and Telecommunications Commission (CRTC), *Communications Monitoring Report 2019*, Reports, by Canadian Radio-television and Telecommunications Commission (CRTC) Government of Canada, Reports BC9-9E-PDF (The Canadian Radio-television and Telecommunications Commission (CRTC), 2020).

Grillo, Matt, “Some Quebec landlords frustrated at having to clean disaster apartments post moving day”, (7 July 2022), online: *CTV News - Montreal* <<https://montreal.ctvnews.ca/some-quebec-landlords-frustrated-at-having-to-clean-disaster-apartments-post-moving-day-1.5979162>>.

Guo, Yufeng, “The 7 Steps of Machine Learning”, (31 August 2017), online: *Towards Data Science* <<https://towardsdatascience.com/the-7-steps-of-machine-learning-2877d7e5548e>>.

Hagan, Margaret, “What is Legal Design?”, (26 January 2015), online: *Law By Design* <<https://lawbydesign.co/legal-design/>>.

Hao, Karen, “The AI technique that could imbue machines with the ability to reason”, (12 July 2019), online: *MIT Technology Review* <<https://www.technologyreview.com/2019/07/12/65579/the-next-ai-revolution-will-come-from-machine-learnings-most-underrated-form/>>.

Kanstrén, Teemu, “A Look at Precision, Recall, and F1-Score”, (11 September 2020), online: *Towards Data Science* <<https://towardsdatascience.com/a-look-at-precision-recall-and-f1-score-36b5fd0dd3ec>>.

Karpathy, Andrej, “Software 2.0”, (11 November 2017), online: *Medium* <<https://medium.com/@karpathy/software-2-0-a64152b37c35>>.

Kemp, Simon, “Digital 2022: Global Overview Report”, (26 January 2022), online: *DataReportal – Global Digital Insights* <<https://datareportal.com/reports/digital-2022-global-overview-report>>.

Knight, Will, “An AI Pioneer Wants His Algorithms to Understand the ‘Why’”, (10 August 2019), online: *Wired* <<https://www.wired.com/story/ai-pioneer-algorithms-understand-why/>>.

Law, Jonathan & Elizabeth A Martin, “public law” in *A Dictionary of Law*, 7th ed (Oxford University Press).

Lee, Dave, “Why Big Tech pays poor Kenyans to programme self-driving cars”, (3 November 2018), online: *BBC News* <<https://www.bbc.com/news/technology-46055595>>.

McLachlan, Thea, “The Montreal apartment building on the front lines of Canada’s rental crisis”, (10 May 2021), online: *Ricochet* <<https://ricochet.media/en/3642>>.

Mendicino, The Hon Marco EL, *2020 Annual Report to Parliament on Immigration*, by The Hon Marco EL Mendicino (Immigration, Refugees and Citizenship Canada, 2020).

Metz, Cade, “In Two Moves, AlphaGo and Lee Sedol Redefined the Future”, (16 March 2016), online: *Wired* <<https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/>>.

Metz, Rachel, “No, Google’s AI is not sentient”, (14 June 2022), online: *CNN Business* <<https://www.cnn.com/2022/06/13/tech/google-ai-not-sentient/index.html>>.

Mohney, Michelle, “How ChatGPT Could Impact Law and Legal Services Delivery”, (24 January 2023), online: *Northwestern Engineering* <<https://www.mccormick.northwestern.edu/news/articles/2023/01/how-chatgpt-could-impact-law-and-legal-services-delivery/>>.

Mostaque, Emad, “Stable Diffusion Public Release”, (22 August 2022), online: *Stability AI* <<https://stability.ai/blog/stable-diffusion-public-release>>.

Muehlhauser, Luke, “What is AGI?”, (11 August 2013), online: *Machine Intelligence Research Institute* <<https://intelligence.org/2013/08/11/what-is-agi/>>.

Nantasenamat, Chanin, “How to Build a Machine Learning Model”, (25 July 2020), online: *Towards Data Science* <<https://towardsdatascience.com/how-to-build-a-machine-learning-model-439ab8fb3fb1>>.

Nelson, Joseph, “Why should I do pre-processing and augmentation on my computer vision datasets?”, (26 January 2020), online: *Roboflow Blog* <<https://blog.roboflow.com/why-preprocess-augment/>>.

Nieri, Adam, “AI-written Scenario for Dungeons & Dragons Is Actually Quite Good”, (18 July 2020), online: *Mind Matters* <<https://mindmatters.ai/2020/07/ai-written-scenario-for-dungeons-dragons-is-actually-quite-good/>>.

Niiler, Eric, “Can AI Be a Fair Judge in Court? Estonia Thinks So”, (25 March 2019), online: *Wired* <<https://www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so/>>.

Oltean, Mihai, “Fruits 360”, (18 May 2020), online: *Kaggle* <<https://www.kaggle.com/datasets/moltean/fruits>>.

OpenAI, *GPT-4 Technical Report* (arXiv, 2023).

OpenAI & Ashley Philipiszyn, “GPT-3 Powers the Next Generation of Apps”, (25 March 2021), online: *OpenAI* <<https://openai.com/blog/gpt-3-apps/>>.

Parnham, Richard, Mario Sako & John Armour, *AI-assisted lawtech: its impact on law firms*, by Richard Parnham, Mario Sako & John Armour (Oxford: University of Oxford, 2021).

Piekniowski, Filip, “AI winter is well on its way”, (29 May 2018), online: *Piekniowski’s blog* <<https://blog.piekniowski.info/2018/05/28/ai-winter-is-well-on-its-way/>>.

Pomerleau, Julie, “Punaises de lit au logement”, (19 April 2018), online: *Blogue SOQUIJ - Actualités juridiques et judiciaires du Québec* <<https://blogue.soquij.qc.ca/2018/04/19/punaises-de-lit-logement/>>.

Rebourg, Muriel, “Le droit mis en barèmes ? / Sayn Isabelle (dir.)”, (6 February 2015), online: *Droit & Société* <<https://ds.hypotheses.org/770>>.

Reiser, Chloe, “Migrants bear the brunt of Canada’s worsening housing crisis”, (23 September 2021), online: *openDemocracy* <<https://www.opendemocracy.net/en/pandemic-border/migrants-bear-the-brunt-of-canadas-worsening-housing-crisis/>>.

Riera, Kevin, Anne-Laure Rousseau & Clément Baudelaire, “Doctor GPT-3: hype or reality?”, (27 October 2020), online: *Nabla* <<https://nabla.com/blog/gpt-3/>>.

Rozworski, Michal, “The roots of our housing crisis: Austerity, debt and extreme speculation”, (14 June 2019), online: *Policy Note* <<https://www.policynote.ca/the-roots-of-our-housing-crisis-austerity-debt-and-extreme-speculation/>>.

Sam Altman [@sama], *ChatGPT launched on wednesday. today it crossed 1 million users!* (2022).

Settles, Burr, *Active Learning Literature Survey*, Technical Report, by Burr Settles, minds.wisconsin.edu, Technical Report TR1648 (University of Wisconsin-Madison Department of Computer Sciences, 2009).

Shaywitz, David A, “‘The Alignment Problem’ Review: When Machines Miss the Point”, (25 October 2020), online: *Wall Street Journal* <<https://www.wsj.com/articles/the-alignment-problem-review-when-machines-miss-the-point-11603659140>>.

Silver, David & Demis Hassabis, “AlphaGo: Mastering the ancient game of Go with Machine Learning”, (27 January 2016), online: *Google AI Blog* <<http://ai.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.html>>.

Simard, Patrick, *Rapport annuel de gestion 2020-2021*, by Patrick Simard (Tribunal administratif du logement, 2021).

Soffer, Virginie, “Un robot répond à vos questions sur le droit du logement”, (26 July 2022), online: <<https://nouvelles.umontreal.ca/article/2021/07/26/un-robot-repond-a-vos-questions-sur-le-droit-du-logement/>>.

Thomas, Julie, “Agressions sexuelles : Lila, le chatbot qui aide les victimes”, (16 July 2019), online: *Revis* <<https://revisherault.org/agressions-sexuelles-lila-le-chatbot-qui-aide-les-victimes/>>.

Thomas, Katelyn, “Rental hell in Montreal: Caught in a renoviction nightmare”, (3 July 2021), online: *Montreal Gazette* <<https://montrealgazette.com/news/local-news/rental-hell-in-montreal-caught-in-a-renoviction-nightmare>>.

Tiku, Nitasha, “The Google engineer who thinks the company’s AI has come to life”, (17 June 2022), online: *Washington Post* <<https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>>.

Tromans, Richard, “Kira Systems Bags \$50m Investment, Largest Ever for a Legal AI Company”, (5 September 2018), online: *Artificial Lawyer* <<https://www.artificiallawyer.com/2018/09/05/kira-systems-bags-50m-investment-largest-ever-for-a-legal-ai-company/>>.

Vasdani, Tara, “Robot justice: China’s use of Internet courts”, (5 February 2020), online: *The Lawyer’s Daily* <<https://www.thelawyersdaily.ca/articles/17741/robot-justice-china-s-use-of-internet-courts>>.

Weinberg, Justin, “Philosophers On GPT-3 (updated with replies by GPT-3)”, (30 July 2020), online: *Daily Nous* <<https://dailynous.com/2020/07/30/philosophers-gpt-3/>>.

Wery, Etienne, “Le règlement des litiges liés au e-commerce se fait de plus en plus souvent en ligne”, (16 May 2018), online: *Droit & Technologies* <<https://www.droit-technologie.org/actualites/reglement-litiges-lies-e-commerce-se-de-plus-plus-souvent-ligne/>>.

Westermann, Hannes, “An Interactive Guide To The Blockchain, Part I-V”, (May 2020), online: *Cyberjustice Laboratory* <<https://www.cyberjustice.ca/mots-cles/blockchain/>>.

Wiso, *Simple neural network* (2008).

Yuan, Tianyu, *Structured relation – Increasing decision efficiency through collaborative decision creation* (Groningen, Netherlands, 2018).

Zhang, Daniel et al, *The AI Index 2022 Annual Report*, by Daniel Zhang et al (AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University).

A fifty-year look at divorces in Canada, 1970 to 2020, The Daily 11-001–X (Statistics Canada, 2022).

“À propos du tribunal spécialisé”, (14 September 2022), online: *Gouvernement du Québec* <<https://www.quebec.ca/justice-et-etat-civil/systeme-judiciaire/processus-judiciaire/tribunal-specialise-violence-sexuelle-violence-conjugale/a-propos>>.

“A2J Author”, online: <<https://www.a2jauthor.org/>>.

“About us”, online: *Juripop* <<https://juripop.org/en/about-us/>>.

“About Us”, online: *Éducaloi* <<https://educaloi.qc.ca/en/about/>>.

“Access to the Official Journal - EUR-Lex”, online: <<https://eur-lex.europa.eu/oj/direct-access.html>>.

“Amicable divorce”, online: *Ministère de la Justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/amicable-divorce/>>.

Brottsoffermyndighetens Referatsamling 2016 (Brottsoffermyndigheten, 2016).

“Call-in Legal Clinic”, online: *Jeune Barreau de Montréal* <<https://ajbm.qc.ca/en/public-services/annual-legal-helpline/>>.

Canada’s Tax System: What’s so Wrong and Why it Matters (Chartered Professional Accountants of Canada, 2018).

“Centres de justice de proximité”, online: <<https://www.justicedeproximite.qc.ca/en/>>.

“Clinique juridique du Mile End | Mile End Legal Clinic”, online: <<http://justicemontreal.org/>>.

“Common Voice by Mozilla”, online: <<https://commonvoice.mozilla.org/>>.

“Communiqué de presse - Lancement du JusticeBot”, (21 July 2021), online: *Laboratoire de cyberjustice* <<http://web.archive.org/web/20210721123006/https://www.cyberjustice.ca/2021/07/20/communiquede-presse-lancement-du-justicebot/>>.

“Conciliation process”, online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/conciliation-between-lessor-and-lessee/conciliation-process>>.

Contre la violence sexuelle, la violence conjugale et Rebâtir la confiance, Stratégie Gouvernementale Intégrée 2022-2027, Stratégie Gouvernementale Intégrée 2022-2027 (Gouvernement du Québec).

“DALL·E 2”, online: *OpenAI* <<https://openai.com/dall-e-2/>>.

“Dealing with a debt collector”, (5 August 2022), online: *Financial Consumer Agency of Canada* <<https://www.canada.ca/en/financial-consumer-agency/services/debt/collection-agency.html>>.

“Declaration on Free Access to Law”, online: <<http://www.falm.info/declaration/>>.

“Deploying Transformers on the Apple Neural Engine”, online: *Apple Machine Learning Research* <<https://machinelearning.apple.com/research/neural-engine-transformers>>.

Distinction entre donner un avis ou une opinion juridique et donner une information juridique (Barreau du Québec, 2013).

“Divorce by Agreement”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/divorce-by-agreement/>>.

“Docassemble”, online: *Docassemble* <<http://docassemble.org/>>.

“Düsseldorfer Tabelle”, (1 January 2019), online: *Oberlandesgericht Düsseldorf* <http://www.olg-duesseldorf.nrw.de/infos/Duesseldorfer_Tabelle/index.php>.

“Éducaloi - Your starting point for legal information”, online: *Éducaloi* <<https://educaloi.qc.ca/en/>>.

“Eligibility criteria (Housing Application)”, online: *Office municipal d’habitation de Montréal* <<https://www.omhm.qc.ca/en/submit-application/eligibility-criteria>>.

Entre fraude et spéculation - Enquêtes sur les reprises et évictions de logements (Comité de logement de la Petite Patrie, 2020).

European ethical Charter on the use of Artificial Intelligence in judicial systems and their environment (Strasbourg: Council of Europe European Commission for the efficiency of justice (CEPEJ), 2018).

“Filing an application”, online: *Ministère de la justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/amicable-divorce/application-for-divorce-based-on-a-draft-agreement/filing-an-application>>.

“Gender Based Violence in Canada”, (1 June 2022), online: *Canadian Women’s Foundation* <<https://canadianwomen.org/the-facts/gender-based-violence/>>.

“Gender-based Violence”, online: *UNHCR* <<https://www.unhcr.org/gender-based-violence.html>>.

“Government Licenses” in *West’s Encyclopedia of American Law*, 2d ed (The Gale Group, 2008).

“Hearings at the Tribunal administratif du logement”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/hearings-at-the-tribunal-administratif-du-logement/>>.

“Hearsay” in *Wex* (LII / Legal Information Institute, Cornell University).

“Home”, online: *CLEO (Community Legal Education Ontario / Éducation juridique communautaire Ontario)* <<https://www.cleo.on.ca/en>>.

“How to apply for divorce?”, online: *Ministère de la Justice* <<https://www.justice.gouv.qc.ca/en/couples-and-families/separation-and-divorce/marriage-or-civil-union/divorce/how-to-apply-for-divorce/>>.

“I Need a Lawyer - Eligibility”, online: *Juripop* <<https://juripop.org/en/i-need-a-lawyer-eligibility/>>.

“Income Assistance”, (24 August 2019), online: *Employment and Social Development Canada* <<https://www.canada.ca/en/employment-social-development/services/benefits/income-assistance.html>>.

“Individual income tax return statistics for the 2022 tax-filing season”, (24 January 2023), online: *Canada Revenue Agency* <<https://www.canada.ca/en/revenue-agency/corporate/about-canada-revenue-agency-cra/individual-income-tax-return-statistics.html>>.

“Introduction to smart contracts”, (23 January 2023), online: *ethereum.org* <<https://ethereum.org/en/smart-contracts/>>.

“JuridiQC”, online: *JuridiQC* <<https://juridiqc.gouv.qc.ca/>>.

Justice For All - Final Report (New York: Center on International Cooperation: The Task Force on Justice, 2019).

“JusticeBot – Interactive legal information tool”, (12 August 2021), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/justicebot-interactive-legal-information-tool>>.

“JusticeBot répond maintenant à vos questions sur les « rénovictions » !”, (17 September 2021), online: *Laboratoire de cyberjustice* <<https://web.archive.org/web/20210917130831/https://www.cyberjustice.ca/2021/09/16/justicebot-repond-maintenant-a-vos-questions-sur-les-renovictions/>>.

“JusticeBot, une application au service des locataires et propriétaires québécois!”, (16 November 2021), online: *Université de Montréal* <<https://droit.umontreal.ca/en/faculty/communications/news-details/news/detail/News/justicebot-une-application-au-service-des-locataires-et-proprietaires-quebecois/>>.

Lancement de la plateforme JusticeBot : Entrevue avec Karim Benyekhlef (Radio Canada, 2021).

“Legal Aid - Am I financially eligible?”, online: *Commissions des services juridiques* <<https://www.csj.qc.ca/commission-des-services-juridiques/aide-juridique/volet-gratuit-aj/en>>.

“Legal Aid - What is legal aid?”, online: *Comissions des services juridiques* <<https://www.csj.qc.ca/commission-des-services-juridiques/aide-juridique/Quest-ce-que-aide-juridique/en>>.

Legal Clinics in Canada: Exploring Service Delivery and Legal Outcomes Among Vulnerable Populations in the Context of COVID-19, J2-542/2022E-PDF (Department of Justice Canada, 2021).

Les comités logement et associations de locataires de l'île de Montréal (Regroupement des comités logement et associations de locataires du Québec, 2019).

Les loyers explosent - Enquête sur le prix des logements à louer au Québec (Regroupement des comités logement et associations de locataires du Québec, 2021).

Les punaises de lit : État de situation à Montréal (2019).

L'habitation en bref 2020 (Société d'habitation du Québec, 2020).

“L’implantation et l’utilisation des outils d’intelligence artificielle dans l’administration publique”, online: *Laboratoire de cyberjustice* <<https://www.cyberjustice.ca/2022/06/21/limplantation-et-lutilisation-des-outils-dintelligence-artificielle-dans-ladministration-publique/>>.

“microsoft/resnet-50 · Hugging Face”, online: <<https://huggingface.co/microsoft/resnet-50>>.

“Mobile Legal Clinic”, online: *cji-mlc* <<https://www.cji-mlc.org>>.

“New guide for plain language in law”, (31 January 2019), online: *The Access to Justice Lab* <<https://a2jlab.org/new-guide-for-plain-language-in-law/>>.

“Noise”, (18 December 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/the-dwelling/noise>>.

“Notification of an application to the other party”, (4 January 2016), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/filing-an-application/notification-of-an-application-to-the-other-party>>.

“Number of internet users in Canada 2022”, online: *Statista* <<https://www.statista.com/statistics/243808/number-of-internet-users-in-canada/>>.

“NVIVO - Qualitative Data Analysis Software for Researchers”, online: <<https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software/home>>.

“OpenAI Five Defeats Dota 2 World Champions”, (15 April 2019), online: *OpenAI* <<https://openai.com/blog/openai-five-defeats-dota-2-world-champions/>>.

“Operational instructions and guidelines — Immigration, Refugees and Citizenship Canada”, (5 December 2017), online: *Immigration, Refugees and Citizenship Canada (IRCC)* <<https://www.canada.ca/en/immigration-refugees-citizenship/corporate/publications-manuals/operational-bulletins-manuals.html>>.

“Our Mission”, online: *Justice Pro Bono* <<https://justiceprobono.ca/en/aboutus/our-mission/>>.

“Our Team”, online: *Éducaloi* <<https://educaloi.qc.ca/en/our-team/>>.

“PARLe”, (6 December 2018), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/parle-3/nos-solutions-logicielles/parle-2/>>.

“PARLe – Commission des normes, de l’équité, de la santé et de la sécurité du travail”, (7 April 2020), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/2020/04/07/the-parle-project-social-mediation-with-the-cnesst/>>.

“PARLe – Ontario Condominium Authority Tribunal”, (7 April 2020), online: *Laboratoire de cyberjustice* <<https://cyberjustice.openum.ca/2020/04/07/the-parle-project-the-condominium-authority-tribunal/>>.

“Plain language – essential for real access to justice”, (18 July 2017), online: *Provincial Court of British Columbia* <<https://www.provincialcourt.bc.ca/enews/enews-18-07-2017>>.

“Police Ethics”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/police-ethics/>>.

“probative value” in *Wex* (LII / Legal Information Institute, 2020).

“Procedures for filing an application”, (4 January 2016), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/filing-an-application/procedures-for-filing-an-application>>.

Project Genesis Annual Report 2020-2021 (Project Genesis, 2021).

“Qualitative Data Analysis Software, Mixed Methods Research Tool”, online: *Provalis Research* <<https://provalisresearch.com/products/qualitative-data-analysis-software/>>.

“Quasar Framework - Build high-performance VueJS user interfaces in record time”, online: *Quasar Framework* <<https://quasar.dev/>>.

“Qu’est-ce que PARLe?”, (1 June 2022), online: *Office de la protection du consommateur* <<https://www.opc.gouv.qc.ca/a-propos/parle/description/>>.

“Questions fréquentes”, online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/fr/questions-frequentes>>.

Rapport annuel 2020 (Office municipal d'habitation de Montréal, 2021).

Rapport annuel de gestion 2021 (Commission des normes, de l'équité, de la santé et de la sécurité du travail).

“Réalizations de l'Office en 2020-2021”, (1 March 2022), online: *Office de la protection du consommateur*

<<https://web.archive.org/web/20220302062245/https://www.opc.gouv.qc.ca/a-propos/office/realisations-20-21/>>.

“Renting”, online: *Éducaloi* <<https://educaloi.qc.ca/en/categories/renting/>>.

Report of The Summit on the Use of Technology to Expand Access to Justice (Legal Services Corporation, 2013).

“Responsibilities of Landlords”, online: *Éducaloi*

<<https://educaloi.qc.ca/en/capsules/responsibilities-of-landlords/>>.

“Rights and obligations of the lessor”, (20 November 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/being-a-lessor/rights-and-obligations-of-the-lessor>>.

“Sexual Assault And Harassment in Canada”, (22 November 2022), online: *Canadian Women's Foundation* <<https://canadianwomen.org/the-facts/sexual-assault-harassment/>>.

“Small claims”, (27 January 2023), online: *Gouvernement du Québec*

<<https://www.quebec.ca/en/justice-and-civil-status/small-claims>>.

“Spot A.I. tool for documenting harassment and discrimination receives major updates”, (20 October 2020), online: *Yukon Human Rights Commission*

<<https://yukonhumanrights.ca/2020/10/spot-a-i-tool-for-documenting-harassment-and-discrimination-receives-major-updates/>>.

spotify/annoy (Spotify, 2022).

“Subproject 8 - Tools for Self-Represented Litigants”, online: *Projet AJC | ACT Project*

<<https://www.ajcact.org/en/organisation/chantier-8-outils-dautorepresentation-pour-plaideurs-citoyens/>>.

“Teaching AI to translate 100s of spoken and written languages in real time”, (23 February 2022), online: *Meta AI* <<https://ai.facebook.com/blog/teaching-ai-to-translate-100s-of-spoken-and-written-languages-in-real-time/>>.

The Privacy Expert's Guide To Artificial Intelligence and Machine Learning (Future of Privacy forum, 2018).

“The Tribunal administratif du logement (TAL or rental board)”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/the-tribunal-administratif-du-logement-rental-board/>>.

Transforming our world: the 2030 Agenda for Sustainable Development, A/RES/70/1 (United Nations General Assembly).

Travel between Canada and other countries, December 2019, The Daily (Statistics Canada, 2020).

“TurboTax® Official Site: File Taxes Online, Tax Filing Made Easy”, online: <<https://turbotax.intuit.com/>>.

“Types of housing”, online: *Office municipal d’habitation de Montréal* <<https://www.omhm.qc.ca/en/about-us/types-housing>>.

“Understand Neighbour Law”, (October 2017), online: *Dial-A-Law* <<https://dialalaw.peopleslawschool.ca/neighbour-law/>>.

“Universal Sentence Encoder Multilingual”, online: <<https://tfhub.dev/google/universal-sentence-encoder-multilingual/3>>.

“Verify the authenticity of an electronic edition of the Official Journal - EUR-Lex”, online: <<https://eur-lex.europa.eu/content/help/oj/authenticity-eOJ.html>>.

“Vue.js - The Progressive JavaScript Framework”, online: *Vue.js* <<https://vuejs.org/>>.

“What is a lease?”, (18 December 2015), online: *Tribunal administratif du logement* <<https://www.tal.gouv.qc.ca/en/signing-a-lease/what-is-a-lease>>.

“What is Access to Justice?”, online: *Alberta Civil Liberties Research Centre* <<http://www.aclrc.com/what-is-access-to-justice>>.

“What is gender-based violence?”, (31 March 2021), online: *Government of Canada* <<https://women-gender-equality.canada.ca/en/gender-based-violence-knowledge-centre/about-gender-based-violence.html#women>>.

“What is the CCJM?”, online: *Legal Aid Montreal | Laval* <<https://www.aidejuridiquedemontreal.ca/en/legal-aid-montreal-laval/>>.

“Workplace Protections in Quebec”, online: *Éducaloi* <<https://educaloi.qc.ca/en/capsules/workplace-protections-in-quebec/>>.