

Université de Montréal

**Towards Simulating the Emergence of Environmentally Responsible
Behavior Among Natural Resource Users: An Integration of Complex
Systems Theory, Machine Learning and Geographic Information Science**

par

Saeed Harati Asl

Département de géographie
Faculté des arts et des sciences

Thèse présentée en vue de l'obtention du grade de
Philosophiæ Doctor (Ph.D.)
en géographie

20 décembre 2021

© Saeed Harati Asl, 2021

Université de Montréal

Faculté des études supérieures et postdoctorales

Cette thèse intitulée

A hybrid agent-based approach to simulation of emergence of environmentally
responsible behavior among natural resource users

présentée par

Saeed Harati Asl

A été évaluée par un jury composé des personnes suivantes :

Oliver Sonnentag (Université de Montréal)

(président-rapporteur)

Liliana Perez (Université de Montréal)

(directrice de recherche)

Roberto Molowny-Horas (CREAF, Universitat Autònoma de Barcelona)

(codirecteur de recherche)

Tyler Bonnell (University of Lethbridge)

(membre du jury)

Élise Filotas (Université TÉLUQ)

(examinatrice externe)

Résumé

La gouvernance pour le développement durable implique de nombreux défis multidimensionnels, notamment en raison de la complexité des systèmes socio-écologiques (SSE). L'un des aspects de la gouvernance dans ce domaine est la protection des ressources écologiques dont dépendent les sociétés, ce qui devient particulièrement difficile dans une perturbation écologique. Souvent, les gouvernements ont des capacités limitées pour s'attaquer seuls à ces problèmes. Par conséquent, pour les gouvernements, il est idéal que leurs sociétés prennent des mesures et adoptent des comportements spécifiques qui contribuent à l'objectif de protection des ressources. Cependant, il existe des obstacles à la réalisation de cette situation idéale dans laquelle les sociétés coopèrent avec les gouvernements pour la protection des ressources écologiques. D'une part, les individus sont souvent motivés par leur intérêt personnel plutôt que par altruisme, et ils sont susceptibles de refuser l'appel à l'action du gouvernement si une telle action est coûteuse. Par conséquent, les individus manquent la motivation de soutenir le gouvernement dans des tâches difficiles. D'un autre côté, l'expérience montre que les gouvernements ne peuvent pas toujours faire respecter leur autorité et obliger leurs sociétés à adopter un comportement respectueux de l'environnement si leurs sociétés ne sont pas prêtes. L'expérience montre également que les efforts du gouvernement pour encourager les gens en leur offrant des motivations financières peuvent fonctionner aussi longtemps que les motivations sont accordées, mais pas nécessairement après. En d'autres termes, l'application des réglementations et des motivations financières n'est pas toujours couronnée de succès, voire pas du tout. Du point de vue d'un gouvernement ayant pour objectif de protéger une ressource écologique, il est idéal que la société coopère avec le gouvernement en adoptant volontairement le comportement écologiquement responsable prévu par le gouvernement, sans que le gouvernement ait besoin de recourir à la force ou à des motivations financières. Il existe un défi différent qui ajoute encore une autre dimension à la complexité de ces situations : souvent le gouvernement ne dispose pas d'informations suffisantes sur les conséquences possibles de son intervention envisagée dans le système écologique. L'objectif du gouvernement est bien sûr de protéger la ressource écologique, mais étant donné l'incertitude et la complexité inhérentes aux systèmes écologiques, le gouvernement ne peut pas savoir ce qu'il advient de la ressource même si la société coopère pleinement avec le gouvernement et exécute le comportement écoresponsable demandé. Par conséquent, ces défis de gouvernance sont soumis à des

complexités sociales et écologiques. En ce sens, la situation idéale pour le gouvernement est que la société coopère avec lui et exécute volontairement l'action que le gouvernement prescrit, et cette action aboutit finalement à la protection de la ressource écologique. La description ci-dessus soulève la question de la possibilité d'une telle situation idéale.

Dans cette thèse, j'ai abordé la question ci-dessus dans le contexte d'un cas d'infestation d'insectes forestiers. J'ai analysé un cas où une forêt est attaquée par des infestations et une entité gouvernante essaie d'obtenir l'aide des utilisateurs de la forêt pour contrôler la propagation des infestations. Comme expliqué, ce problème a deux dimensions, une qui implique des interactions entre le gouvernement et les individus – dans ce cas, les utilisateurs des ressources – et une autre qui implique une complexité écologique. Afin d'aborder ce problème complexe de SSE, je l'ai d'abord divisé en ses composantes sociales et écologiques pour étudier chacune d'elles séparément. Ce faisant, j'ai fait des hypothèses simplificatrices et des abstractions quand nécessaire, et j'ai construit des modèles qui représentent les composantes sociales et écologiques du problème. J'ai adopté une approche de modélisation car dans de tels problèmes, l'apprentissage par essais et erreurs peut conduire à des effets imprévus qui sont peut-être destructeurs et irréversibles. Dans de tels problèmes, où l'expérimentation peut être coûteuse et les conditions et les paramètres d'expérimentation ne peuvent pas toujours être reproduits avec précision, la modélisation est une approche alternative qui offre une opportunité d'en apprendre davantage sur les systèmes et leurs complexités. Après avoir développé les modèles sociaux et écologiques, je les ai couplés pour construire un modèle du SSE. Ensuite, j'ai effectué des tests en utilisant le modèle SSE couplé.

L'une des approches possibles du problème de la promotion d'un comportement respectueux de l'environnement repose sur l'idée que les individus sont motivés par l'intérêt personnel et que l'un des aspects de l'intérêt personnel est la réputation. Les individus ont le désir d'avoir une bonne réputation dans leurs sociétés, et ce désir peut contribuer à leur motivation pour un comportement respectueux de l'environnement. Les individus se soucient également de la visibilité de leurs actions. Autrement dit, sachant que leurs actions sont observées par la société, les individus ont tendance à considérer leurs choix avant d'adopter un comportement. Avec ces considérations, dans cette thèse, j'ai proposé et analysé un mécanisme dans lequel le gouvernement récompense le comportement responsable par la reconnaissance, et

les individus dans la société évaluent la valeur d'être reconnus comme des membres responsables de la société dans leurs actions. J'ai appliqué ce mécanisme dans la construction du modèle social de cette étude et utilisé le modèle pour en savoir plus sur les complexités qui découlent de ce mécanisme. Plus précisément, j'ai évalué si ce mécanisme conduit potentiellement à l'émergence d'un nouveau comportement dans la société.

Une approche courante de l'étude des problèmes écologiques qui surviennent à de grandes échelles spatiales est la modélisation des changements des terres. Dans l'analyse de la composante écologique du problème, j'ai utilisé un modèle de changement des terres qui est construit sur des données d'observation de la propagation d'une infestation d'insectes forestiers. Avec ce modèle, qui s'appuie fortement sur des techniques d'apprentissage automatique supervisé, j'ai réalisé des simulations et acquis des connaissances sur la dynamique spatiale de la perturbation écologique dans la ressource forestière. J'ai notamment travaillé sur la calibration et la validation de ce modèle avec des données empiriques pour m'assurer qu'il produit des simulations réalistes, et pour comprendre les limites du modèle. Cela impliquait de comparer les résultats de simulation avec des données de référence, et dans le cas de cette étude doctorale, il s'agissait de méthodes d'évaluation de modèles par comparaison de cartes.

Dans la dernière partie de ce travail doctoral, j'ai couplé les modèles sociaux et écologiques ensemble, et construit un modèle d'un système plus complexe dans lequel les deux modèles interagissent de manière itérative : le modèle écologique simule une perturbation, et le modèle social apporte des modifications à la paysage du modèle écologique. Les perturbations futures dans le modèle écologique dépendent en partie de la propagation antérieure des perturbations et en partie de la couverture paysagère du modèle, qui est modifiée par le modèle social. D'autre part, les actions menées par le modèle social dépendent en partie de la dynamique interne du modèle social et en partie de l'état de santé de la ressource écologique. Cet ensemble d'interactions entrelacées fournit un laboratoire virtuel dans lequel j'ai testé divers scénarios pour en savoir plus sur la dynamique du SSE simulé.

Ce travail interdisciplinaire a abouti à plusieurs découvertes importantes sur le contexte du SSE décrit ci-dessus. Tout d'abord, il a montré que le mécanisme de reconnaissance proposé est influent dans l'émergence de nouveaux comportements dans la société. En particulier, avec un ensemble approprié de décisions, le gouvernement est capable, en mesure d'utiliser le mécanisme

de reconnaissance, de promouvoir un comportement respectueux de l'environnement et éventuellement protéger une ressource écologique. Cependant, il a également été démontré que sans une prise de décision appropriée par le gouvernement, son intervention peut entraîner des effets négatifs. Un autre point important qui a été trouvé était qu'avec suffisamment de temps avant les perturbations écologiques, le gouvernement peut utiliser le mécanisme de reconnaissance pour préparer la société à une action respectueuse de l'environnement, de sorte que lorsque le besoin s'en fait sentir, la société se joint rapidement au gouvernement pour protéger sa ressource écologique. En ce sens, cette étude a fourni un aperçu de la façon dont une société peut évoluer avant qu'elle ne soit prête à participer à l'action environnementale, ou à accepter de nouvelles réglementations environnementales.

En plus de la conclusion ci-dessus, l'approche démontrée dans cette étude est utile dans plusieurs questions connexes. Notamment, le modèle développé peut être utilisé comme un outil qui permet d'apprendre les conséquences possibles d'une série des interventions dans un écosystème. Cet outil est un ajustement simplifié du modèle SSE ci-dessus dans un scénario hypothétique dans lequel les utilisateurs effectuent toujours la tâche souhaitée par le gouvernement. En d'autres termes, il s'agit d'une supposition d'imposition du comportement environnemental prévu par le gouvernement. Différentes actions d'intervention peuvent être simulées et testées avec cet outil. De plus, le modèle SSE couplé de cette étude peut être utilisé dans d'autres contextes écologiques, en branchant le modèle écologique respectif et en le couplant au modèle social de cette étude. De même, d'autres modèles sociaux représentant une variété de phénomènes sociaux peuvent être branchés dans ce modèle de SSE couplé. Dans tous les cas, cette étude sert d'exemple qui montre comment les modèles peuvent être utilisés pour fournir des informations sur le SSE. La gouvernance pour le développement durable peut bénéficier de modèles empiriques ainsi que conceptuels afin d'en apprendre davantage sur les systèmes complexes, sur leur dynamique et sur les conséquences possibles d'une intervention dans ceux-ci.

Mots-clés : Système socio-écologique, gouvernance, comportement écoresponsable, systèmes complexes, modèle basé sur les agents, modèle spatial, évaluation de modèle, apprentissage automatique, apprentissage par renforcement

Abstract

Governance for sustainable development involves many multidisciplinary challenges, especially due to the complexities of social-ecological systems (SES). One of the aspects of governance in this domain is the protection of ecological resources that societies rely on, which becomes particularly challenging in times of ecological disturbance. Often, governments have limited capability in addressing such problems alone. Therefore, for governments, it is ideal that their societies take action and perform specific behaviors that contribute to the objective of resource protection. However, there are obstacles to the realization of that ideal situation in which societies cooperate with governments towards protection of ecological resources. For one thing, individuals are often driven by self-interest rather than altruism, and they are likely to refuse the government's call for action if such action is costly. Therefore, individuals may not have the motivation to support the government in difficult tasks. On the other hand, experience shows that governments cannot always enforce their authority and oblige their societies to do environmentally responsible behavior if their societies are not ready. Experience also shows that government efforts to motivate people by offering them financial incentives may work for as long as the incentives are given, but not necessarily afterwards. In other words, enforcement of regulations and financial incentives are not always successful, if possible at all. From the viewpoint of a government with an objective to protect an ecological resource, it is ideal that the society cooperates with the government by performing the government's intended environmentally responsible behavior voluntarily, without the need for the government to resort to force or financial incentives. There exists a different challenge that adds yet another dimension to the complexity of these situations: often the government does not have sufficient information about the possible consequences of its intended intervention in the ecological system. The objective of the government is of course to save the ecological resource, but given the uncertainty and complexity that is inherent in ecological systems, the government may not know what happens to the resource even if the society fully cooperates with the government and performs the requested environmentally responsible behavior. Therefore, these governance challenges are subject to social as well as ecological complexities. In this sense, the ideal situation for the government is that the society cooperates with it and voluntarily performs the action that the government prescribes, and that action eventually results in the protection of the

ecological resource. The above description gives rise to the question of possibility of such ideal situation.

In this thesis I addressed the above question in the context of a case of forest insect infestations. I analyzed a case where a forest is attacked by infestations and a governing entity tries to elicit help from forest users to control the spread of infestations. As explained, this problem has two dimensions, one that involves interactions between the government and individuals – in this case, resource users – and another that involves ecological complexity. In order to approach this complex SES problem, I first broke it into its social and ecological components to study each of them separate from the other. In doing so I made simplifying assumptions and abstractions where necessary, and I built models that represent the social and the ecological components of the problem. I took a modelling approach because in such problems, learning by trial-and-error may lead to unanticipated effects that are possibly destructive and irreversible. In such problems, where experimentation may be costly and experiment conditions and settings may not be reproducible precisely, modelling is an alternative approach that offers an opportunity to learn about the systems and their complexities. Once the social and the ecological models were developed, I coupled them together to build a model of the SES. Then I performed tests using the coupled SES model.

One of the possible approaches to the problem of promotion of environmentally responsible behavior relies on the assumptions that individuals are driven by self-interest, and that one of the aspects of self-interest is reputation. Individuals have a desire for good reputation in their societies, and this desire may contribute to their motivation for environmentally responsible behavior. Individuals also care about the visibility of their actions. That is, knowing that their actions are observed by the society, individuals tend to consider their choices before performing a behavior. With these considerations, in this thesis I proposed and analyzed a mechanism in which the government rewards responsible behavior with recognition, and individuals in the society assess the value of being recognized as responsible members of the society in their actions. I applied this mechanism in building the social model of this study and used the model to learn about the complexities that arise from this mechanism. Specifically, I assessed whether this mechanism potentially leads to the emergence of a new behavior in the society.

A common approach to the study of ecological problems that occur in large spatial scales is land change modelling. In the analysis of the ecological component of the problem, I used a land change model that is built on observation data of the spread of a forest insect infestation. With this model, which relies heavily on supervised machine learning techniques, I made simulations and gained insight about the spatial dynamics of the ecological disturbance in the forest resource. I particularly worked on calibrating and validating this model with empirical data to make sure that it produces realistic simulations, and to understand the model's limits. This required comparing simulation results with reference data, and in the case of this doctoral study it involved methods of model assessment by map comparison.

In the final part of this doctoral work I coupled the social and the ecological models together, and built a model of a more complex system in which the two models interact iteratively: the ecological model simulates a disturbance, and the social model makes alterations to the landscape of the ecological model. The future disturbances in the ecological model depend partly on previous spread of disturbances and partly on the model's landscape cover, which is changed by the social model. On the other hand, the actions taken by the social model depend partly on the internal dynamics of the social model and partly on the state of health of the ecological resource. This intertwined set of interactions provides a virtual laboratory in which I tested various scenarios to learn about the dynamics of the simulated SES.

This interdisciplinary work resulted in several important findings about the context of the SES described above. First and foremost, it showed that the proposed recognition mechanism is influential in the emergence of new behavior in the society. Particularly, with an appropriate set of decisions, the government is able to use the recognition mechanism to promote environmentally responsible behavior and eventually save an ecological resource. However, it was also shown that without appropriate decision making by the government, its intervention efforts may result in adverse effects. Another important matter that was found was that given enough time prior to ecological disturbance, the government can use the recognition mechanism to prepare the society for environmentally responsible action, such that when the need arises, the society quickly joins the government in protecting their ecological resource. In this sense, this study provided an insight about how a society may evolve before it is ready to participate in environmental action, or to accept new environmental regulations.

In addition to the above finding, the approach demonstrated in this study is useful in several related issues. Notably, the developed model can be used as a tool that allows to learn about the possible consequences of interventions in an ecosystem through a series of steps. This tool is a simplified adjustment of the above SES model in a hypothetical scenario in which users always perform the task that the government desires. In other words, this is an assumption of enforcement of the intended environmental behavior by the government. Various intervention actions can be simulated and tested with this tool. Moreover, the coupled SES model of this study can be used in other ecological contexts, by plugging in the respective ecological model and coupling with the social model of this study. Similarly, other social models representing a variety of social phenomena can be plugged in this coupled SES model. In any case, this study serves as an example that shows how models can be used to provide insight about SES. Governance for sustainable development can benefit from empirical as well as conceptual models in order to learn about complex systems, their dynamics, and the possible consequences of intervening in them.

Keywords: Social-ecological system, governance, environmentally responsible behavior, complex systems, Agent Based Model, spatial model, model assessment, Machine Learning, Reinforcement Learning

Table of contents

Résumé.....	3
Abstract.....	7
Table of contents.....	11
List of tables.....	16
List of figures.....	18
List of symbols and abbreviations	23
Symbols.....	23
Abbreviations.....	23
Acknowledgements.....	26
Introduction.....	28
Presentation of the problem	28
Sustainable development: a very brief overview	28
A problem in sustainable development.....	30
Conceptual and methodological background.....	33
Complex systems	33
Social-Ecological Systems.....	36
Social norms	40
Complex Systems Modelling.....	41
Model assessment	45
Machine Learning.....	49
Research setting, questions and objectives	55
Included publications	59
References.....	60
Chapter 1.....	70
Presentation of the article.....	70
Promoting the Emergence of Behavior Norms in a Principal–Agent Problem—An Agent-Based Modeling Approach Using Reinforcement Learning	71
Abstract.....	71
1.1. Introduction.....	72

1.2. Materials and Methods.....	76
1.2.1. Overview, Design Principles, Details	77
1.2.2. Model Parameters	84
1.2.3. Simulation Experiments.....	85
1.2.4. Implementation	86
1.3. Results.....	86
1.3.1. Simulations without RL	86
1.3.2. Simulations with RL	89
1.4. Discussion	97
1.5. Conclusions.....	102
Appendix 1.A. Model Parameters.....	104
Appendix 1.B. Parameter Combination Codes	104
Appendix 1.C. Flowcharts of Selected Algorithms	105
Appendix 1.D. Post-publication remarks.....	106
References.....	109
Chapter 2.....	115
Presentation of the article.....	115
Integrating Neighborhood Effect and Supervised Machine Learning Techniques to Model and Simulate Forest Insect Outbreaks in British Columbia, Canada.....	116
Abstract	116
2.1. Introduction.....	117
2.2. Materials and Methods.....	119
2.2.1. Study Area	119
2.2.2. Pine Mortality Dataset	120
2.2.3. Predictor Variables	121
2.2.4. Approaches to Model and Simulate Land Cover Changes	124
2.2.5. Model Calibration and Validation	126
2.2.6. Software	128
2.3. Results.....	129
2.4. Discussion	134
2.4.1. Model Validation	134

2.4.2. Model Predictions	135
2.4.3. Limitations of the Study	136
2.5. Conclusions.....	136
Appendix 2.A. Calculation of Threshold Value on Cumulative <i>Pinus contorta</i> Mortality Data	138
Appendix 2.B. Plots of Average Mortality vs. Predictors	143
Appendix 2.C. Graphic Representation of the Four Different Neighborhood Types Implemented	145
Appendix 2.D. Model Parameterization	145
Appendix 2.E. Code and Data Availability	147
Appendix 2.F. Post-publication remarks	147
References.....	149
Chapter 3.....	155
Presentation of the article.....	155
Validating models of one-way land change: an example case of forest insect disturbance	156
Abstract.....	156
3.1. Introduction.....	157
3.1.1 Validation of land change models	157
3.1.2 Scope and objectives.....	159
3.1.3 Case study	163
3.2 Methods.....	164
3.2.1 Distance analysis and proximity suitability map	164
3.2.2 Total Operating Characteristic curves	165
3.2.3 Three-map comparison at multiple resolutions	166
3.2.4 Four-map comparison and analysis of components of change in partitioned study area	168
3.2.5 Data	169
3.2.6 Software tools	170
3.3 Results.....	170
3.4 Discussion.....	175
3.4.1 Insights about the case of study	175

3.4.2. Limitations and implications of the study.....	179
3.5 Conclusions.....	181
Appendix 3.A. Post-publication remarks.....	182
References.....	184
Chapter 4.....	190
Presentation of the article.....	190
Emergence of environmentally responsible behavior through a recognition mechanism – insights from a conceptual social-ecological model.....	191
Abstract.....	191
4.1. Introduction.....	192
4.1.1. A governance problem in sustainable development	192
4.1.2. Background from multiple disciplines.....	193
4.1.3. Setting, questions and objectives	195
4.2. Methods.....	198
4.2.1. Overview.....	198
4.2.2. Design concepts	203
4.2.3. Details	208
4.2.4. Simulation scenarios	212
4.3. Results.....	214
4.4. Discussion	219
4.4.1. Insights about the case of study	219
4.4.2. Insights about governance of SES	222
4.4.3. Limitations of the study and prospects for future works	223
4.5. Conclusion	224
References.....	225
Conclusions.....	230
Novelty and findings.....	230
Chapter 1.....	230
Chapter 2.....	231
Chapter 3.....	233
Chapter 4.....	234

Remarks	236
Challenges and opportunities for future research	239
Computational challenges.....	239
Other challenges and opportunities.....	239
Final words.....	241
References.....	242

List of tables

Table S1.1. Model parameters.....	104
Table S1.2. Codes and combinations of problem parameters. These codes correspond to column numbers of the heatmap figures in the text. μ and σ are the mean and standard deviation of User agent thresholds, respectively. n is the number of user agents. γ is the future discount rate.....	104
Table S1.3. Codes and combinations of algorithm settings. These codes correspond to row numbers of the heatmap figures in the text. Algorithms are Double Expected SARSA (DXS), Double Q-Learning (DQ), Double SARSA (DS), Expected SARSA (ES), Q-Learning (Q), and SARSA (S). ε is the exploration rate. α is the learning rate.....	105
Table 2.1. List of predictor variables used for modeling. The Acronym column identifies the corresponding variable in Equations (2.2a) and (2.2b). The Time column indicates whether the variable is calculated at $t1p$ or $t1$ (see text for an explanation).....	122
Table 2.2. Comparison of hits, misses, and false alarms for different algorithms of simulation of change in study area from 2008 to 2014. GLM stands for Generalized Linear regression Model.	131
Table 2.3. Estimates of cumulative percentages of pine volume killed in MPB attacks by 2014 for three algorithms.....	132
Table 2.4. Estimates of cumulative percentages of pine volume killed in MPB attacks by 2020 for three algorithms.....	133
Table S2.1. Parameterization statistics for the GLM1 model.	145
Table S2.2. Parameterization statistics for the GLM2 model.	146
Table S2.3. RF model variable ranks and relative importance scores rounded to two decimal digits.....	146
Table 3.1. Definitions.....	159
Table 3.2. Percentiles of shortest Manhattan distances to initial points of spread.....	171
Table 4.1. Model entities and their state variables.....	200
Table 4.2. Model parameters.....	212
Table 4.3. Mean and standard deviation of the final proportions of healthy, infested and cut areas in simulations. Each scenario was run 50 times, except for Enforce, which run once.	217

Table 4.4. Wilcoxon signed rank test statistic and p-value for comparison of remaining healthy forest proportions of scenarios with BAU and with Enforce (n=50)..... 219

Table 4.5. Wilcoxon signed rank test statistic and p-value for comparison of remaining healthy forest proportions between Random and Suggest scenarios (n=50)..... 219

List of figures

Figure I.1. Schematic overview of Ostrom’s SES framework (Ostrom, 2009). The dashed line indicates the border of the SES.	38
Figure I.2. Von Neumann and Moore neighborhoods (red) about a central cell (gray) in a rectangular grid	42
Figure I.3. An example ABM in abstraction. Dotted line represents system boundary. Each agent may be in one-way or two-way interaction with other agents in the system, or with the system’s environment.	43
Figure I.4. Hypothetical example of supervised classification. Each box represents a class of objects. Correct labels for classes are given; the algorithm is expected to identify to which class each new object belongs.	52
Figure I.5. Hypothetical example of clustering as an unsupervised learning problem. The algorithm is expected to identify groups of similar object based only on the attributes of the objects	54
Figure I.6. Hypothetical example of Reinforcement Learning with a game of dart. Numbered crosses indicate three throws at the target. After the first throw, performance assessment was “above target” and the policy was “aim lower”. After the second throw, performance assessment was “too far below target” and the policy was “aim higher”.....	55
Figure 1.1. Interactions between three classes of agents in one time-step. Class names are highlighted. Variable names are shown in italics.	78
Figure 1.2. Histograms of overall scores of episodes in reference dataset, ordered by number and mean decision threshold of user agents. Each histogram summarizes 6,553,600 data points. Each data point is an episode of 17 time-steps. The score of each episode is its mean reward per time-step. Rewards are real numbers between -1 and 0 , and they are calculated based on the ratio of user agents that cooperate with the governing agent in each time-step.....	88
Figure 1.3. Heatmap of mean scores of simulations. Each row is a unique combination of algorithm settings. Each column is a unique combination of problem parameters. Rows and columns are ordered using hierarchical clustering, as shown in their respective dendrograms. Each pixel represents the mean score of 50 simulations with its respective row and column settings. Simulation scores are mean rewards per time-step.....	90

Figure 1.4. Heatmap of ranks of simulations. Each row is a unique combination of algorithm settings. Each column is a unique combination of problem parameters. For each column, the row with the highest mean score of simulations is given the highest rank. The order of rows and columns in the ranks heatmap is the same as that of the mean scores heatmap.	92
Figure 1.5. Rewards versus time-steps for Double SARSA and Double Expected SARSA algorithms. For each of the two algorithms, 810 curves are shown. Each curve represents a parameter setting for its respective algorithm. For each parameter setting, 50 simulations were run, their mean score was plotted at each time-step, and a line segment was drawn between the score points of consecutive time steps to produce a curve.	94
Figure 1.6. Histograms of scores at the final time step for RL algorithms Double SARSA and Double Expected SARSA as well as a random baseline. Each histogram summarizes 40,500 data points.....	95
Figure 1.7. Boxplots of scores at the final time step for RL algorithms Double SARSA and Double Expected SARSA as well as a random baseline, classified by mean value of user agents' decision threshold. Each boxplot shows quartile ranges of scores of 13,500 data points.	96
Figure S1.1. Double SARSA flowchart.....	106
Figure S1.2. Double Expected SARSA flowchart.....	106
Figure S1.3. Simulation ranks heatmap with zones associated to the learning algorithms.	109
Figure 2.1. The Province of British Columbia in Canada is mainly dominated by a Lodgepole pine (<i>Pinus contorta</i>) forest, which by 2014 had been decimated by almost 50% due to mountain pine beetle (MPB) infestation.	120
Figure 2.2. Schematic representation of the integrative approach to parametrize and calibrate our model to simulate forest insect disturbances.	128
Figure 2.3. Comparison of model simulations with reference observations of changes in the study area from 2008 to 2014. (a) Mountain pine beetle outbreak change in British Columbia (BC) recorded by satellite imagery and aerial surveys; mountain pine beetle outbreak change simulated with algorithms: (b) random forests (RF), (c) binomial regression (GLM1), (d) binomial regression with parabolic elevation (GLM2).....	129
Figure 2.4. Validation analysis by comparison of simulated and observed changes in study area from 2008 to 2014. White color indicates no data. Algorithms: (a) RF, (b) GLM1, (c) GLM2. Hits are correct simulations of change. False alarm errors are persistence simulated as change.	

Miss errors are change simulated as persistence. Correct rejections are correct simulations of persistence. Infestations prior to the study are excluded from the study area.	131
Figure 2.5. Predictions of changes in study area from 2014 to 2020. Algorithms: (a) random forest (RF); (b) binomial regression (GLM1); (c) binomial regression with parabolic elevation (GLM2).	133
Figure S2.1. Three sample curves of cumulative (left panels) and annual (right panels) infestation percentages.	138
Figure S2.2. Offset-corrected cumulative infestation data points.	140
Figure S2.3. Mean observed and offset-corrected cumulative infestation (red dots), normal cumulative distribution function It (CDF, in blue, scaled to 0–100), whose standard deviation is explained in text, and its derivative $I't$, i.e., the normal probability density function (in green).	141
Figure S2.4. Cumulative distribution function (red) of the normal distribution and its 2nd derivative (blue). The location of the maximum of the 2nd derivative and the corresponding ordinate is indicated with a dashed line.	143
Figure S2.5. Plots of average mortality vs. predictors in three transition periods: 2001–2002 (circles), 2007–2008 (triangles), and 2013–2014 (plus signs).	144
Figure S2.6. The four neighborhood types implemented in this model are (a) no-weight: neighborhood effect is uniformly distributed; (b) linear: neighborhood effect decreases linearly with increasing distance; (c) inverse-distance: neighborhood effect decreases inversely proportional to distance; and (d) squared-inverse-distance: neighborhood effect decreases inversely proportional to distance squared.	145
Figure 3.1. Outline of a contingency table.	159
Figure 3.2. Implications of reversibility and irreversibility in interpretation of assessment data: (a) in a one-way process, a change of state during the study period can be identified by comparing reference data at the beginning and end of the period; (b) in a reversible process, change of state may happen during the study period without a trace in the reference data at the beginning and end of the period.	161
Figure 3.3. A hypothetical TOC curve. In this plot, hits, misses, false alarms, and correct rejections are measures of area. Abundance is the area of reference change.	166

Figure 3.4. A hypothetical example showing reduction of allocation error in coarsening of resolution: (a) two fine-resolution pixels contain a miss and a false alarm; (b) when aggregated in the same coarse pixel, they indicate a hit and a correct rejection. ‘A’ and ‘B’ are states..... 167

Figure 3.5. Flowchart of the proposed validation method including definition of a proximity baseline and use of the baseline in assessment of simulations. 169

Figure 3.6. Total Operating Characteristic curves for two simulations and a proximity model. Hits and false alarms are pixel counts. Areas Under Curve are 0.79 for Logistic Regression, 0.83 for Random Forest, and 0.86 for Proximity..... 172

Figure 3.7. Maps of infestation suitability from 2008 to 2014 for two simulations and a proximity model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008..... 173

Figure 3.8. Comparison of maps of reference in 2008, reference in 2014, and prediction in 2014 for two simulations and a baseline model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008..... 173

Figure 3.9. Multiple-resolution analysis of components of agreement and disagreement between maps of reference in 2008, reference in 2014, and prediction in 2014..... 174

Figure 3.10. Partition of the study area into *near* and *far* strata using the baseline model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008..... 174

Figure 3.11. Components of change in near and far strata of the study area..... 175

Figure S3.1. Hypothetical proximity suitability map and baseline. (a) Reference (b) Proximity suitability with Manhattan distances from initial point of spread (c) Baseline. Initial point of spread is marked in gray, reference change in red, suitability in shades of blue, and baseline with diagonal stripes. Baseline marks the same number of cells as reference change. 183

Figure S3.2. Critical threshold selection with TOC curve. The vertical line drawn down from the upper left corner of the parallelogram meets the TOC curve at the critical threshold point (shown with a star)..... 183

Figure S3.3. Hypothetical four-map comparison. The four maps of the first row are converted to three maps of change on the second row, then overlaid to identify components of agreement and disagreement on the third row..... 184

Figure 4.1. Study area of forest insect infestations simulation in British Columbia, Canada. Locations of active infestations at the beginning of the simulation period are marked as ‘Insects at start’. Locations that were infested before the start of simulations are marked as ‘Previously infested’ 201

Figure 4.2. The flip-flop mechanism. Each of the two models requires input from the other model. The models communicate with each other via inbox directories. Arrows show direction of data transfer. Numbers beside arrows show the order of operations. 210

Figure 4.3. The coupled model. In each time-step, the ecological model begins with reading previous infestation maps and ends with writing a newly simulated infestation map; the social model begins with waiting for a message from the ecological model and ends with sending a message to the ecological model..... 210

Figure 4.4. Maps of mean remaining infestation after the final time-step of simulations. The Enforce scenario was run once, and all other scenarios were run 50 times. For each scenario, 'High' infestation in a cell means the presence of infestation in the cell in all runs of that scenario; 'Low' infestation means the absence of infestation in the cell in all runs of that scenario. 215

Figure 4.5. Mean cooperation ratio over time for Suggest and Random scenarios with 0, 10 and 20 preparation steps. Each scenario was run 50 times..... 216

Figure 4.6. Proportions of healthy, infested and cut areas in BAU and Enforce baseline scenarios..... 216

Figure 4.7. Proportions of healthy, infested and cut areas in Random baseline and Suggest scenarios. Each plot represents 50 runs. 217

List of symbols and abbreviations

Symbols

n	Number of user agents
μ	Mean decision threshold of user agents
σ	Standard deviation of decision thresholds of user agents
γ	Future discounting rate in a Reinforcement Learning algorithm
ε	Exploration rate in a Reinforcement Learning algorithm
α	Learning rate in a Reinforcement Learning algorithm
e	Elevation
r	Surface ruggedness
a	Aspect (cosine)
s	Slope
Z_{iden}	Distance function with identical intensity
Z_{lin}	Distance function with linear decline
Z_{inv}	Distance function with inverse linear decline
Z_{squ}	Distance function with inverse squared decline

Abbreviations

ABM	Agent-Based Model
AUC	Area Under Curve
BAU	Business As Usual
BC	British Columbia
CA	Cellular Automata
GIS	Geographical Information System
GLM	Generalized Linear Model
LR	Logistic Regression

LUCC	Land Use and Cover Change
ML	Machine Learning
MPB	Mountain Pine Beetle
ODD	Overview, Design, Details
ODD+D	Overview, Design, Details and Decisions
RF	Random Forest
RL	Reinforcement Learning
ROC	Receiver Operating Characteristic
SES	Social-Ecological System
SL	Supervised Learning
TD	Temporal Difference
TOC	Total Operating Characteristic
UL	Unsupervised Learning

To my loving family

Acknowledgements

First and foremost, I would like to express my deepest and most sincere thanks to my supervisors, Dr. Liliana Perez of the department of geography of Université de Montréal (UdeM), and Dr. Roberto Molowny-Horas of CREAf, Universitat Autònoma de Barcelona, for their invaluable support all the way in my doctoral studies. Dr. Perez, my direct supervisor, especially dedicated her time and effort generously to my scientific success and academic development, for which I can not thank her enough. Dr. Molowny-Horas, my co-supervisor from Spain, provided me with substantial help and advice countless times, for which I am utterly grateful. In addition to supervising my work, Dr. Perez provided a significant part of the funding for my doctoral studies through the NSERC Discovery Grant number RGPIN/05396 available to her, for which I am most thankful. I also received financial support, scholarships and awards from the UdeM Faculty of Graduate and Postdoctoral Studies and Faculty of Arts and Sciences, especially from the department of geography, as well as from Institut EDDEC, TD Bank, and the Canadian Association of Geographers. I am very thankful to all these institutions for their generous support. I am also thankful to the UdeM and especially to the academic and administrative staff of its department of geography. I would next like to express my sincere gratitude to Professor Robert Gilmore Pontius Jr of Clark University for his highly valuable collaboration in the publication that is presented in Chapter 3 of this thesis. It was indeed a great honor for me to work with such a distinguished scientist. I would like to extend my gratitude to the distinguished members of the jury for patiently and carefully reading my thesis and providing me with valuable comments and generous encouragement. I wish to thank my friends and colleagues in the Laboratoire de Géosimulation Environnementale of UdeM for the good memories they gave me and for their kind support of my work. Especially, I am thankful to Jonathan Gaudreau and Guillaume Arnoux-Hebert for the opportunity to collaborate with them in two scientific publications. I wish to thank Christina Lyster, Bijan Arfa, Dr. Mandana Arfa, Dr. Afshin Parsi, Daniel Daemi, Ali Bihamta, Shadan Sabzevari, Khodayar Jeiroodi, Anahita Atashbiz, Neda Asadi, Mathieu Chenette, the Simards, the Nabizadas, Dr. Hassan Sharifi, Dr. Faramarz Moattar, Dr. Parvin Nassiri, Kent Sjölander, and Iraj Fazeli. Finally, I thank my family, Esmat Fazeli, Abolfazl Harati, and Mohammad Harati for their unconditional love, which has been my never-ending source of motivation and inspiration.

Introduction

Presentation of the problem

Sustainable development: a very brief overview

Sustainable development is now the subject of a plethora of academic works, social debates, and policy and governance issues. The ensemble of ideas that comprise the concepts of sustainable development have evolved through time. In 1798, Thomas Malthus noted that population and resource access curves do not grow at the same rate. Specifically, he argued that while population increases exponentially, resource availability increases linearly. Therefore, there will inevitably be a critical point where the available resource will no longer be sufficient for the needs of the population (Malthus, 1998). This concern – about future state of the resources that a population depends upon – was echoed in other scholarly works later in the 19th century. For example, Stanley Jevons argued that the economic progress in Britain was dependent on coal as its energy resource, and warned about challenges that arise when this finite resource is exhausted (Jevons, 1865). In addition to the problem of resource depletion, broader environmental degradation issues received attention in mid 19th century. Specifically, in his book titled ‘Man and Nature: on Physical Geography as Modified by Human Action’, George Perkins Marsh argued that the impact of humans on the environment is not small and negligible. Drawing on historical examples, he warned that destruction of the environment can lead to collapse of civilizations (Marsh, 1864).

In the 20th century, resource and environmental issues received more attention and it became clear that they involve problems of collective action. This is particularly the case where a common resource is shared among several users. Arguably, each user may worry that other users over-consume the resource for their own prosperity in a competitive economy. In order not to be left behind others, each user is thus motivated to consume the resource as much as they can, which leads to fast depletion of the resource beyond the levels that the resource can be renewed. This unfortunate dynamic is known as the tragedy of the commons (Hardin, 1968). In order to avoid the tragedy of the commons, one of the possible solutions is to regulate resource use by a governing entity. Another possible solution is that a governing entity mediates and facilitates communication among users. With acknowledgement of the important role of governments in resource and environmental problems, sustainability issues increasingly became subjects of governmental and intergovernmental debates. By 1972, when the United Nations Conference on

Human Environment was held in Stockholm, many states were concerned about environmental degradation and pollution caused by industrial and economic activities. This conference ended with a declaration highlighting issues such as protection of the environment, reduction of pollutants, conservation of biodiversity, environment and development, and planning (United Nations Conference on the Human Environment, 1972). In 1987, the World Commission on Environment and Development issued a report titled ‘Our Common Future’ – also known as the Brundtland Report – in which sustainable development was defined as a development that provides the needs of the present time without compromising the ability to provide the needs of the future (Brundtland, 1987). This report addressed a variety of sustainability related issues including development, world economy, population, food security, ecosystems, energy, industries, cities, public goods and heritage, peace, and security (idem). To date, a large number of works have been dedicated to the study of sustainability, and a multiplicity of definitions and models have been given for sustainable development. Nevertheless, it is generally accepted that the concept of sustainable development includes the three pillars of environment, economy, and society (Purvis et al., 2019). This three-pillar model can be used in identifying challenges under each of the three classes of problems as well as their intersections. In this sense, it contributes to planning for sustainable development.

Given the multifaceted nature of the concept of sustainable development, implementing it involves a multitude of problems. The role of local and state governments is implicit in the the above descriptions of sustainability (Brundtland, 1987; United Nations Conference on the Human Environment, 1972). In this sense, sustainable development poses many challenges to governments. Some of those challenges are efficiency and representativeness of the government (Mancebo, 2010). There are often many stakeholders involved in sustainability problems – each with their own views and interests – and the government might need to function as an arbitrator between them in order to implement a sustainability agenda. This requires to represent the viewpoints of all stakeholders in government decision process. With any government decision, there may be stakeholders who do not consider it legitimate, which puts the government in a situation of choosing between use of authority for efficient management on the one hand, and reconsideration of its course of action on the other hand (idem). In any case, it is very important to engage the society in the implementation of sustainable development plans, otherwise some

groups in the society may think that the government is biased towards the interests of other groups (Mancebo, 2010; Robertson and Hull, 2003).

In brief, over time, the idea of sustainable development evolved from a concern about resource availability to a more comprehensive framework with economic, environmental and social objectives. Such evolution continues as more knowledge is gained through experience and efforts to implement the transition to sustainable development. As a result, the field now has frontiers such as corporate responsibility (Li and Toppinen, 2011; Maas and Reniers, 2014) and citizen responsibility (Macnaghten and Jacobs, 1997; Parag and Darby, 2009), which fall under the umbrella of engagement of stakeholders (Mancebo, 2010). The problems and challenges of sustainable development are plenty. Among them, this thesis addresses a particular one about government-stakeholder interactions.

A problem in sustainable development

Many problems in sustainable development involve a governing entity in charge of conservation of a natural resource, where that resource is at risk due to human or natural causes. By duty, the governing entity is interested in keeping the resource in a state of health, but its capacity is limited. In addition to the governing entity, the context of the problem involves resource users. If the users cooperate with the governing entity, together they may be able to save the resource and keep the resource system within sustainable limits. However, the problem is that the users often act based on self-interest (De Young, 2000), and they do not cooperate with the government if such cooperation involves costs or loss of revenue for them. Even though these actions are expected to have long term benefits for all, they are not associated with direct and immediate benefits for the individuals who do them. Therefore, one of the challenges of these situations is the lack of motivation among users for cooperation with the governing entity.

From the viewpoint of conservation and environmental sustainability, the ideal situation is that the users find the motivation for environmentally responsible behavior. Such ideal situation may not happen out of the users' sense of altruism (Kaplan, 2000). As an alternative, the governing entity may offer financial incentives to users who cooperate with it. However, that alternative is costly and puts the governing entity in a conflicting situation: on the one hand the governing entity desires that users cooperate with it, and on the other hand it desires to reduce expenditure. In other words, it will be in the economic interest of the governing entity that the

users do not cooperate with it, so that it does not have to pay incentives to them. That would lead the governing entity into a situation of conflict of interest. Alone, financial incentives may not lead to continuous environmentally responsible behavior (De Young, 2000; Katzev and Johnson, 1987). Another alternative for the governing entity is to impose its authority and force the users to perform the environmentally responsible behavior. However, experience shows that this alternative may fail as well (Blundell and Gullison, 2003; Feeny et al., 1990; Wagner, 2004; Wittemyer et al., 2011) if the society is not ready to accept the new regulations. Therefore, the ideal situation from the point of view of the governing entity is where the society of the users is ready to behave in a certain way that the governing entity wishes, without any need for the governing entity to use legal force or financial incentives. Furthermore, in the ideal situation, the prescribed behavior should lead to the conservation of the natural resource.

The above setting gives rise to a number of questions about the possibility of the said ideal situation. Is it possible for the governing entity to encourage the users, without force and/or financial incentives, to change their initial state of inaction? Is it possible that the users voluntarily cooperate with the governing entity with a motivation other than altruism or profit? Is it possible to create a new norm of environmentally responsible behavior in the society of users, simply by encouragement? Is it possible to save the natural resource by the prescribed behavior, if it is fully adopted by the society of users? What happens to the resource if the prescribed behavior is followed by only a fraction of the society? More fundamentally, how should these problems be approached?

These problems involve multiple complexities. For one thing, interactions within the social system create complex dynamics that are hard to predict. Decisions by the governing entity may influence users' behavior. In turn, users' behavior may influence the governing entity's decisions, as well as the behavior of other users. Besides, ecological systems are composed of elements with interactions among them in ways that are neither linear nor homogeneous. As such, it is difficult to predict the dynamics of ecological systems. In addition to all of this, new complexities arise when a social system is connected to an ecological system. The society, based on the needs and levels of access of its members, uses services that the ecosystem provides. In so doing, the society makes changes in the ecosystem. The used services flow in the society and contribute to its dynamics, bringing it to a new state of needs and access.

On the other hand, those changes in the ecosystem influence ecosystem dynamics, which may result in a new state of availability of services. Therefore, the coupling of social and ecological systems adds a new dimension of complexity to the problem.

Besides the above-mentioned conceptual complexities, there are other challenges in practice. Managers and decision makers have to consider the uncertainty that comes with the complexity of these systems. There exists a risk that intervention in a complex social-ecological system leads to unanticipated, costly, and irreversible changes. Therefore, a challenge in management of these systems is to gain insight and knowledge about the dynamics of the system without the risk of trial-and-error. Moreover, even if there exists scientific insight about a complex system, justifying the applicability of that scientific insight for decision makers remains a challenge.

This thesis focuses on a particular type of the above-mentioned sustainability problems, where the ecological system is a forest resource that is attacked by a pest, and the social system includes users of the resource and a governing entity. As described below, this problem is of multi-disciplinary nature. In it, changes in the ecological system can be viewed from the lense of land change science (Lambin et al., 2006). Formation of a common behavior among users is a problem in the domains of collective action (Nyborg et al., 2016; Ostrom, 2009) and social norms (Farrow et al., 2017; Nyborg et al., 2016). The study of the dynamics that emerge through coupling the society with the forest is the subject of social-ecological systems (Liu et al., 2007; Ostrom, 2009). In a broader view, this problem and the question of how to address it are in the realm of complex systems (Cosens et al., 2021; Filotas et al., 2014). Many studies of complex systems involve building and using models that replicate some aspects of those systems (Railsback and Grimm, 2012; Wolfram, 2002). Some other domains that related to these problems are complex systems modeling (Railsback and Grimm, 2012; Wilensky and Rand, 2015), model assessment (Rykiel, 1996; van Vliet et al., 2016), and Machine Learning algorithms (Alpaydin, 2020; van Vliet et al., 2016). These fields and approaches are briefly introduced in the following section. Then a detailed definition of the problem is given and objectives of the thesis are stated.

Conceptual and methodological background

Complex systems

When studying a particular issue it often helps to limit the scope of consideration to a certain set of objects or phenomena. These, in turn, are often related to one another and other objects or phenomena beyond the defined scope, as real world problems are normally independent of the way we see them. Therefore, such conceptualization will reduce the study subjects to a select number of components, their relations with each other, as well as their relations with their surroundings. The term *system* is used to refer to such conceptualization. An example of definition of system is “a representation of an entity as a complex whole open to feedback from its environment” (Ryan, 2008). Systems can be defined in diverse settings such as physics, chemistry, biology, information and organization, to name a few. Developments in each of these settings contributed to better understanding of the dynamics and complexity of systems.

One of the early domains of application of the concept of systems was thermodynamics. This field of study was built upon measurements of macroscopic properties of matter, such as temperature and pressure. Classical thermodynamics was able to describe certain processes, such as compression of a gas, in terms of changes in macroscopic properties. Next, another discipline was born to describe how the microscopic behavior of matter results in macroscopic properties. This newer discipline was the kinetic theory of gases (Demtröder, 2017), which became the foundation for statistical mechanics (Helrich, 2009). Kinetic theory regards gases not as microscopic entities, but as systems composed of numerous microscopic components – molecules – each of which behaves based on known physical laws. The theories including such physical laws were well established and supported with strong mathematical analytical tools. Therefore, the kinetic theory of gases and subsequently statistical mechanics were among the earliest domains that addressed problems of complex systems (Bertin, 2012).

Reflections on the notion of system gave rise to a new scientific discipline called General System Theory (GST) (Von Bertalanffy, 1950). This discipline emphasizes wholeness of systems rather than reducing them to their components, and seeks to identify organizations rather than disordered events in the physical world. In the point of view of this discipline, individual, isolated behavior of components of a system may be described by simple laws, but the integrated behavior of those components in the whole system is different and may not follow the same laws. Austrian biologist Ludwig Von Bertalanffy defines GST as a logical-mathematical basic

science discipline dedicated to the formulation and extraction of principles that are applicable to systems, regardless of the nature of their components and the nature of the interactions among them (*idem*). Von Bertalanffy notes similar – or isomorphic – laws in various domains, from physics to biology to medicine to psychology to social sciences and even philosophy; and he argues that such isomorphism of scientific law may be due to similarities in the configurations of various systems in the examples named above. In this sense, GST is a general discipline that is applicable to all other sciences, as long as they involve systems (*idem*).

In parallel with GST, important scientific developments in the 20th century lead to the establishment of fields of information theory and cybernetics. Information theory is best known by the works of Claude Shannon (1948), who in turn gave credit to prior works by Hartley (1928) and Nyquist (1924). Shannon, Hartly and Nyquist addressed the problem of communication of information via a medium with finite capacity. They were interested in optimizing communication such that the maximum amount of information could be transmitted through the medium. To that end, a fundamental step was to quantify information. A major contribution of information theory to science is the definition of quantity of information and entropy (Shannon, 1948). Information theory also accounts for the concept of redundancy of information (*idem*). Soon after, building on the idea of quantification of information, the field of cybernetics was developed for the analysis of communication and control in machines as well as living organisms (Wiener, 1950). With regard to control in those systems, cybernetics uses the notion of feedback – or sending a system’s output to its input (*idem*). In this sense, cybernetics enriched the studies of systems.

It makes sense for a simple system’s output change to be proportional to its input change. In such a case, a linear relationship exists between input and output of that system. Conversely, there are systems that are not so simple, and their output is not linearly related to their input. These systems are non-linear. Dynamics of non-linear systems has been the subject of analytical studies since before the 20th century. A well-known problem in classical physics was the calculation of orbits of celestial systems. These are systems composed of more than two bodies subject to forces of gravity due to masses of all the bodies. As in many other physical systems whose change depends on their previous state, these many-body systems have been described with differential equations. However, solution of those differential equations remained a

challenge. In 1890, Henri Poincaré published an important work on the three-body problem, in which he analyzed stability of results of systems of non-linear differential equations, identified limits to problems that can be integrated, and presented a generalization of his findings to the many-body problem (Poincaré, 2017). The works of Poincaré and others who followed him laid the foundations for a theory of non-linear systems (Mira, 1997). Of particular importance among them were the works of Aleksander Andronov about bifurcations in dynamic systems (idem). The field of dynamics of non-linear systems that was developed on analytical works of these scientists has also been influential in other domains where analytical methods based on differential equations cannot be applied, because findings of non-linear dynamics research have shed light on behaviors that may emerge from complex systems.

Building on the works of Poincaré on non-linear dynamics, and in parallel with Shannon's information theory, a closely related field of science was developed with a focus on the study of chaotic processes. Chaos theory distinguishes between random variation and chaotic behavior. It demonstrates how systems that are analytically described can end in very different states if their initial conditions change only slightly. It also introduces fractals, or self-similar structures, as abstract mathematical entities that can contribute to models of various applications (Bolotin et al., 2009).

Another scientific development that made a remarkable contribution to the science and study of systems came later in the 20th century in physical chemistry. Dissipative structures (Prigogine and Lefever, 1973) are formations that emerge in systems in open exchange with their environments and particularly far from equilibrium. Some order exists in these structures, but such order can not be described by principles of equilibrium. Ilya Prigogine won the Noble Prize in chemistry for the contribution of his theory of dissipative structures to thermodynamics and physical chemistry. However, the implications of his work are not limited to physical chemistry. Dissipative structures contribute to our understanding of self-organisation (Prigogine and Nicolis, 1985) of complex systems.

Studies of complexity led to a more clearly defined scope for systems science as well. Weinberg (1975) considers two ranges of complexity and randomness for subjects of study. He uses the expression 'unorganized complexity' for problems beyond a certain level of randomness, and recommends statistical treatment for them. Then, for problems of limited

randomness and limited complexity, he uses the expression ‘organized simplicity’ and recommends analytical treatment. For the remaining problems he uses the expression ‘organized complexity’. This latter group of problems is the domain of systems (Weinberg, 1975). Weinberg’s views are also shared by more recent authors such as Easterling and Polsky (2004) and Leveson (2016).

As a consequence of the above-mentioned developments, the notion of system complexity emerged and complex systems became a new field of study. A complex system, as defined by Batty and Torrens (2005), is “an entity, coherent in some recognizable way but whose elements, interactions, and dynamics generate structures and admit surprise and novelty that cannot be defined a priori.” Goldstein (1999) names four characteristics for complex systems: nonlinearity, self-organization, going beyond equilibrium, and existence of attractors other than a state of equilibrium. He mentions that these characteristics can cause emergence of new patterns in complex systems. Manson (2001) uses the expression ‘aggregate complexity’ to refer to the interaction of system components that results in holism and synergy. He highlights internal relationships, internal structure (subsystems), relationships with the environment, learning and memory, emergence, and change and evolution as key attributes of such aggregate complexity.

Complex systems science has evolved with contributions from a multitude of other scientific fields, and is therefore multidisciplinary by nature. Moreover, applications of complex systems appear in plenty of other disciplines, which leads to a stronger emphasis on the multidisciplinary nature of this domain. This is because, as argued by Von Bertalanffy (1950), what is of interest in this field is not the nature of the elements and interactions in a system, but the structure and configuration of those elements and interactions.

Social-Ecological Systems

The relations between human and resource systems and degradation of environmental resources have been key motivating themes of a wide range of interdisciplinary studies known as political ecology (Walker, 2005), which connects cultural ecology (Steward, 1972), community ecology (Putman, 1994), cybernetics (Bateson, 2000) and systems theory (Odum, 1983). Such studies have been various in scale, from human-environment interactions in small communities to global environmental issues. Some of the early works in the field tried to explain the accelerated degradation of the environment using analytical tools of both natural and social

sciences (Blaikie and Brookfield, 2015). In the works of the 1980s and early 1990s much attention was paid to biogeophysical research, in order to provide valid and reliable models of environmental systems and their resilience and sustainability (Walker, 2005). In subsequent years, questions of power inequality and competition over access to resource led to more political debates, to the extent that some theorists expressed their concern about the future of the field becoming 'politics without ecology' (Vayda and Walters, 1999). Later works marked the importance of integrating natural and social realms - as opposed to partitioning them - in analysis and in policy making. Liu et al. (2007) stated that the systems that involve both human and natural components demonstrate complexities that are hidden from the lens of social or ecological studies alone. Claiming that such complexity has not been well understood, they called for departure from existing approaches and endorsing integrated studies of Coupled Human And Natural Systems. Ostrom (2009), highlighting the challenge of combining social and ecological knowledge, proposed a framework to analyze the sustainability of complex Social-Ecological Systems (SES). In Ostrom's framework, four sub-systems compose a SES: governance system and resource system at larger scales, plus users and resource units at smaller scales. Within each of the social and ecological systems, changes can happen through top-down as well as bottom-up processes. Between the social and ecological systems, change can be caused by interactions. At the core of interactions between social and ecological systems are services that societies receive from ecosystems. These are known as ecosystem services (Daily, 2000; Millennium Ecosystem Assessment, 2003). With this view, many sustainability applications can be analyzed as coupled SES in which ecosystem services flow from the ecological system to the social system, while the social system cause environmental change in the ecological system. Figure I.1 demonstrates the core subsystems of Ostrom's framework and their interactions in a SES model.

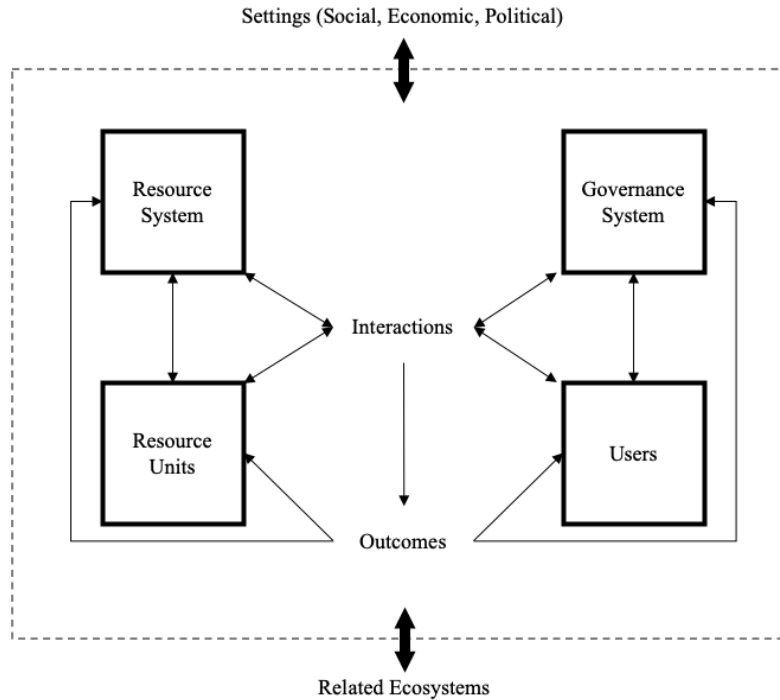


Figure I.1. Schematic overview of Ostrom’s SES framework (Ostrom, 2009). The dashed line indicates the border of the SES.

Ostrom’s framework defines second tier variables for each of the main subsystems and their interactions and outcomes (Ostrom, 2009). However, not all of those variables are pertinent in all applications (Basurto et al., 2013). Indeed, the relative importance of variables of each SES component differs from one case of study to another (Ostrom, 2009). Moreover, depending on the application, variables of third, fourth, and lower tiers may be defined to clarify the framework (Hinkel et al., 2015; McGinnis and Ostrom, 2014). In any case, it is important to note that constructing the SES framework for a given application is not the end but rather the beginning of the study on its respective SES governance questions (Basurto et al., 2013). Building on the premise of importance of correct diagnosis prior to intervention, the SES framework is seen as a tool that helps experts diagnose factors that determine the sustainability of a given SES. This tool does not replace expert judgement, but rather it supports expert judgement by providing a sketch for knowledge organization (McGinnis and Ostrom, 2014). SES literature emphasizes the importance of strong scholarly work that should accompany the use of the framework; otherwise, even with the SES framework, analysis may be misleading and accumulation of knowledge will be slow (Basurto et al., 2013).

The first step in implementing the SES framework for a given application involves identification of interactions and outcomes related to the resource system and resource units, relevant agents and governance systems (McGinnis and Ostrom, 2014). At this stage, care should be taken to clearly distinguish between resource systems and resource units. SES literature proposes a procedure to clarify this distinction with a set of questions about the agents that are involved in the SES, the benefits that those agents gain from the SES, and the situations of collective action where those benefits are produced (Hinkel et al., 2015).

Once the interactions and outcomes related to a SES are identified, the SES framework helps researchers select the pertinent variables that should be measured and watched. Although this selection can be done in various ways, using the SES framework reduces the risk of neglecting important variables (McGinnis and Ostrom, 2014), as the framework provides a basic heuristic tool for classification and organization of related knowledge (Basurto et al., 2013).

In addition to building a foundation for organizing knowledge related to a SES, the framework can be helpful in interpretation of the observations. As a diagnostic tool, the framework allows to clarify the distinction between lessons of general models on the one hand, and particular case analyses. At the same time, the framework allows to link and integrate the above in an overall insight (Basurto et al., 2013). Furthermore, by providing a common language, the framework facilitates communication between researchers from various domains of science (McGinnis and Ostrom, 2014). Such communication is especially important in the field of SES because of the multidisciplinary nature of this field. With effective communication, researchers can use and build on previous works of their peers. This, in turn, speeds up the process of knowledge development.

Despite the progress that has been made in the field of SES, there remain major challenges yet to be resolved. In particular, proper representation of complex SESs – i.e. systems with multiple dependencies between multiple agents and groups of agents – is highly important in order to understand or model them (Hinkel et al., 2015; McGinnis and Ostrom, 2014). The SES framework of Ostrom is not sufficient for capturing all the interdependencies in such complex systems. Moreover, states of agents and groups and subsystems in a SES are time-dependent. The dynamic aspects of these entities determine the evolution of the whole system; hence they are also very important to comprehend and to include in the model of the system.

This is the second major challenge of SES studies, for which the framework cannot provide a complete solution (Hinkel et al., 2015). These challenges call for more studies on specific cases as well as more general and theoretical work on SES problems.

Social norms

The problem of emergence of behaviors in the context of the society is a subject in the realm of social norms. Norms have been defined in various studies. According to Ross (1973), in a society, norms are cultural rules guiding people's behavior. Savarimuthu and Cranefield (2011) consider norms as social rules that govern how certain behaviors are encouraged or condemned. In the context of institutions, Ostrom (1990) writes that norms show the valuations of the actions of individuals in a society, regardless of the immediate consequences of those actions. According to Crawford and Ostrom (1995), norms are part of institutions, and deviating from them has unknown or undefined consequences. North (1990) states that institutions are able to formalize norms into laws, and enforce them legally. Literature reviews report that many of the previous studies associate norms with social sanctions, or with the punishment of individuals who do not follow norms (Hollander and Wu, 2011; Savarimuthu and Cranefield, 2011). However, the term 'norms' has also been used in studies of the emergence of behavioral expectations that do not involve sanctions (Savarimuthu and Cranefield, 2011). Cialdini et al. (1990) distinguish two types of norms, which they refer to as 'descriptive' and 'injunctive'. Descriptive norms inform the individual of what others in the society do. While injunctive norms urge the individual to do what others in the society approve of. According to Cialdini et al. these two types of norms come from different concepts and different motivations. Therefore, although what people do and what people approve of are often the same, separating these two norms is important in the study of normative influence. Injunctive norms are associated with social sanctions, whereas descriptive norms are not. In this study, our interest is in a setting without social sanctions. Therefore, in the rest of this research we focus on descriptive norms. In the context of sustainable development, a literature review by Bourceret et al., (2021) emphasizes that social norms are known to largely influence environmentally responsible behavior, and recommends policy-makers to take note of this matter (Bourceret et al., 2021). It is also known that people's decisions on environmental behavior are influenced by what others in the society do and approve (Nyborg et al., 2016), which is in line with the above-given definitions of descriptive norms and injunctive norms (Cialdini et al., 1990).

Complex Systems Modelling

Before trying to answer how to analyze such complex systems, it is worth noting that there are typically two objectives pursued in sustainable development studies and analyses: one is to understand why a system works the way it does, and the other is to foresee how it will respond to a certain intervention. The first type contributes to knowledge and the second type provides what-if analyses, which can be helpful for decision making. In the domain of sustainability, decision making is increasingly finding itself in need of knowledge of systems. Therefore, it can be said that the first objective - knowledge - serves the second objective - decision making. However, in making decisions that affect social and ecological systems there may be a tendency to avoid trial and error when and where possible, because some actions may lead to irreversible damage. In order to gain knowledge for decision-support without risking irreversible consequences of trials and errors, one approach is to conceptualize the system under study, including its components and interactions, in a model, and analyze such conceptual model to see the possible effects of the considered intervention. This, of course, requires for the model to be as close in behavior to the real system.

Model development generally involves mathematical analysis. This can be a challenge for models of complex systems with plenty of interactions and continuous change and evolution in their nature. Rather than analytical solutions, complex systems models are typically built using computers. Two widely used computational methods for the analysis of complex systems are known as Cellular Automata and Agent-Based Modelling.

Cellular Automata

Cellular automata (CA) can be described as grids containing cellular elements arranged adjacent to one another, where the state of each cell at each instant is defined based on a *transition rule* taking into account its previous state as well as its interactions with neighboring cells and the environment. Such interactions can be in form of flow of matter, energy, or information. Because of their neighborhood-based structure, CA can be efficient in modelling systems and phenomena that are spatially spread, systems in which there is a transfer of matter, energy, or information. Examples of applications of CA include models of urban growth, (Batty et al., 1999; Clarke et al., 1997; de Almeida et al., 2003) and forest disturbance (Bone et al., 2006; Gaudreau et al., 2016) among others.

The field of CA modelling was created as a result of the works of Von Neumann (1966) on automata theory. He was curious to find a computational mechanism capable of reproducing itself (Hoekstra et al., 2010). Von Neumann's automaton was a machine that would assume a state from a set of finite states, undergo a transition, and arrive at a new state. Based on the automaton, Von Neumann and Ulam developed their CA as a spatial arrangement of finite state machines where interactions were limited by spatial adjacency. In Von Neumann's design, each machine was placed in a cell in a rectangular grid, and the adjacency where direct interactions between machines were allowed was orthogonal. An alternative possible manner of defining the adjacency or neighborhood of interactions is to include diagonally adjacent grid cells in addition to orthogonally adjacent ones. This adjacency is called the Moore neighborhood (idem). Figure I.2 shows Von Neumann and Moore neighborhoods.




Figure I.2. Von Neumann and Moore neighborhoods (red) about a central cell (gray) in a rectangular grid

CA models became more popular as computational power became more easily accessible. One of the CA models that became well known was John Conway's Game of Life, in which each grid cell had a binary state of alive or dead, and the future state of each cell would be determined based on simple rules taking into account the number of alive and dead cells in its Moore neighborhood as well as its own state (Gardner, 1970). The Game of Life was a conceptual model and it solely demonstrated a CA example. Iterative changes of states of cells and the rules of the game would produce dynamic patterns that appeared, moved, merged with other patterns, froze, or dissolved in the 2-dimensional grid.

In the original works of Von Neumann and Ulam, CA model was defined in a 2-dimensional grid of rectangular cells (Hoekstra et al., 2010). It is possible to construct alternative

kinds of CA by modifying some of the configurations set by Von Neumann and Ulam. For one example, as mentioned above, Moore’s neighborhood was an alternative definition of adjacency. Indeed, many types of interaction neighborhoods can be defined for CA models. Larger Von Neumann and Moore neighborhoods, circular neighborhoods, and weighted neighborhoods are some possible examples. In the latter, interactions from neighboring cells are weighted by their distance from the central cell. Moreover, models can be constructed with non-rectangular cells. An example is hexagonal cells. As yet another remark, models do not have to be 2-dimensional. It is possible to build 3-dimensional CA models as well (idem). All these innovative modifications have made CA a more versatile method for modelling complex systems.

Agent-Based Models

Agent-Based Models (ABM) consist of simulated individuals or agents that take actions based on interactions with other agents and the environment. Agents have a *decision rule*. The actions of agents may in turn change the environment for other individuals. Together, the individuals - which may be numerous - define the behavior of the entire system. It may be that a simple decision rule - when repeatedly applied to plenty of individuals at various initial conditions - results in emergence of new and surprising patterns of behavior. ABM can be applied to a variety of problems. Examples of research carried out using this method include epidemiology (Perez and Dragicevic, 2009), crime reduction (Malleon et al., 2010), land development (Pooyandeh and Marceau, 2013) and forest disturbance (Katan and Perez, 2021; Perez and Dragicevic, 2010) among others. Figure I.3 depicts the concept of ABM with an example.

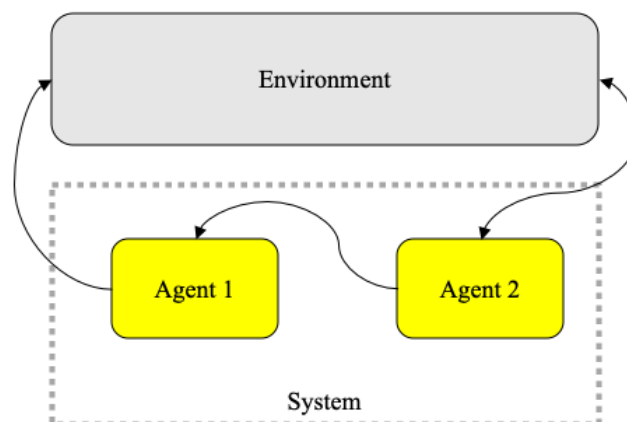


Figure I.3. An example ABM in abstraction. Dotted line represents system boundary. Each agent may be in one-way or two-way interaction with other agents in the system, or with the system’s environment.

It is noteworthy that models of complex systems, and especially ABMs, can be used in abstract simulations that aim to provide better understanding of situations and phenomena. A well-known example of such models was Thomas Schelling's (1971) segregation model. Schelling assumed an abstract society with individuals having some quantified level of tendency to reside in adjacency of neighbours that are similar to them. If an individual sees too few similar neighbours (too many neighbours that are different from it), it will move to a new place in the city. Schelling demonstrated that such a simple rule of decision results in the segregation of an initially randomly distributed population into clusters, each containing the same type of residents (*idem*). It is not easy to foresee the emergence of such a pattern without using an ABM. This is due to the inherent complexity of the setting of the problem. Schelling's model of segregation provides insight into such complexity.

In another example of abstract works, Robert Axelrod initiated and carried out a series of studies on the formation of cooperation using ABM (Axelrod and Dion, 1988). He, too, analyzed a hypothetical simplified situation, known as the prisoner's dilemma, in which two or more people should make decisions about cooperating with each other. The settings of the problem are such that cooperation is beneficial if adopted by all parties, but there is benefit in selfish action as well. When assessing the choice, one issue that everyone considers is that the benefit of selfish action is a certain gain, whereas the prize of cooperative action is only achievable if the other party or parties choose to cooperate as well, and involves therefore an element of risk. An ABM framework is well-suited for the analysis of such situations. Axelrod performed various tests by running the model with different setup conditions in order to find out the effects of different factors on the emergence of cooperation in finite iterations, that is, running the model with agents keeping their memory of previous iterations. These studies led to important insights about cooperative behaviour, particularly on the role of reciprocity. In fact, one of the decision strategies that was successful in many experiments was simply to make a cooperative choice at the first iteration and from then on, to repeat the other party's action of the last iteration; it was called TIT-FOR-TAT. Here, too, the simplified abstraction of the model creates a distance between the model and reality. Efforts have been made to add more features to the model and release some simplifying conditions. Further studies have been carried out on the role of factors such as number of agents, range of choices, and levels of benefit associated with each choice, to

name a few. The authors highlight that more complementary studies are needed to make the model more comprehensive and more realistic (idem).

The above two examples were important scientific contributions in their respective fields, even though they were developed and run in completely hypothetical settings and without reference to empirical data. This is important because it means that even abstract models can serve as laboratories for performing hypothetical experiments and gaining fundamental insight about a subject, as was the case in the above examples.

One similarity between ABM and CA models of complex systems is that they both involve interactions and a decision - or transition - rule. Defining these decision - or transition - rules is a key part of model development. One difference between the two types of models is that while cellular automata is typically applied to problems with a fixed spatial aspect, interactions in agent-based models need not be bound to certain spatial neighborhoods or otherwise spatial relationships. Another difference between the two is that in CA the states of all cells are updated at the same time, whereas ABM agents are event-driven and respond to signals or stimuli that they receive from other agents or the environment, in an asynchronous fashion.

Model assessment

In an overall view, model development can be considered to include four stages of planning, design, assessment, and documentation. In the planning stage, modellers should clarify research questions, collect relevant information and data, and make sure that necessary information for answering the research questions exists. If that is not the case, then research questions should be revised or otherwise the missing information should be provided. The design stage involves identification of a hierarchical structure of entities for the model and their linkages. This hierarchy will be such that higher level entities include lower level entities. Of particular importance is to identify the links between adjacent levels of the hierarchy. Model assessment should demonstrate that the logic behind the model has been correctly implemented, and it should present how well model results match reference information or otherwise make sense. Finally, documentation involves communicating the above to the scientific community so that they can provide further critique of the model as well as interpretation of its output (Aumann, 2007). Given the importance of model assessment in this process, it will be described in further detail in this section.

After the first version of a model has been built and the model is functioning and capable of producing results, it needs to be evaluated. The evaluation stage of the model is critical to ensuring the credibility of model results. This is particularly important for models that are intended to support policy makers in their decisions about complex problems, such as the sustainability of Social-Ecological Systems. By justifying the credibility of simulations, model assessment facilitates communication between scientists and policy-makers, which can lead to adoption of science-based policies.

A model is an abstraction that represents some features of an issue under study. The representation will demonstrate some relationships between its components. It is expected that the model should explain or otherwise behave in a way that is similar to something in the real world. Such similarity may be checked indirectly by means of indicators of results of processes of the real system and its model. However, it should be taken into consideration that sometimes similar results may be obtained from different processes. Evaluation of models is necessary before their results are applied in decision making or in the construction of knowledge.

Rykiel (1996) describes model evaluation as composed of the five steps of verification, calibration, validation, credibility and qualification. In his work, which has been widely appreciated and used in later literature (Ford, 2009; Grimm and Railsback, 2013; Guisan and Zimmermann, 2000; Pontius Jr et al., 2004; Refsgaard and Henriksen, 2004), he defines the steps of model evaluation as follows. Verification is the demonstration of correctness of the model's internal logic and its computer implementation. Calibration is the adjustment of model parameter values in such a way that model results are in best accordance with reference data. Validation is a demonstration that a model is acceptable for use in its domain. Credibility is a subjective judgment of the quality and sufficiency of the model for use as a basis of scientific and management decisions. Qualification refers to the generalization of model results to conditions beyond those of its definition. Among these steps, calibration and validation typically involve computational methods.

In their extensive literature review, van Vliet et al. (2016) identify five main calibration approaches: calibration based on expert knowledge, manual calibration, automated calibration, calibration using statistical analysis, and calibration based on other applications. Expert knowledge can be influential in realistic identification of CA transition rules or ABM decision

rules (Castle and Crooks, 2006; Verburg et al., 2006). Manual calibration involves running the model several times with different sets of parameters to find the set which maximizes the similarity of model output with real data - obviously a time consuming work. Automated calibration also involves testing different parameter sets but by a software. It is particularly useful when large numbers of tests are to be done. Computational power can be a limit to this approach. Statistical analysis involves extraction of parameters from real data using regression and other statistical techniques (Pontius Jr et al., 2004). This approach eliminates the need for searching the large state-space of model parameters for an optimum set. Finally, some modellers choose to import their parameter values from other applications. van Vliet et al. (2016) report that statistical analysis and automated procedures are the most commonly used calibration approaches among the reviewed models.

van Vliet et al. (2016) identified four main validation approaches based on: location accuracy, pattern accuracy, sensitivity analysis, and uncertainty analysis. Testing for location accuracy involves pixel-by-pixel comparison of simulated and observed maps (Pontius Jr, 2000; Pontius Jr et al., 2004). Testing for pattern accuracy is justified by the argument that due to complexity and stochasticity of modelled processes, location accuracy may not be a proper indicator of conformity of simulation with reality, but the model should be able to produce realistic patterns, which can be measured by a number of metrics such as number of clusters (Clarke et al., 1996). Sensitivity analysis involves testing the model outputs in different runs with slightly varied parameters (Fonoberova et al., 2013). Uncertainty analysis involves testing the model outputs in different runs with varied level of certainty of variables - particularly, by replacing fixed values with a distribution of possible values and corresponding probabilities (Railsback and Grimm, 2012). van Vliet et al. (2016) report that location accuracy is the base of the most commonly used validation approaches among the reviewed models.

Model validation is a subject worth more detailed consideration. In principle, this kind of model assessment requires comparing the output of the model with some reference. In practice, many – if not all – complex system models are used for the study of change processes that involve a temporal dimension. In this sense, the model's output indicates the simulation of the final state of a system. It is only via comparison with the initial state of the system that the simulated *change* can be identified. Similarly, in order to identify reference change, it is

necessary to compare reference information pertaining to the initial and final states of the system. Here, the terms *initial* and *final* refer to the before and after the simulation time interval, respectively. The implication of the above line of arguments is that validation of a model of change requires not two, but three sets of data: a reference indicator of the initial state of the system, a reference indicator of the final state, and a simulation of the final state (Pontius Jr et al., 2004).

The three datasets of reference and simulated states make it possible to produce two datasets of reference and simulated *change*. The latter two datasets can then be compared. In the simplest case, the state of each unit in the data is binary. In this sense, each unit of the system can be identified by its initial state and whether or not it has changed. The binary state constraint in this case means that the final state of the unit will be known. Then, four possibilities exist in the comparison of reference change with simulated change for each unit of the system:

- 1- Simulation and reference both indicate change. This case is called a *hit*, which means change was simulated correctly.
- 2- Simulation indicates no change, whereas reference indicates change. This is called a *miss*, which is an error.
- 3- Simulation indicates change, whereas reference indicates no change. This is called a *false alarm*, which is also an error.
- 4- Simulation and reference both indicate no change. This is called a *correct rejection*, which means persistence of state was simulated correctly.

By counting the total number of hits, misses and false alarms, useful information can be obtained about the performance of the model (Pontius Jr et al., 2004). Note that the sum of hits, misses, false alarms and correct rejections is equal to the total number of system units for which comparison has been made between simulation and reference. Therefore, it suffices to know only three of those four metrics.

A more complicated situation arises where the state is not binary, but a categorical variable with more than two modalities. In validation of these models, in addition to the four cases named above, a fifth possibility is that the simulation and reference both predict change,

however into different states. This error term should then be added in the analyses of model assessment (Pontius Jr et al., 2008).

In the above notes simulation and reference change information were assumed as simple sets of data, and comparison was made between respective elements of the datasets. A particular group of complex system models, especially in geography, include spatial data. In these cases, simulation and reference datasets are maps, and model validation calls for methods of map comparison. As described, location accuracy is the basis of the most common validation approaches for such models (van Vliet et al., 2016). Assessment of location accuracy typically begins with cell-by-cell comparison of simulation and reference change maps. The spatial nature of data in these applications makes more in-depth performance analysis possible. Specifically, sometimes the modeller may judge that some errors are justifiable if the model simulates change not at the exact locations of reference change, but at short distances from them. These are cases where a change was missed in simulation, and instead a nearby location was wrongly simulated as changed, hence creating a pair of miss and false alarm errors at a short distance from each other. If the modeller judges that such distance is short enough, then this case can be called a *near hit*. In contrast, if the distance between those errors is large, then it cannot be regarded as a small error or a near hit. The distinction between near hits and other errors can be made by using multiple resolution comparisons (Pontius Jr, 2002; Pontius Jr et al., 2008, 2004; Pontius Jr and Millones, 2011). In multiple resolution comparison, the maps are coarsened at various resolutions, and some of the near hits of the fine-resolution maps become hits and correct rejections in the coarse-resolution maps.

Machine Learning

Design, calibration and validation of models of complex systems often involve the application of methods of Machine Learning (ML), which is a dynamic front of research in Artificial Intelligence today. By merging statistics and computer science, ML seeks to enable computers to use available data to solve problems (Alpaydin, 2020). Up until the 1990s, research on machine learning was more focused on knowledge-driven approaches. But since then, the dominant approaches have been data-driven.

In 1956 in a conference in Dartmouth, U.K., the name ‘Artificial Intelligence’ (AI) was officially given to the field of research that aimed at creating capabilities similar to those of the

human brain, building on the foundations set by Godel and Turing (McCarthy et al., 2006). Early efforts based on first-order logic demonstrated success in tasks such as proving geometry theorems, solving algebra problems, and some processing of natural language using semantic nets (Ertel, 2011). The prospects for achieving the great goal of AI seemed very near. Funding was abundant and scientists explored and expanded various fronts of research. Many of the algorithms that were developed were based on searching a space of possible solutions. In the years that followed the search space for most applications expanded exponentially in a combinatorial explosion, and those algorithms were not able to perform in the new and large scales. This failure to meet expectations disappointed funding and the older wave of AI research slowed down dramatically by the 1980s (Ertel, 2011).

ML was an already established discipline under AI. Its earlier evolution, according to Shavlik & Dietterich (1990), had passed through stages of exploration in the 1950s and 1960s, algorithm development in the 1970s, and expansion from the 1980s. Some of the highlights of the exploration stage are the simulation of neurons and particularly the development of the perceptron by Rosenblatt in 1959; the simulation of evolution, random mutations, and natural selection; and research on non-supervised learning and particularly the development of a program that was capable of playing and improving at checkers by Samuel in 1959. Some of the highlights of the algorithm development stage are the introduction of blocks-world learning by Winston in 1970; and the development of induction learning guided by knowledge - that is, learning algorithms with knowledge-based rules - and particularly the work of Anderson in 1977. The 1980s witnessed a massive growth in ML research, particularly on learning theory, symbolic learning, neural networks, clustering, explanation-based learning, knowledge-guided induction learning, case-based reasoning, and genetic algorithms (Shavlik and Dietterich, 1990).

ML was one of the fronts that received increasingly more attention after the decline of the first-order-logic line of AI research. The use of techniques such as Bayesian nets, decision tree learning, and back-propagation neural networks increased remarkably from the 1980s, marking a new era in the history of AI (Alpaydin, 2020; Ertel, 2011). It is noteworthy that although the foundations of artificial neural networks were long developed - particularly with Rosenblatt's 'perceptron' in 1959 - research on this field gained popularity only in the 1980s, possibly due to the strong criticism by Minsky and Papert in 1969 (Shavlik and Dietterich, 1990).

Present day data is massive in terms of both number and attributes of observation. In most situations it is assumed that the variation in the data is not random but follows rules or patterns, though such rules and patterns are not known. In many real-world problems, the input and the type of output are known but the algorithm to connect the two is not, due to limited knowledge. Today's ML approaches these problems with the assumption that there are simple mechanisms that generate complex data, and aims to discover them. This can be used for description of systems, or for prediction of their future states (Alpaydin, 2020). ML techniques can be categorized into supervised, unsupervised, and reinforcement learning.

Supervised Learning

Supervised Learning (SL) techniques try to identify functions that relate given inputs to outputs - independent to dependent variables. Once they learn from the given data, they can be used to predict the output of new data. Two important groups of these techniques are classification learning and regression. If the type of the output is nominal - that is, if the sought function places the input data into categories - classification learning is used. If the output is numeric, regression can be used (Marsland, 2009).

Classification learning algorithms assign a category to each data item, learning from a given dataset that is labeled (Kubat, 2017). The smallest number of categories is 2. In that case, the output can be considered as binary. In cases with more categories, it is possible to represent outputs as binary arrays, with each element of the array corresponding to a category. In all cases the given dataset should be looked into and criteria for distinguishing among categories must be identified. That means it should be decided what attributes of the data are pertinent, and then it should be identified what range of values of those attributes relates to the considered category. The choice of attributes can be done using existing knowledge, and the range of values for each attribute should be found by looking into the given data. The latter will require a criterion for acceptability, which can be an error function to be minimized. If n attributes are considered, an n -dimensional interval will be identified as the location of the items of a certain category. This location is what the algorithm learns. Later, it should be tested with further independent data (Legendre and Legendre, 2012). A good learning algorithm should be able to categorize the new data items with least error. In the case of classification, a possible definition of error is the count of assigned labels that are not correct. Figure I.4 depicts a problem of supervised classification schematically.

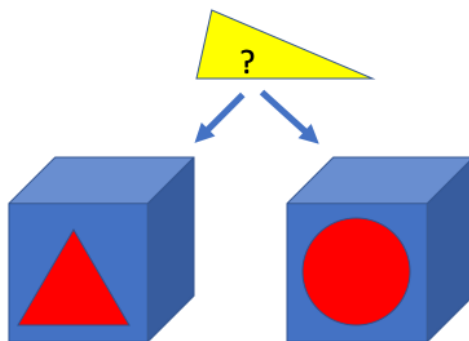


Figure I.4. Hypothetical example of supervised classification. Each box represents a class of objects. Correct labels for classes are given; the algorithm is expected to identify to which class each new object belongs.

Regression is similar to classification learning in basic principles, though it typically involves real functions (Zielesny, 2016). The given dataset should be looked into so that a relationship between the independent and dependent variables can be guessed. Then that guess should be tested with further data. The initial family of guesses - for example, linear functions - is selected using existing knowledge. Then within that family, one guess that shows the best fit with the given data is identified. This requires a measure of goodness of fit to be chosen as well. The same measure is used to report the goodness of fit of the chosen function with the test data. It is noteworthy that in regression it is possible to define error as a function of distance - the difference between predicted and observed numbers. Since the variables of the function are real numbers, it is possible to perform real analysis operations such as differentiation in order to find the parameter set that minimizes the error.

In both classification learning and regression, certain assumptions are made. Sometimes it is possible to analyze all attributes of the given data items. However, if only a selection of attributes is taken into consideration, that selection is an assumption. The family of initial guesses is another assumption. For example, in linear regression it is taken for granted that the dependent variable is a linear function of the independent variable(s). Furthermore, the definition of the error function and the process of finding the optimum guess are also assumptions. If any of these assumptions are made differently, the result of classification or regression may be different. On the other hand, if any of these is missing, the algorithms will be incomplete and cannot work. Therefore, such inductive bias always exists in this type of learning. It is important to select the

right bias, such that the model is neither too simple nor too sophisticated - neither underfit nor overfit with training data (Alpaydin, 2020).

Unsupervised Learning

Unsupervised Learning (UL) problems differ from SL applications in that they are not given any classification labels or values as the correct output to assign to the given input (Marsland, 2009). Hence as the name implies, there is no supervision in them - no provision of hints or reference information. On the contrary, there is only input data, with no dependent variables. The objective of UL techniques is to reveal the structure of the data and find its regularities - or anomalies (Alpaydin, 2020). Some of the most important groups of UL problems are clustering and reduction of dimensions of data.

Clustering techniques seek to find similarities among data items and to put similar ones together in clusters. The guiding principle in these problems is that items in the same cluster should be closer to (less distant from) one another than other clusters. This requires some measure of similarity (or distance) to be defined. One possible way to form clusters is to calculate the distance between all pairs of data items, sort the pairs by their distance, and make seeds of clusters by shortest-distance pairs, and grow them gradually by adding items that are closer to them than to other clusters - if a pair has no member in any clusters it can be the seed of a new cluster. This method establishes links between clusters and newly read items. If these links are defined as the shortest distance with any member of a cluster, the method is called single-linkage or nearest neighbor clustering (Legendre and Legendre, 2012). Another way to form clusters is to choose a given number of seeds - or to begin with a division of items into a given number of clusters - and iteratively assign items to the nearest cluster - that is, the cluster with the nearest centroid to the item - until there are no new exchanges in clusters between consecutive iterations. This method is known as k-means clustering (Legendre and Legendre, 2012), where k is the number of clusters and is decided by the user. The result of k-means clustering may vary depending on the initial selection of clusters. Figure I.5 depicts a problem of clustering schematically.




Figure I.5. Hypothetical example of clustering as an unsupervised learning problem. The algorithm is expected to identify groups of similar object based only on the attributes of the objects

Dimensionality reduction techniques read datasets with multiple attributes and identify a few important attributes (or a combination of them) along which the data shows most variation (Legendre and Legendre, 2012). Two broad classes of these techniques are feature selection and feature extraction. Feature selection algorithms identify a number of the multiple features - or attributes, or dimensions - of the data as the most informative ones. Feature extraction algorithms identify a number of combinations of the dimensions of the data (Alpaydin, 2020). Different techniques exist for different types of data. One of the most important feature extraction techniques for numerical data is Principal Component Analysis. It basically involves elimination of dependent dimensions (if any) and transformation and rotation of coordinates such that the axes of the new coordinate system display the most variance in the data; in other words, moving along those axes can display a large part of variations in the data (Alpaydin, 2020). In addition to possible advantages in terms of data storage or computational cost, dimensionality reduction is particularly useful for representation of data in 2 or 3 dimensions.

Reinforcement Learning

Reinforcement Learning (RL) techniques use a reward function to iteratively assess and improve their own performance (Sutton and Barto, 2018). In this sense RL is somewhat similar to SL because it includes some correction, but it is also different from SL since there is no given correct input-output pairs. RL is particularly good for problems where the final outcome is the result of a sequence of decisions. Such sequence of decisions is called a policy, and the goodness of a decision depends on the policy it is part of. In other words, the same single decision can be

made in different policies and the game may have different ends. Therefore RL algorithms should be able to assess effectiveness of policies (Alpaydin, 2020). Figure I.6 shows a RL problem schematically.



Figure I.6. Hypothetical example of Reinforcement Learning with a game of dart. Numbered crosses indicate three throws at the target. After the first throw, performance assessment was “above target” and the policy was “aim lower”. After the second throw, performance assessment was “too far below target” and the policy was “aim higher”.

In RL problems an agent tries to find a good sequence of actions. Let us assume that the set A of all possible actions and the set S of all states are finite. Transition model $T(s_t, a_t, s_{t+1})$ gives the probability of arriving at state s_{t+1} given state s_t and action a_t . The expected reward for this transition is written as $R(s_t, a_t, s_{t+1})$. Policy $\pi(s)$ suggests action a for state s . The RL problem then is formulated as the search for the optimum policy π^* which maximizes the expected cumulative reward throughout the game (Alpaydin, 2020).

Research setting, questions and objectives

Review of literature from the above disciplines, especially on the importance of personal motivations such as good reputation in voluntary action (Omoto and Snyder, 1995; Stern et al., 1993), the importance of visibility of one’s actions in one’s behavioral choices in society (Mosler, 1993; Nyborg et al., 2016), and the use of conceptual models in order to gain insight about complex SES (Janssen and Ostrom, 2006) inspired me to develop a conceptual setting in order to gain insight about the prospects of engagement of stakeholders in conservation of a

natural resource by using the stakeholders' desire for good reputation. This conceptual setting is described in the next paragraph.

A SES is comprised of a forest resource, forest users, and a governing entity. In an ecological disturbance, the forest resource is attacked by an insect infestation, which spreads progressively in the forest and kills trees every year. The governing entity has a goal to save the forest from the ecological disturbance. To that end, the governing entity needs the forest users to cooperate in implementing a management plan. Such cooperation is costly to the users. Therefore, the users are initially reluctant to cooperate with the governing entity. The governing entity does not give financial incentives to the users, and does not enforce its authority on the users. Instead, the governing entity introduces a scheme to recognize forest users who voluntarily cooperate with it in a task that the governing entity defines every year. Such recognition is done by giving the cooperating users a 'responsible user' label, which is visible to all users. Every year, the users consider the difficulty of the requested task and the desirability of the 'responsible user' label in their response to the governing entity. On the other hand, the function of the governing entity in this setting is to decide, every year, whether it should request the users to perform a difficult task or an easy, zero-cost task. In so doing, the governing entity considers previous responses of the users to its requests, as well as the state of the ecosystem. As such, over a period of several years, the work of the governing entity is summarized in a series of binary decisions. At the end of the period, the governing entity will be successful if, as a result of its decisions, the users cooperate with it in implementing the management plan, and if the management plan controls the ecological disturbance and saves the forest resource.

In the above setting, it is not known when or if users will cooperate with the governing entity, and it is not known if the management plan succeeds in saving the forest resource. Building a conceptual model with the said interactions allows to try several interventions in a simulated environment and perform 'what-if' analysis, while avoiding the risks associated with the uncertainties in the responses of the social and ecological systems.

In order to do 'what-if' analysis and gain insight about the ecosystem's response to various interventions, an ecological model is built and assessed. Then, the validated model is assumed to represent the reality with which the society interacts. The ecological model is coupled with a social model and used as a virtual laboratory, which allows hypothetical tests of

different scenarios. The ecosystem of this study is a forest in western Canada that is attacked by an insect infestation. In the ecological model of this study, the spread of the infestation in the forest is simulated.

Since the 1990s, forests of the Canadian Province of British Columbia (BC) were attacked by the Mountain Pine Beetle, a wood boring insect that primarily infested and killed lodgepole pines in the province. In the 2000s the infestations crossed the natural barrier of the Rocky Mountains and spread to the neighboring province of Alberta (Strohm et al., 2016). The life-cycle of the insect is one year. Each summer, Mountain Pine Beetles fly in search of new hosts to infest. Female insects find hosts and release pheromones that attract more females and males to their location. Females lay their eggs in galleries in the bark of the host tree. Over the next months eggs hatch and larvae further excavate inside the bark of the tree. The next summer, they emerge as fully developed insects and leave the tree. The beetles are carriers of the blue stain fungi, which they leave inside the host tree as they excavate its bark. The beetles and the larvae eat the tree's phloem, and together with the fungi they disrupt the tree's functioning and water and nutrient transport system, eventually killing the tree. Usually after one year, dead trees change color and become red, and in a few years their foliage falls and they appear gray. These phases are called red attack and gray attack, respectively. Tree stands in red and gray attacks can be identified in aerial images (Natural Resources Canada, 2015).

It is estimated that the insect has attacked over 25 million hectares of BC forests (Bone and Nelson, 2019) and damaged over half of the commercial pine volume in the province (Natural Resources Canada, 2019). In order to control the infestations, BC increased allowable annual cut in its forests. The initial motivation for this decision was to suppress the infestations, which was not successful. The province continued with its increased allowable annual cut in later years for harvesting wood from killed trees before they lose their economic value (Forest Practices Board, 2009, 2007). There were concerns that the policy of increasing the allowable annual cut would fail to control the infestation and only make the forest ecosystem vulnerable in the long term. The rationale for the concern of inefficacy of increased harvest rates was that even with the increased cuts, it would take decades to harvest the wood that was infested by mid 2000s (Forest Practices Board, 2009).

It must be noted that the Mountain Pine Beetle infestation – and forest disturbances in general – are not necessarily disastrous events in all situations. In particular, the Mountain Pine Beetle has been a native insect in the forests of western North America before its outbreaks became epidemic in the 1990s and 2000s (Natural Resources Canada, 2019). In general, species and processes that limit the life or growth of other species are integrated parts of complex adaptive ecological systems (Gunderson and Holling, 2002). Indeed, from time to time, regimes in forests as complex adaptive systems collapse due to disturbances, then resources that were exclusively used by previously dominant species become available to a vast variety of species, new processes are established, and new regimes emerge in the systems (idem). In this sense, disturbances play an important role that facilitates the continuation of life in such systems. Nevertheless, in a SES where livelihoods and economies in societies depend on natural resources, adaptation to change in the short term becomes a challenge for humans.

The overall goal of this study is to gain insight about the particular setting of SES elements and interactions described above in this section. This setting raises general questions regarding the possibility of success of the governing entity's recognition scheme in convincing the users to cooperate with it; and, assuming cooperation by the users, the possibility of success of the governing entity's management plan in saving the forest resource. Because of the complexity of this conceptual SES, a third question arises about the overall possibility of success of the governing entity in saving the forest resource, considering the combined effect of the recognition scheme and the management plan. Answering these questions has been the motivation for the present doctoral work. To that end, this doctoral work seeks a method to produce a sequence of decisions for the governing entity to promote a new behavior norm among the users society, and a method to shed light on implementation of various management plans amidst the ecological disturbance. As such, the objectives of this doctoral work are:

- To build and assess a conceptual model to simulate the emergence of a behavior norm in the users society under guidance by the governing entity, and answer the following question:
 - In a hypothetical social setting with egoist self-interested agents, is it possible to encourage those agents to perform a costly behavior voluntarily and without expectation of economic return?

- To build a land change model to simulate the spread of forest insect infestations, and answer the following question:
 - Given data on previous infestations, where does the insect infest next?
- To assess the land change model using reference observation data, and answer the following question:
 - Is where errors occur important in assessment of a land change model?
- To build a SES model by coupling the above two models, to run hypothetical experiments by implementing management plans in the SES simulations, and to interpret the outcome of such experiments. Specifically, to answer the following questions:
 - In a hypothetical social-ecological setting under disturbance, with the goal of saving the ecosystem from the disturbance, and with egoist self-interested social agents, is it possible to encourage those agents to perform a costly environmentally responsible behavior voluntarily and without expectation of economic return? Will the adoption of that behavior lead to saving the ecosystem from the disturbance?

Included publications

Corresponding to the above objectives, four research papers have been prepared in the course of this doctoral work. Each of those papers constitutes a chapter in this thesis, and is preceded by a linking paragraph for better integration in the thesis. As of the date of initial submission of this thesis, three of the four papers that are presented in the next chapters have been published in peer-reviewed journals. Remarks on comments received on each of these papers after its publication have been inserted in its corresponding chapters as a post-publication appendix.

References

- Alpaydin, E., 2020. Introduction to Machine Learning, 4th ed, Adaptive Computation and Machine Learning series. MIT Press.
- Aumann, C.A., 2007. A methodology for developing simulation models of complex systems. *Ecol. Modell.* 202, 385–396. <https://doi.org/10.1016/j.ecolmodel.2006.11.005>
- Axelrod, R., Dion, D., 1988. The Further Evolution of Cooperation. *Science* (80-.). 242, 1385–1390. <https://doi.org/10.1126/science.242.4884.1385>
- Basurto, X., Gelcich, S., Ostrom, E., 2013. The social–ecological system framework as a knowledge classificatory system for benthic small-scale fisheries. *Glob. Environ. Chang.* 23, 1366–1380. <https://doi.org/10.1016/j.gloenvcha.2013.08.001>
- Bateson, G., 2000. Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology. University of Chicago Press, Chicago.
- Batty, M., Torrens, P.M., 2005. Modelling and prediction in a complex world. *Futures* 37, 745–766. <https://doi.org/10.1016/j.futures.2004.11.003>
- Batty, M., Xie, Y., Sun, Z., 1999. Modeling urban dynamics through GIS-based cellular automata. *Comput. Environ. Urban Syst.* 23, 205–233. [https://doi.org/10.1016/S0198-9715\(99\)00015-0](https://doi.org/10.1016/S0198-9715(99)00015-0)
- Bertin, E., 2012. A Concise Introduction to the Statistical Physics of Complex Systems, SpringerBriefs in Complexity. Springer, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-23923-6>
- Blaikie, P., Brookfield, H., 2015. Land Degradation and Society. Routledge. <https://doi.org/10.4324/9781315685366>
- Blundell, A.G., Gullison, R.E., 2003. Poor regulatory capacity limits the ability of science to influence the management of mahogany. *For. Policy Econ.* 5, 395–405. [https://doi.org/10.1016/S1389-9341\(03\)00038-8](https://doi.org/10.1016/S1389-9341(03)00038-8)
- Bolotin, Y., Tur, A., Yanovsky, V., 2009. Chaos: Concepts, Control and Constructive Use, Understanding Complex Systems. Springer Berlin Heidelberg, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-00937-2>
- Bone, C., Dragicevic, S., Roberts, A., 2006. A fuzzy-constrained cellular automata model of forest insect infestations. *Ecol. Modell.* 192, 107–125. <https://doi.org/10.1016/j.ecolmodel.2005.09.013>

- Bone, Nelson, 2019. Improving Mountain Pine Beetle Survival Predictions Using Multi-Year Temperatures Across the Western USA. *Forests* 10, 866. <https://doi.org/10.3390/f10100866>
- Bourceret, A., Amblard, L., Mathias, J.D., 2021. Governance in social-ecological agent-based models: A review. *Ecol. Soc.* 26. <https://doi.org/10.5751/ES-12440-260238>
- Brundtland, G.H., 1987. Report of the World Commission on environment and development: "our common future".
- Castle, C., Crooks, A., 2006. Principles and concepts of agent-based modelling for developing geospatial simulations. Centre for Advanced Spatial Analysis (UCL), London, UK.
- Cialdini, R.B., Reno, R.R., Kallgren, C.A., 1990. A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *J. Pers. Soc. Psychol.* 58, 1015–1026. <https://doi.org/10.1037/0022-3514.58.6.1015>
- Clarke, K.C., Hoppen, S., Gaydos, L., 1997. A self-modifying cellular automaton model of historical urbanization in the San Francisco Bay area. *Environ. Plan. B Plan. Des.* 24, 247–261. <https://doi.org/10.1068/b240247>
- Clarke, K.C., Hoppen, S., Gaydos, L.J., 1996. Methods And Techniques for Rigorous Calibration of a Cellular Automaton Model of Urban Growth. *Third Int. Conf. Integr. GIS Environ. Model.* 21–25.
- Cosens, B., Ruhl, J.B., Soininen, N., Gunderson, L., Belinskij, A., Blenckner, T., Camacho, A.E., Chaffin, B.C., Craig, R.K., Doremus, H., Glicksman, R., Heiskanen, A.-S., Larson, R., Similä, J., 2021. Governing complexity: Integrating science, governance, and law to manage accelerating change in the globalized commons. *Proc. Natl. Acad. Sci.* 118, e2102798118. <https://doi.org/10.1073/pnas.2102798118>
- Crawford, S.E., Ostrom, E., 1995. A grammar of institutions. *Am. Polit. Sci. Rev.* 89, 582–600.
- Daily, G.C., 2000. Management objectives for the protection of ecosystem services. *Environ. Sci. Policy* 3, 333–339. [https://doi.org/10.1016/S1462-9011\(00\)00102-7](https://doi.org/10.1016/S1462-9011(00)00102-7)
- de Almeida, C.M., Batty, M., Monteiro, A.M.V., Câmara, G., Soares-Filho, B.S., Cerqueira, G.C., Pennachin, C.L., 2003. Stochastic cellular automata modeling of urban land use dynamics: Empirical development and estimation. *Comput. Environ. Urban Syst.* 27, 481–509. [https://doi.org/10.1016/S0198-9715\(02\)00042-X](https://doi.org/10.1016/S0198-9715(02)00042-X)
- De Young, R., 2000. New Ways to Promote Proenvironmental Behavior: Expanding and Evaluating Motives for Environmentally Responsible Behavior. *J. Soc. Issues* 56, 509–526.

<https://doi.org/10.1111/0022-4537.00181>

- Demtröder, W., 2017. *Mechanics and Thermodynamics, Undergraduate Lecture Notes in Physics*. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-27877-3>
- Easterling, W.E., Polsky, C., 2004. *Crossing the Complex Divide: Linking Scales for Understanding Coupled Human-Environment Systems*, in: McMaster, R., Sheppard, E. (Eds.), *Scale and Geographic Inquiry*. Blackwell Publishing Inc., Oxford, pp. 55–64.
- Ertel, W., 2011. *Introduction to Artificial Intelligence, Undergraduate Topics in Computer Science*. Springer London, London. <https://doi.org/10.1007/978-0-85729-299-5>
- Farrow, K., Grolleau, G., Ibanez, L., 2017. *Social Norms and Pro-environmental Behavior: A Review of the Evidence*. *Ecol. Econ.* <https://doi.org/10.1016/j.ecolecon.2017.04.017>
- Feeny, D., Berkes, F., McCay, B.J., Acheson, J.M., 1990. *The Tragedy of the Commons: Twenty-Two Years Later*. *Hum. Ecol.* 18, 1–19.
- Filotas, E., Parrott, L., Burton, P.J., Chazdon, R.L., Coates, K.D., Coll, L., Haeussler, S., Martin, K., Nocentini, S., Puettmann, K.J., Putz, F.E., Simard, S.W., Messier, C., 2014. *Viewing forests through the lens of complex systems science*. *Ecosphere* 5, 1–23. <https://doi.org/10.1890/ES13-00182.1>
- Fonoberova, M., Fonoberov, V.A., Mezić, I., 2013. *Global sensitivity/uncertainty analysis for agent-based models*. *Reliab. Eng. Syst. Saf.* 118, 8–17. <https://doi.org/10.1016/j.res.2013.04.004>
- Ford, A., 2009. *Modeling the Environment*. Island Press, Washington, D.C.
- Forest Practices Board, 2009. *Biodiversity conservation during salvage logging in the Central Interior of BC*.
- Forest Practices Board, 2007. *Tree Species Harvested In Areas Affected By Mountain Pine Beetles Special Report Table of Contents*.
- Gardner, M., 1970. *Mathematical Games*. *Sci. Am.* 223, 120–123. <https://doi.org/10.1038/scientificamerican1070-120>
- Gaudreau, J., Perez, L., Drapeau, P., 2016. *BorealFireSim: A GIS-based cellular automata model of wildfires for the boreal forest of Quebec in a climate change paradigm*. *Ecol. Inform.* 32, 12–27. <https://doi.org/10.1016/j.ecoinf.2015.12.006>
- Goldstein, J., 1999. *Emergence as a Construct: History and Issues*. *Emergence* 1, 49–72.

- https://doi.org/10.1207/s15327000em0101_4
- Grimm, V., Railsback, S.F., 2013. *Individual-based Modeling and Ecology*. Princeton University Press.
- Guisan, A., Zimmermann, N.E., 2000. Predictive habitat distribution models in ecology. *Ecol. Modell.* 135, 147–186.
- Gunderson, L.H., Holling, C.S., 2002. Panarchy: Understanding Transformations in Human and Natural Systems, *Angewandte Chemie International Edition*, 6(11), 951–952. Island Press, Washington, D.C.
- Hardin, G., 1968. The Tragedy of the Commons. *Science* (80-). 162, 1243–1248.
<https://doi.org/10.1126/science.162.3859.1243>
- Hartley, R.V.L., 1928. Transmission of Information 1. *Bell Syst. Tech. J.* 7, 535–563.
<https://doi.org/10.1002/j.1538-7305.1928.tb01236.x>
- Helrich, C.S., 2009. *Modern Thermodynamics with Statistical Mechanics*. Springer, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-540-85418-0>
- Hinkel, J., Cox, M.E., Schluter, M., Binder, C.R., Falk, T., 2015. A diagnostic procedure for applying the social-ecological systems framework in diverse cases. *Ecol. Soc.* 20, art32.
<https://doi.org/10.5751/ES-07023-200132>
- Hoekstra, A.G., Kroc, J., Sloom, P.M.A. (Eds.), 2010. *Simulating Complex Systems by Cellular Automata, Understanding Complex Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-12203-3>
- Hollander, C.D., Wu, A.S., 2011. The Current State of Normative Agent-Based Systems. *J. Artif. Soc. Soc. Simul.* 14. <https://doi.org/10.18564/jasss.1750>
- Janssen, M.A., Ostrom, E., 2006. Governing Social-Ecological Systems, in: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics* (Vol. 2). Elsevier B.V., pp. 1465–1509. [https://doi.org/10.1016/S1574-0021\(05\)02030-7](https://doi.org/10.1016/S1574-0021(05)02030-7)
- Jevons, W.S., 1865. *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal Mines*. MacMillan and Co., London.
- Kaplan, S., 2000. New Ways to Promote Proenvironmental Behavior: Human Nature and Environmentally Responsible Behavior. *J. Soc. Issues* 56, 491–508.
<https://doi.org/10.1111/0022-4537.00180>
- Katan, J., Perez, L., 2021. ABWiSE v1.0: Toward and Agent-Based Approach to Simulating

- Wildfire Spread. *Nat. Hazards Earth Syst. Sci.* 1–27. <https://doi.org/10.5194/nhess-2021-179>
- Katzev, R.D., Johnson, T.R., 1987. Promoting energy conservation: An analysis of behavioral research. Westview Press, Boulder, CO, United States.
- Kubat, M., 2017. An Introduction to Machine Learning. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-63913-0>
- Lambin, E.F., Geist, H., Rindfuss, R.R., 2006. Introduction: Local Processes with Global Impacts, in: Lambin, E.F., Geist, H. (Eds.), *Land-Use and Land-Cover Change*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1–8. https://doi.org/10.1007/3-540-32202-7_1
- Legendre, P., Legendre, L.F., 2012. *Numerical Ecology*, Third Engl. ed. Elsevier, Amsterdam, The Netherlands.
- Leveson, N.G., 2016. *Engineering a safer world: Systems thinking applied to safety*. MIT Press.
- Li, N., Toppinen, A., 2011. Corporate responsibility and sustainable competitive advantage in forest-based industry: Complementary or conflicting goals? *For. Policy Econ.* 13, 113–123. <https://doi.org/10.1016/j.forpol.2010.06.002>
- Liu, J., Dietz, T., Carpenter, S.R., Alberti, M., Folke, C., Moran, E., Pell, A.N., Deadman, P., Kratz, T., Lubchenco, J., Ostrom, E., Ouyang, Z., Provencher, W., Redman, C.L., Schneider, S.H., Taylor, W.W., 2007. Complexity of Coupled Human and Natural Systems. *Science* (80-.). 317, 1513–1516. <https://doi.org/10.1126/science.1144004>
- Maas, S., Reniers, G., 2014. Development of a CSR model for practice: connecting five inherent areas of sustainable business. *J. Clean. Prod.* 64, 104–114. <https://doi.org/10.1016/j.jclepro.2013.07.039>
- Macnaghten, P., Jacobs, M., 1997. Public identification with sustainable development. *Glob. Environ. Chang.* 7, 5–24. [https://doi.org/10.1016/S0959-3780\(96\)00023-4](https://doi.org/10.1016/S0959-3780(96)00023-4)
- Malleson, N., Heppenstall, A., See, L., 2010. Crime reduction through simulation: An agent-based model of burglary. *Comput. Environ. Urban Syst.* 34, 236–250. <https://doi.org/10.1016/j.compenvurbsys.2009.10.005>
- Malthus, T.R., 1998. *An essay on the principle of population* (Original work published 1798). Electronic Scholarly Publishing.
- Mancebo, F., 2010. *Le développement durable*. Armand Colin. <https://doi.org/10.3917/arco.mance.2010.01>

- Manson, S.M., 2001. Simplifying complexity: A review of complexity theory. *Geoforum* 32, 405–414. [https://doi.org/10.1016/S0016-7185\(00\)00035-X](https://doi.org/10.1016/S0016-7185(00)00035-X)
- Marsh, G.P., 1864. *Man and Nature: on Physical Geography as Modified by Human Action*. Charles Scribner, New York.
- Marsland, S., 2009. *Machine learning: an algorithmic perspective*. CRC Press, Boca Raton, FL.
- McCarthy, J., Minsky, M.L., Rochester, N., Shannon, C.E., 2006. A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Mag.* 27, 12. <https://doi.org/10.1609/aimag.v27i4.1904>
- McGinnis, M.D., Ostrom, E., 2014. Social-ecological system framework: initial changes and continuing challenges. *Ecol. Soc.* 19, art30. <https://doi.org/10.5751/ES-06387-190230>
- Millennium Ecosystem Assessment, 2003. Ecosystems and their services, in: *Ecosystems and Human Well-Being: A Framework for Assessment*. Island Press, Washington, DC, pp. 49–70. <https://doi.org/10.1007/s13398-014-0173-7.2>
- Mira, C., 1997. Some historical aspects of nonlinear dynamics—Possible trends for the future. *J. Franklin Inst.* 334, 1075–1113. [https://doi.org/10.1016/S0016-0032\(97\)00030-6](https://doi.org/10.1016/S0016-0032(97)00030-6)
- Mosler, H.-J., 1993. Self-dissemination of environmentally-responsible behavior: The influence of trust in a commons dilemma game. *J. Environ. Psychol.* 13, 111–123. [https://doi.org/10.1016/S0272-4944\(05\)80144-6](https://doi.org/10.1016/S0272-4944(05)80144-6)
- Natural Resources Canada, 2019. Mountain pine beetle [WWW Document]. URL <https://www.nrcan.gc.ca/our-natural-resources/forests-forestry/wildland-fires-insects-disturban/top-forest-insects-diseases-cana/mountain-pine-beetle/13381> (accessed 6.2.20).
- Natural Resources Canada, 2015. Mountain pine beetle factsheet [WWW Document]. URL <https://tidcf.nrcan.gc.ca/en/insects/factsheet/2816> (accessed 6.2.20).
- North, D.C., 1990. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/CBO9780511808678>
- Nyborg, K., Anderies, J.M., Dannenberg, A., Lindahl, T., Schill, C., Schlüter, M., Adger, W.N., Arrow, K.J., Barrett, S., Carpenter, S., Chapin, F.S., Crépin, A.-S., Daily, G., Ehrlich, P., Folke, C., Jager, W., Kautsky, N., Levin, S.A., Madsen, O.J., Polasky, S., Scheffer, M., Walker, B., Weber, E.U., Wilen, J., Xepapadeas, A., de Zeeuw, A., 2016. Social norms as solutions. *Science* (80-.). 354, 42–43. <https://doi.org/10.1126/science.aaf8317>
- Nyquist, H., 1924. Certain Factors Affecting Telegraph Speed. *Trans. Am. Inst. Electr. Eng.*

- XLIII, 412–422. <https://doi.org/10.1109/T-AIEE.1924.5060996>
- Odum, H.T., 1983. *Systems ecology: an introduction*. Wiley, New York, NY.
- Omoto, A.M., Snyder, M., 1995. Sustained helping without obligation: Motivation, longevity of service, and perceived attitude change among AIDS volunteers. *J. Pers. Soc. Psychol.* 68, 671–686. <https://doi.org/10.1037/0022-3514.68.4.671>
- Ostrom, E., 2009. A General Framework for Analyzing Sustainability of Social-Ecological Systems. *Science* (80-.). 325, 419–422. <https://doi.org/10.1126/science.1172133>
- Ostrom, E., 1990. *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press, Cambridge, UK.
- Parag, Y., Darby, S., 2009. Consumer–supplier–government triangular relations: Rethinking the UK policy path for carbon emissions reduction from the UK residential sector. *Energy Policy* 37, 3984–3992. <https://doi.org/10.1016/j.enpol.2009.04.044>
- Perez, L., Dragicevic, S., 2010. Modeling mountain pine beetle infestation with an agent-based approach at two spatial scales. *Environ. Model. Softw.* 25, 223–236. <https://doi.org/10.1016/j.envsoft.2009.08.004>
- Perez, L., Dragicevic, S., 2009. An agent-based approach for modeling dynamics of contagious disease spread. *Int. J. Health Geogr.* 8, 50. <https://doi.org/10.1186/1476-072X-8-50>
- Poincaré, H., 2017. *The Three-Body Problem and the Equations of Dynamics*, Original F. ed, Astrophysics and Space Science Library. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-52899-1>
- Pontius Jr, R.G., 2002. Statistical methods to partition effects of quantity and location during comparison of categorical maps at multiple resolutions. *Photogramm. Eng. Remote Sensing* 68, 1041–1050.
- Pontius Jr, R.G., 2000. Quantification error versus location error in comparison of categorical maps. *Photogramm. Eng. Remote Sens.* 66, 1011–1016.
- Pontius Jr, R.G., Huffaker, D., Denman, K., 2004. Useful techniques of validation for spatially explicit land-change models. *Ecol. Modell.* 179, 445–461. <https://doi.org/10.1016/j.ecolmodel.2004.05.010>
- Pontius Jr, R.G., Millones, M., 2011. Death to Kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *Int. J. Remote Sens.* 32, 4407–4429. <https://doi.org/10.1080/01431161.2011.552923>

- Pontius Jr, R.G., Thontteh, O., Chen, H., 2008. Components of information for multiple resolution comparison between maps that share a real variable. *Environ. Ecol. Stat.* 15, 111–142. <https://doi.org/10.1007/s10651-007-0043-y>
- Pooyandeh, M., Marceau, D.J., 2013. A spatial web/agent-based model to support stakeholders' negotiation regarding land development. *J. Environ. Manage.* 129, 309–323. <https://doi.org/10.1016/j.jenvman.2013.07.028>
- Prigogine, I., Lefever, R., 1973. *Theory of Dissipative Structures*, in: *Synergetics*. Vieweg+Teubner Verlag, Wiesbaden, pp. 124–135. https://doi.org/10.1007/978-3-663-01511-6_10
- Prigogine, I., Nicolis, G., 1985. Self-Organisation in Nonequilibrium Systems: Towards A Dynamics of Complexity, in: Hazewinkel, M., Jurkovich, R., Paelinck, J.H.P. (Eds.), *Bifurcation Analysis*. Springer Netherlands, Dordrecht, pp. 3–12. https://doi.org/10.1007/978-94-009-6239-2_1
- Purvis, B., Mao, Y., Robinson, D., 2019. Three pillars of sustainability: in search of conceptual origins. *Sustain. Sci.* 14, 681–695. <https://doi.org/10.1007/s11625-018-0627-5>
- Putman, R.J., 1994. *Community Ecology*. Kluwer, Dordrecht, Netherlands.
- Railsback, S.F., Grimm, V., 2012. *Agent-based and Individual-based Modeling: A Practical Introduction*. Princeton University Press.
- Refsgaard, J.C., Henriksen, H.J., 2004. Modelling guidelines—terminology and guiding principles. *Adv. Water Resour.* 27, 71–82. <https://doi.org/10.1016/j.advwatres.2003.08.006>
- Robertson, D.P., Hull, R.B., 2003. Public ecology: an environmental science and policy for global society. *Environ. Sci. Policy* 6, 399–410. [https://doi.org/10.1016/S1462-9011\(03\)00077-7](https://doi.org/10.1016/S1462-9011(03)00077-7)
- Ross, H.L., 1973. *Perspectives on social order*. McGraw-Hill, New York.
- Ryan, A.J., 2008. What is a Systems Approach?
- Rykiel, E.J., 1996. Testing ecological models: The meaning of validation. *Ecol. Modell.* 90, 229–244. [https://doi.org/10.1016/0304-3800\(95\)00152-2](https://doi.org/10.1016/0304-3800(95)00152-2)
- Savarimuthu, B.T.R., Cranefield, S., 2011. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent Grid Syst.* 7, 21–54. <https://doi.org/10.3233/MGS-2011-0167>
- Schelling, T.C., 1971. Dynamic models of segregation. *J. Math. Sociol.* 1, 143–186.

- <https://doi.org/10.1080/0022250X.1971.9989794>
- Shannon, C.E., 1948. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* 27, 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- Shavlik, J.W., Dietterich, T., 1990. *Readings in Machine Learning*. Morgan Kaufmann.
- Stern, P.C., Dietz, T., Kalof, L., 1993. Value Orientations, Gender, and Environmental Concern. *Environ. Behav.* 25, 322–348. <https://doi.org/10.1177/0013916593255002>
- Steward, J.H., 1972. *Theory of Culture Change: The Methodology of Multilinear Evolution*. University of Illinois Press, Chicago.
- Strohm, S., Reid, M.L., Tyson, R.C., 2016. Impacts of management on Mountain Pine Beetle spread and damage: A process-rich model. *Ecol. Modell.* 337, 241–252. <https://doi.org/10.1016/J.ECOLMODEL.2016.07.010>
- Sutton, R.S., Barto, A.G., 2018. *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, Cambridge, Massachusetts.
- United Nations Conference on the Human Environment, 1972. *Declaration of the United Nations Conference on the Human Environment*.
- van Vliet, J., Bregt, A.K., Brown, D.G., van Delden, H., Heckbert, S., Verburg, P.H., 2016. A review of current calibration and validation practices in land-change modeling. *Environ. Model. Softw.* 82, 174–182. <https://doi.org/10.1016/j.envsoft.2016.04.017>
- Vayda, A.P., Walters, B.B., 1999. Against Political Ecology. *Hum. Ecol.* 27, 167–179.
- Verburg, P.H., Kok, K., Pontius Jr, R.G., Veldkamp, A., 2006. Modeling Land-Use and Land-Cover Change, in: Lambin, E.F., Geist, H. (Eds.), *Land-Use and Land-Cover Change*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 117–135. https://doi.org/10.1007/3-540-32202-7_5
- Von Bertalanffy, L., 1950. An Outline of General System Theory. *Br. J. Philos. Sci.* 1, 134–165. <https://doi.org/10.1093/bjps/I.2.134>
- Von Neumann, J., 1966. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana and London.
- Wagner, W., 2004. Commons Ignorance: The Failure of Environmental Law to Produce Needed Information on Health and the Environment. *Duke Law J.* 53, 1619–1746.
- Walker, P.A., 2005. Political ecology: Where is the ecology? *Prog. Hum. Geogr.* 29, 73–82.
- Weinberg, G.M., 1975. *An introduction to general systems thinking*. Wiley, New York, NY,

USA.

Wiener, N., 1950. Cybernetics. Bull. Am. Acad. Arts Sci. 3, 2–4.

<https://doi.org/10.2307/3822945>

Wilensky, U., Rand, W., 2015. An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo, The MIT Press. MIT Press.

Wittemyer, G., Daballen, D., Douglas-Hamilton, I., 2011. Rising ivory prices threaten elephants. Nature 476, 282–283. <https://doi.org/10.1038/476282c>

Wolfram, S., 2002. A new kind of science. Wolfram media, Champaign, IL, United States.

Zielesny, A., 2016. From Curve Fitting to Machine Learning, Intelligent Systems Reference Library. Springer International Publishing, Cham. [https://doi.org/10.1007/978-3-319-32545-](https://doi.org/10.1007/978-3-319-32545-3)

3

Chapter 1

Presentation of the article

The first chapter of the thesis addresses an abstract problem of elicitation of cooperation of members of a society, which I call users, by a governing entity. In this problem, the governing entity desires to promote a new behavior in the society. In the eyes of the governing entity, success in this problem happens when the majority of users perform that new behavior. The governing entity's challenges are that firstly, users act upon self-interest, secondly, the intended new behavior is costly to users, and thirdly, it is assumed that the governing entity cannot apply force or use financial incentives. In this chapter a mechanism is introduced by which the governing entity offers a 'responsible user' label to users in exchange for voluntarily performing an action that the governing entity decides. The governing entity declares criteria for obtaining the 'responsible user' label, and users decide about doing what the governing entity requests. This scheme of interactions is iterated in simulations produced by a conceptual model. I use this conceptual model to answer if such a setting may result in the emergence of the behavior intended by the governing entity among the majority of the users.

This chapter has been published in the peer-reviewed journal "Applied Sciences" in 2021. My coauthors in this publication were my supervisors, Dr. Liliana Perez and Dr. Roberto Molowny-Horas. The chapter as it appears in this thesis involves modifications in the layout and style of the published paper, and slight modifications in figure and table numbers. Other than those, this chapter includes no changes in the content of the published paper.

Citation:

Harati, S., Perez, L., Molowny-Horas, R. Promoting the Emergence of Behavior Norms in a Principal-Agent Problem – An Agent-Based Modeling Approach Using Reinforcement Learning. *Appl. Sci.* **2021**, *11*, 8368. <https://doi.org/10.3390/app11188368>

Promoting the Emergence of Behavior Norms in a Principal–Agent Problem—An Agent-Based Modeling Approach Using Reinforcement Learning

Saeed Harati ^{1,*}, Liliana Perez ¹ and Roberto Molowny-Horas ²

¹ Laboratory of Environmental Geosimulation (LEDGE), Department of Geography, Université de Montréal, 1375 Ave. Thérèse-Lavoie-Roux, Montreal, H2V 0B3 QC, Canada; l.perez@umontreal.ca

² CREAM, Centre de Recerca Ecològica i Aplicacions Forestals, E08193 Bellaterra (Cerdanyola de Vallès), Catalonia, Spain; roberto@creaf.uab.es

* Correspondence: saeed.harati.asl@umontreal.ca

Harati, S., Perez, L., Molowny-Horas, R. Promoting the Emergence of Behavior Norms in a Principal–Agent Problem—An Agent-Based Modeling Approach Using Reinforcement Learning. *Appl. Sci.* **2021**, *11*, 8368. <https://doi.org/10.3390/app11188368>

Abstract

One of the complexities of social systems is the emergence of behavior norms that are costly for individuals. Study of such complexities is of interest in diverse fields ranging from marketing to sustainability. In this study we built a conceptual Agent-Based Model to simulate interactions between a group of agents and a governing agent, where the governing agent encourages other agents to perform, in exchange for recognition, an action that is beneficial for the governing agent but costly for the individual agents. We equipped the governing agent with six Temporal Difference Reinforcement Learning algorithms to find sequences of decisions that successfully encourage the group of agents to perform the desired action. Our results show that if the individual agents' perceived cost of the action is low, then the desired action can become a trend in the society without the use of learning algorithms by the governing agent. If the perceived cost to individual agents is high, then the desired output may become rare in the space of all possible outcomes but can be found by appropriate algorithms. We found that Double

Learning algorithms perform better than other algorithms we used. Through comparison with a baseline, we showed that our algorithms made a substantial difference in the rewards that can be obtained in the simulations.

Keywords: complex systems; emergence; reinforcement learning; temporal difference learning; social status

1.1. Introduction

One of the challenges of management in general, and sustainable development management in particular, is to gain the support of the individuals who are being managed. The use of incentives can be costly to managers and governments, and the use of authority is not always successful (Blundell & Gullison, 2003; Feeny, Berkes, McCay, & Acheson, 1990; Wagner, 2004; Wittemyer, Daballen, & Douglas-Hamilton, 2011). These problems, where a Principal (or several Principals) wishes to make an Agent (or several Agents) behave in a certain way are known as Principal–Agent problems (Braun & Guston, 2003).

In this study, we are interested in learning if the Principal can use recognition and the offer of good reputation to promote a new behavioral norm among the Agents. We are particularly interested in the complexities that emerge with the new norm, as the norm influences and is influenced by the decisions of the Agents. In this regard, social science literature describes a focus theory of normative conduct (Cialdini, Reno, & Kallgren, 1990), which suggests that in making decisions, individuals consider what others do and what others approve of. We illustrate an implication of this theory in a Principal–Agent setting. We would like to see if the Agents’ regard for their image in their society can lead to the emergence of a behavior norm that the Principal desires. Specifically, it is interesting for us to learn if, in absence of social sanctions and other forms of enforcement, good reputation can be a sufficient motivation for Agents to cooperate with the Principal. We would also like to see if the Principal’s intervention can hasten the emergence of this norm. This study is an effort to gain insight into the complexities that arise in an abstract Principal–Agent setting with the added consideration of normative conduct. We are curious about and intrigued by the complexities related to the abstract structure of entities, motivations, and interactions in the above setting.

Norms have been defined in various studies. For example, according to Ross (Ross, 1973), in a society, norms are cultural rules guiding people's behavior. Savarimuthu and Cranefield (Savarimuthu & Cranefield, 2011) consider norms as social rules that govern how certain behaviors are encouraged or condemned. In the context of institutions, Ostrom (Ostrom, 1990) writes that norms show the valuations of the actions of individuals in a society, regardless of the immediate consequences of those actions. In Crawford and Ostrom's view (Crawford & Ostrom, 1995), norms are part of institutions, and deviating from them has unknown or undefined consequences. North (D. C. North, 1990) states that institutions are able to formalize norms into laws, and enforce them legally. Literature reviews report that many of the previous studies associate norms with social sanctions, or with the punishment of individuals who do not follow norms (Hollander & Wu, 2011; Savarimuthu & Cranefield, 2011). However, the term 'norms' has also been used in studies of the emergence of behavior expectations that do not involve sanctions (Savarimuthu & Cranefield, 2011). Cialdini et al. (Cialdini et al., 1990) distinguish two types of norms, which they refer to as 'descriptive' and 'injunctive'. Descriptive norms inform the individual of what others in the society do. Injunctive norms urge the individual to do what others in the society approve of, and to avoid things of which others disapprove. According to Cialdini et al. these two types of norms come from different concepts and different motivations. Therefore, although what people do and what people approve of are often the same, separating these two norms is important in the study of normative influence. In order to avoid confusion regarding social sanctions, we follow the recommendation of Cialdini et al. Injunctive norms are associated with social sanctions, whereas descriptive norms are not. In this study, our interest is in a setting without social sanctions. Therefore, in the rest of this paper we focus on descriptive norms.

We take a complex systems approach to analyze the above problem. Complex systems are structures composed of elements, interactions and dynamics in such a way that they produce novel configurations and demonstrate surprising emerging behavior (Batty & Torrens, 2005). Some characteristics for complex systems are nonlinearity, self-organization, going beyond equilibrium, and existence of attractors other than a state of equilibrium, such that the combination of these characteristics can cause the emergence of new patterns in complex systems (Goldstein, 1999). Complex systems literature uses the expression 'aggregate complexity' to refer to the interaction of system components that results in holism and synergy,

with key attributes of such aggregate complexity being internal relationships, internal structure (subsystems), relationships with the environment, learning and memory, emergence, and change and evolution (Manson, 2001).

One of the approaches used in the study of complex systems is Agent-Based Modeling. An Agent-Based Model (ABM) consists of multiple agents acting upon their individual objectives (An, 2012). ABMs are constructed in a bottom-up manner and allow us to compute the aggregate and large scale results of the interactions of agents with each other and with their environment (Crooks, Castle, & Batty, 2008). As such, ABMs have been used in a wide variety of disciplines. Some examples of these applications include innovation diffusion (Kiesling, Günther, Stummer, & Wakolbinger, 2012), theory of cooperation (Axelrod, 1997), automated negotiation (Sanchez-Anguix, Tunalı, Aydođan, & Julian, 2021), recommender systems (Amato, Moscato, Moscato, Pascale, & Picariello, 2020), migration (Arnoux Hebert, Perez, & Harati, 2018), urban segregation (T. Anderson, Leung, Dragicevic, & Perez, 2021; T. Anderson, Leung, Perez, & Dragićević, 2021; Perez, Dragicevic, & Gaudreau, 2019) epidemiology (Perez & Dragicevic, 2009), forest ecology (Perez & Dragicevic, 2010), and species distribution (Gaudreau, Perez, & Harati, 2018). ABMs have also been applied in sustainable development studies, such as in urban planning (Li & Liu, 2008), sustainable transportation (Maggi & Vallino, 2021), circular economy (Yazan & Fraccascia, 2020), and in problems related to the tragedy of the commons (Bristow, Fang, & Hipel, 2014; Cialdini et al., 1990).

ABMs have been extensively used in studies of social norms (Savarimuthu & Cranefield, 2011). Some models apply the Belief–Desire–Intention (BDI) framework in the decisions of their agents (Afshar Sedigh, Purvis, Savarimuthu, Frantz, & Purvis, 2021; Afshar Sedigh, Purvis, Savarimuthu, Purvis, & Frantz, 2021). In the BDI framework, agents have mental attributes of belief, desire, and intention, which indicate their state in terms of information, motivation, and deliberation for action, respectively (Bratman, 1987; Rao & Georgeff, n.d.). Some of the mechanisms employed in agent-based normative simulations are leadership (Boman, 1999), learning by imitation (Lindström, Jangard, Selbing, & Olsson, 2018), machine learning and reinforcement learning (Shoham & Tennenholtz, n.d.), norm recognition (Andrighetto, Campenni, Cecconi, & Conte, 2010), and reputation (Hales, 2002). In a review of the literature, Hollander and Wu (2011) identify areas for research and improvement. Some of those areas are

norm creation and ideation, alternatives for social sanction, and the verification and validation of models (Hollander & Wu, 2011).

Given the above context, our focus in this study is on the creation and emergence of a new descriptive norm in a setting with central leadership, where there are rewards for performing the behavior that the leadership promotes, but there are no sanctions or punishments for the non-performers. The reward in this setting is reputation and recognition as a responsible member of the society. We take an agent-based simulation approach to explore the possibility of norm emergence in such a setting.

Our conceptual framework is as follows: several user agents act upon self-interest, while a governing agent requests the user agents to take a costly action. There is no force in the governing agent's request. If the user agents cooperate with the governing agent, then the governing agent acknowledges them by giving them a 'responsible user' label. The governing agent can choose what action it shall request users to do—an action that is easy for all users but useless for the governing agent, or an action that is difficult for users and desirable for the governing agent. In the former case, the governing agent gives free 'responsible user' labels to all users. In the latter case, the user agents estimate the benefit of having the label. They do so by considering if the label makes them unique in their group, and if being unique in owning the label has any value. Such value is zero at first and increases with the exposure of the group to the label over time. Ultimately, user agents compare their estimated benefit of gaining the label with their own perception of the cost of the action they are asked to do. This way, they decide if they will cooperate with the governing agent. These actions and interactions occur in each time step. The definition of our conceptual framework was inspired by a work of Bone and Dragičević (Bone & Dragičević, 2010), wherein user agents are logging companies in a forest, and in each time-step they consider cooperating with a conservationist agent, though with different interactions and algorithms from our model.

In terms of the Belief–Desire–Intention (BDI) framework (Rao & Georgeff, n.d.), our model's user agents' belief is composed of two parts: the information they have about the last known percentage of the users that participated in the costly behavior, and the information they have concerning the number of responsible user labels awarded since the beginning of the run.

The user agents' desire is to have a good reputation while avoiding costly decisions. The user agents' intentions are the decisions they make in response to the governing agent's requests.

Reinforcement Learning (RL) algorithms are a group of Machine Learning algorithms that are based on self-evaluation. RL algorithms do not know the correct answer to the problem at hand, but they can learn to improve themselves from the differences between the results of their own efforts (Alpaydin, 2014). An RL algorithm has a policy that prescribes an action for each state. In this sense the policy is a function. With each action, there comes a reward and a subsequent state. RL algorithms take note of rewards that are gained from various (state, action) pairs, and update their policies in such a way that the sum of rewards weighted by their time-values is maximized (Canese et al., 2021). RL algorithms are suitable for the problem of our study, as our model's governing agent searches for a sequence of decisions to maximize a reward, which in our model is the proportion of user agents that cooperate with the governing agent. Because of their relevance to problems involving repeated decision making, RL algorithms have been used in a variety of simulations of social systems (Angourakis, Santos, Galán, & Balbo, 2015; Chan & Steiglitz, 2008; Okdinawati, Simatupang, & Sunitiyoso, 2017) as well as social-ecological systems (Bohensky, 2014; Bone & Dragičević, 2010; Rasch, Heckelei, Oomen, & Naumann, 2016). In a similar fashion, we used RL algorithms in our model.

Within the above framework, our objectives are to answer the following questions:

1. Can the actions of agents in the above setting result in the emergence of a behavior norm in the user agents, such that the user agents compete for social status and cooperate with the governing agent despite the costly action they are asked to do?
2. How can the governing agent find a sequence of choices that facilitates or hastens the emergence of the above behavior norm?

1.2. Materials and Methods

To answer the questions of this study we adopted an agent-based simulation approach. First, we built a model of interactions of user agents and the governing agent. In the model, we included algorithms for a governing agent to guide the user agents towards the desired norm of behavior. Next, we performed tests on the model, with and without the governing agent's algorithms. To gain insight about the emergence of the intended norm of behavior, we planned

model runs without a purposeful intervention from the governing agent. This allowed us to become familiar with the state of possible outcomes of repeated actions of the user agents. Then, we tested the model with purposeful interventions with a governing agent that was equipped with several algorithms. This allowed us to compare different algorithms against each other and identify algorithms and parameters that lead to the emergence of the desired behavior norms faster than other algorithms and parameters. Finally, we ran the model several times with random interventions by the governing agent to construct a baseline for comparison with the best simulations. In this section we describe the design of the model, the algorithms, and the tests of performance of the simulations.

1.2.1. Overview, Design Principles, Details

The model description follows the ODD (Overview, Design concepts, Details) protocol (Grimm et al., 2006, 2010, 2020), which serves as a standard for communication of information about Agent-Based Models. In addition, the model description is inspired by the ODD + D protocol (Müller et al., 2013), which is an adaptation of the ODD protocol for describing human decisions in Agent-Based Models.

1.2.1.1. Purpose

This ABM is an abstract model of interactions of entities and emergence of a particular social behavior among them. Using this model, we intend to, firstly, obtain an insight into the emergence of social behavior that is costly for individuals, and secondly, examine if such emergence can be facilitated with appropriate learning algorithms.

1.2.1.2. Entities, State Variables, and Scales

Entities of this model are three classes of agents: several user agents, a governing agent, and a registrar agent. State variables of user agents are named *threshold* and *decision*. Each user agent's *threshold* is a real number between 0 and 1, which is predefined at the beginning of each simulation, remains constant throughout the simulation, and is visible to that user agent alone. User agents' *decisions* are binary variables that change throughout the simulation and are visible to all agent classes upon request. State variables of the governing agent are named *signal*, *state*, *Q*, and *policy*, which change throughout the simulation. Except for *signal*, all other variables of the governing agent are known to itself only. *Signal* is a binary variable. *State* is a two-dimensional variable with non-negative integer values. *Q* is a table with a real value for each *state* and *signal* combination. *Policy* is a table with a real number between 0 and 1 for each *state*.

The registrar's state variables are named *nLast* and *nSum*, which are non-zero integers that change throughout the simulation and are visible to all agent classes upon request. User agents and the governing agent are the main entities of the model. Registrar is an auxiliary agent that is meant to make the model easier to understand and serves as a mediator of information. This abstract model does not have a spatial dimension, and time in the model is measured with dimensionless time steps.

1.2.1.3. *Process Overview and Scheduling*

Each simulation run consists of a number of episodes. Each episode consists of a number of time-steps. Time in this ABM is modeled as discrete time-steps. In each time-step, the agents act as described in Figure 1.1.

```

Governing agent:
  Ask Registrar to report nLast and nSum
  Produce signal
  Ask Registrar to run a step

Registrar (step() function):
  Ask Governing agent to report signal
  If signal is 0 then {
    give promotional responsible agent labels to all User agents
  }
  else {
    ask User agents to report their decisions
    count the number of User agents whose decision is 1
  }
  Update nLast and nSum

User agent:
  Ask Registrar to report nLast and nSum
  Produce decision

```

Figure 1.1. Interactions between three classes of agents in one time-step. Class names are highlighted. Variable names are shown in italics.

At the beginning of each new episode of time-steps, the variables *nLast* and *nSum* are set to zero, and user agents forget their memory.

1.2.1.4. *Design Concepts* **Basic Principles**

In this ABM the governing agent does not enforce its authority over user agents. Rather, it offers them ‘responsible user’ labels in return for cooperation with the governing agent. User agents see improved social status as the benefit of being recognized as a ‘responsible user’. The basis for this idea is the assumption that individuals have a motivation for better social status, and that they may take actions that cost them money if their peers and neighbors do so (C. Anderson, Hildreth, & Howland, 2015; Lazaric et al., 2020; Nolan, Schultz, Cialdini, Goldstein, & Griskevicius, 2008; Tascioglu, Eastman, & Iyer, 2017).

Emergence

Decisions of user agents are made individually. Emergence of a pattern of such decisions that is costly to the individuals will be an unexpected phenomenon.

Adaptation

The governing agent adapts its Q based on results of each step, and accordingly calculates a new *policy* for its actions.

Objectives

The objective of the governing agent is to increase the ratio of cooperating users when it makes signals of 1. The target cooperation ratio is 0.5 in this model. The governing agent aims to reach a state with target cooperation ratio as soon as possible. The reward for the target state is defined as 0, and the reward for all other states is defined as unity minus the ratio of cooperating users. Therefore, the governing agent’s reward at each step is between -1 and 0 . User agents react to their perceived conditions by comparing the benefit of the said cooperation against their *thresholds*, which represent the cost to each user of cooperation with the governing agent.

Learning

The governing agent learns to adjust its behavior based on responses that it observes in the user agents. To this end, the governing agent uses RL algorithms. For each RL algorithm there is a separate model. In the RL algorithms, the governing agent stores the value of each action taken at each *state* in its table, Q . From Q it extracts *policy*. *Policy* recommends an action at each *state*. Actions of the governing agent are the *signals* it produces. In the next time-step,

based on the outcome of its action, which is the observed ratio of cooperation of user agents, the governing agent updates its Q and repeats this loop.

Prediction

In each time-step, the governing agent predicts the present value of the sequence of future rewards of the actions that it may take. This prediction is made based on the results of previous time-steps, and it is stored in Q . As such, Q is the basis for both learning and prediction in the governing agent.

Sensing

The governing agent and user agents read $nLast$ and $nSum$ from the registrar. The registrar reads *signal* from the governing agent and *decisions* from user agents.

Interaction

The interaction of the governing agent with user agents is through *signal*. If the governing agent produces a *signal* of 0, it is asking for a task that has no cost to the users, hence giving ‘responsible user’ labels to all users at no cost. If it produces a *signal* of 1, it is asking the users to cooperate in a task that is costly to them. In this case, the response of each user is its *decision*. If the user produces a *decision* of 0, it is not cooperating with the governing agent. If the user produces a *decision* of 1, it is cooperating with the governing agent despite the costly demand of the governing agent.

Stochasticity

The governing agent produces its *signals* using its *policy*, which is stochastic. In the governing agent’s *policy*, the probability of recommendation of each action is the ratio of its estimated value to the sum of estimated values of all possible actions. In addition, *thresholds* of user agents are defined at the beginning of the simulation as random numbers with given mean and standard deviation.

Collectives

Individual *decisions* of each user agent affect future *decisions* of itself and other user agents. Other than that, there is no connection between the user agents.

Observation

Throughout each episode, the governing agent stores in its temporary memory the sequence of rewards that it receives in each time-step. At the beginning of the next episode, this part of its temporary memory is erased. At the end of the final episode of each run, the sequence of rewards is stored in a file as output. The reason for choosing the final episode is that as learning happens throughout the simulation, the governing agent's performance improves in each episode. Therefore, the final episode represents the outcome of learning in the model.

Heterogeneity

User agents are heterogenous in their decision thresholds.

Individual Decision Making

All agents make decisions in the model. In each iteration, the object of decision of the governing agent is to choose between (i) requesting a costly behavior from user agents, and (ii) requesting an easy behavior from user agents. The governing agent gives recognition labels to cooperating user agents. In iterations where the governing agent requests the easy behavior, all user agents unconditionally cooperate with the governing agent and receive the 'responsible user' labels. In iterations where the governing agent requests the costly behavior, the object of decision of the user agents is to choose between accepting and rejecting the governing agent's request. The objective of the governing agent is to encourage at least half of user agents to perform the costly behavior. Decisions of the governing agent affect decisions of user agents and vice versa. Moreover, decisions of each user agent affect future decisions of itself and other user agents. Within the same time-step, the decision of one user agent does not affect decisions of other user agents. The governing agent's decision policy is probabilistic, and in each state recommends an action. The governing agent is equipped with RL algorithms. User agents do not have learning or optimization capabilities. Instead, the basis of decision making of each user agent is a simple if-statement. User agents calculate the utility of the 'responsible user' label by considering (i) the *uniqueness* that the label will give them, and (ii) the *value* of the label in their agent society. They calculate *uniqueness* based on last known cooperation ratio of user agents with the governing agent; and they calculate *value* based on the number of times the 'responsible user' label has been presented in their agent society since the beginning of the run. When some

user agents begin performing the costly action, that behavior might become a norm. This emerging norm influences future decisions of user agents. User agents value being recognized with a ‘responsible user’ label. There are no social sanctions or other punishments for user agents who do not follow the emerging norm.

1.2.1.5. Initialization

At the beginning of each run, the registrar’s $nLast$ and $nSum$ are zero. Additionally, the governing agent’s Q table is filled with random values.

1.2.1.6. Input Data

The model does not use input data to represent time-varying processes.

1.2.1.7. Submodels

In RL algorithms, in order to assess policies and find a pathway to improving them, a function is used that allocates a value to each (state, action) pair. In RL literature this function is known as Q (Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 2018). In turn, Q is used to update the policy. Different RL algorithms are distinguished in their timing and method of updating Q . We used a class of RL algorithms known as Temporal Difference (TD) learning algorithms. In TD algorithms, learning occurs at each time step. That is, with every action that is taken, its reward and its subsequent state are used to update Q and policy, so that the next action is prescribed with improved knowledge of the behavior of the system (Sutton & Barto, 2018). Below, we describe six different TD algorithms which we used. Each of these algorithms defined a submodel in our work. In these descriptions we assume that in state S , the algorithm’s policy p prescribes action A . Taking this action results in reward R and subsequent state S' . The next action will be A' . In all cases, the present value of a future earning is calculated using a future discounting rate, γ . Moreover, a learning rate, α , is applied to the correction term before updating Q . All the descriptions and formulas are from Sutton and Barto (Sutton & Barto, 2018). For flowcharts of these algorithms, see the Data Availability section.

SARSA

In SARSA, the next action A' is identified as $p(S')$. Then, assuming that Q leads the system from pair (S, A) to (S', A') , the previous assessment of Q is corrected. This correction accounts for the reward R as well as the value of the future pair (S', A') . Future discounting rate

γ is used to calculate the present value of that future pair. Equation (1.1) summarizes this description:

$$Q(S, A) = Q(S, A) + \alpha[R + \gamma \times Q(S', A') - Q(S, A)] \quad (1.1)$$

Q-Learning

Another TD algorithm, known as Q-Learning, looks at Q after taking action A and identifying subsequent state S'. Then, among all pairs (S', a) that are registered in Q for the new state S' and all possible actions, the algorithm selects the one with the maximum value, and uses it to correct Q. These operations are summarized in Equation (1.2):

$$Q(S, A) = Q(S, A) + \alpha[R + \gamma \times \max_a \{Q(S', a)\} - Q(S, A)] \quad (1.2)$$

Expected SARSA

Another TD algorithm, known as Expected SARSA, proceeds similar to Q-Learning up to the correction of Q. In that stage, Expected SARSA considers all (S', a) pairs and calculates their average value. Equation (1.3) describes this update process:

$$Q(S, A) = Q(S, A) + \alpha[R + \gamma \times \sum_a p(a|S') \times Q(S', a) - Q(S, A)] \quad (1.3)$$

where the summation is performed over all actions a.

Double Learning Methods

Corresponding to the above three methods, there are more complicated methods that are called Double Learning algorithms, because they involve two Q tables. In each time-step, one of the Q tables is selected randomly and updated using the other one. The following formulas describe this concept. In each set, only one of the two formulas is performed in each time-step. The formulas for update of Q tables of the Double SARSA, Double Q-Learning, and Double Expected SARSA algorithms are as shown in equation pairs (1.4) and (1.5), (1.6) and (1.7), (1.8) and (1.9), respectively.

Double SARSA:

$$Q_1(S, A) = Q_1(S, A) + \alpha[R + \gamma \times Q_2(S', A') - Q_1(S, A)] \quad (1.4)$$

$$Q_2(S, A) = Q_2(S, A) + \alpha[R + \gamma \times Q_1(S', A') - Q_2(S, A)] \quad (1.5)$$

Double Q-Learning:

$$Q_1(S, A) = Q_1(S, A) + \alpha[R + \gamma \times \max_a \{Q_2(S', a)\} - Q_1(S, A)] \quad (1.6)$$

$$Q_2(S, A) = Q_2(S, A) + \alpha[R + \gamma \times \max_a \{Q_1(S', a)\} - Q_2(S, A)] \quad (1.7)$$

Double Expected SARSA:

$$Q_1(S, A) = Q_1(S, A) + \alpha[R + \gamma \times \sum_a p(a|S') \times Q_2(S', a) - Q_1(S, A)] \quad (1.8)$$

$$Q_2(S, A) = Q_2(S, A) + \alpha[R + \gamma \times \sum_a p(a|S') \times Q_1(S', a) - Q_2(S, A)] \quad (1.9)$$

1.2.2. Model Parameters

The model includes several parameters, which we have divided into two groups: those that are parameters of the problem, and those that are parameters of the algorithm. Parameters of the problem are the number of agents (n), mean (μ) and standard deviation (σ) of the decision thresholds of the population from which user agents are selected, and future discounting rate (γ). Parameters of the algorithm are the rate of exploration vs. exploitation (ϵ) and learning rate (α). Appendix 1.A lists these parameters and their values in the simulations. These parameters and values produce 810 different combinations for each of the 6 algorithms. Therefore, a total of 4860 distinct problem and solution/algorithm settings are possible. We produced simulations for each of these settings. In the analysis of the results, we separated the two parameter groups, taking note of 54 combinations of problem parameters and 15 combinations of ϵ and α for each of the 6 algorithms, which produced 90 combinations of solution/algorithm parameters.

Among the problem parameters, n , μ and σ are used in the making of user agents. There are 18 possible combinations of values of these parameters. For each of those combinations, we made 50 sets of user agents. Each set was defined by selecting n thresholds from a normally distributed population with mean μ and standard deviation σ . These sets of thresholds were saved and used in all simulations that shared their respective values of n , μ and σ .

In addition to the above parameters, the model has some parameters that we did not vary in simulations. Specifically, the number of training episodes for each run, which was set to 4000;

the number of time steps per episode, which was set to 17; and the number of levels of the two-dimensional state variable, which was set to $\lceil n/2 \rceil$, or integer ceiling of half of user agents, for each dimension. The rationale for this choice was to enable the governing agent to distinguish states with different levels of cooperating user agents. The target state is when the number of cooperating user agents reaches or exceeds $n/2$.

1.2.3. Simulation Experiments

1.2.3.1. Quantifying Model Performance

In each simulation and for each set of parameters, the model trained itself in 4000 episodes. In the final episode of each run, after long sequences of updates and improvements, the model policy and Q were at their best. Therefore, results of the final episode of each run were used to assess that run. We used two measures to quantify model performance. In most of our analyses we calculated the mean of the rewards of the time steps of an episode as the score of that episode. Our rationale for this choice was that it is a measure of cumulative rewards. In another part of our analyses we used the rewards of the final time step as the score of the episode. The rationale for this choice was that it shows the state of the system at the end of the simulation and allows us to answer questions such as to what extent the desired state was achieved.

1.2.3.2. Space of Outputs

The user agents in our model react to the *signals* produced by the governing agent. In order to better understand the process of the study, we produced the space of all possible outputs of the model, by producing all possible sequences of decisions of the governing agent and feeding those sequences to the user agents. This way, we constructed a large binary tree of all binary strings of length 16. The length of these strings is one less than the length of the episode because in episodes of length 17 the algorithm makes 16 decisions. We kept the size of episodes to this level because for larger episode lengths, the scope of outputs would become exponentially larger and more difficult to manage, from the point of view of computation. Taking note of the score of each of the produced 2^{16} chains, we obtained insight about the space of outputs, which allowed us to realize which scores are rare and what conditions favor higher scores.

1.2.3.3. Comparison of Simulations with Each Other

The combination of 54 problem parameters and 90 algorithm/solution parameters produced 4860 unique combinations of parameters and algorithms to run the model in. We ran

the model 50 times in each of these combinations of parameters and algorithms. We then took note of scores of runs as the mean reward per time-step of the 50 simulations in each run. This produced a dataset that we organized as a matrix with 54 columns and 90 rows. Based on recorded scores, rows and columns of the scores' matrix were analyzed and ordered with hierarchical clustering. We then identified the ranks of the values within each column of the matrix. The ranked matrix showed a comparison of the 90 algorithm/solution parameters against one another. Both matrices were plotted as heatmaps. Using these visualizations enabled us to find groups of simulations with higher scores. This visual finding was confirmed by marginal sums of the matrices. Through this process we were able to select algorithms that performed better than the others in most of the problem parameter settings.

1.2.3.4. Comparison of Simulations with a Reference Baseline

In addition to comparing algorithms with each other, we compared the selected algorithms against a baseline. The reason for this comparison was to note what would happen without the algorithms, and so assess the role of the RL algorithms in the achievement of results. To make this basis for comparison, we ran another series of simulations with the same problem settings and the same thresholds for user agents but without RL algorithms for the governing agent. Instead, in these simulations the governing agent produced *signals* randomly in each time-step. We then compared the results of this new model, which we call the baseline, with the selected RL algorithms.

1.2.4. Implementation

The model was developed and run using Java Repast Symphony 2.7 (M. J. North et al., 2013). Simulation results were analyzed and visualized using R statistical software (R Core Team, 2019) and its packages ggplot2 (Wickham, 2016), scales (Wickham & Seidel, 2020), and signs (Wolfe, 2020). For model code and results, see the Data Availability section.

1.3. Results

1.3.1. Simulations without RL

Figure 1.2 shows histograms obtained by simulating the actions of various sets of user agents given all possible sequences of signals by the governing agent. In these simulations, no RL algorithm was used for the governing agent. Rather, all binary sequences of length 16 were generated and tried on the user agents. As such, the results of these simulations depict the space of outcomes of all possible policies. Each sequence of decisions was tried on 100 sets of user

agents with similar characteristics. Therefore, each histogram shows the distribution of scores of 6,553,600 episodes. The score of each episode is calculated as the mean reward per time-step of that episode.

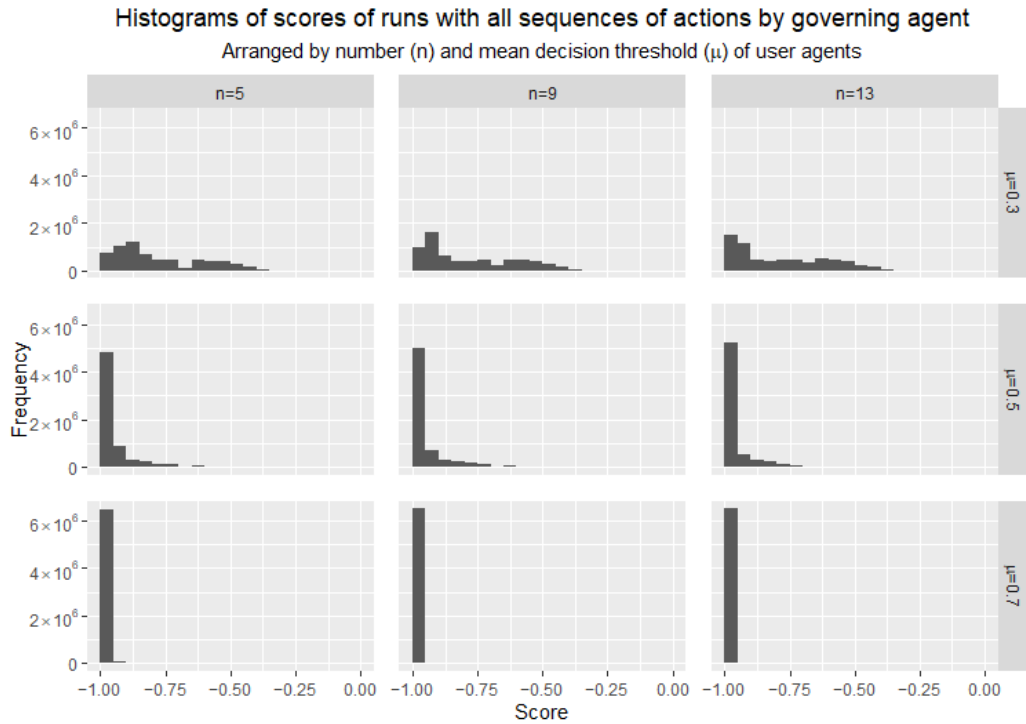


Figure 1.2. Histograms of overall scores of episodes in reference dataset, ordered by number and mean decision threshold of user agents. Each histogram summarizes 6,553,600 data points. Each data point is an episode of 17 time-steps. The score of each episode is its mean reward per time-step. Rewards are real numbers between -1 and 0 , and they are calculated based on the ratio of user agents that cooperate with the governing agent in each time-step.

As shown in Figure 1.2, the emergence of the desired norm of behavior is highly dependent on the mean cost–benefit decision threshold of the user agents. In simple terms, the more costly the behavior, the less likely it is to become a trend in the society. This is especially evident in the simulations with a mean decision threshold of 0.7 for the user agents, where nearly all episodes ended with the minimum score, and cooperation of the user agents with the governing agent was a rarity. To a lesser extent, this happened in the simulations with a mean decision threshold of 0.5 as well. The histograms of these runs show lower peaks and more dispersed distributions of scores, though their modes are still at the minimum score. On the other hand, in the simulations with a mean decision threshold of 0.3 , the scores are distributed more evenly. In two of the three histograms of these simulations, the mode is not at the minimum score. In fact, these histograms show that the number of user agents has an inverse effect on the mode of scores.

1.3.2. Simulations with RL

Figure 1.3 shows a heatmap of scores of 4860 sets of simulations. This heatmap is composed of 54 columns and 90 rows. The columns and rows of this figure correspond to problem parameters and algorithm/solution parameters, respectively. Specifically, each column is for one unique combination of number of user agents, mean and standard deviation of decision threshold of the population of user agents, as well as the future discounting rate. Each row is for a unique combination of the RL algorithm, its rate of exploration vs. exploitation, and its learning rate. As such, each column represents a problem, and each row is a solution to that problem. Appendix 1.B includes parameter combinations and their respective codes, which are assigned to columns and rows of the heatmap figures. Each pixel in the heatmap represents the mean score of 50 simulations with the same problem parameters and algorithm/solution parameters. The score of each simulation is the mean reward per time-step of that simulation. Column numbers and row numbers are printed on the margins of the heatmap. The order of columns and rows was determined through hierarchical clustering, using column sums and row sums, respectively. Dendrograms of the hierarchical clustering of columns and of rows are shown in the figure. Larger copies of these dendrograms as well as lists of parameters of rows and columns are available in the Data Availability section.

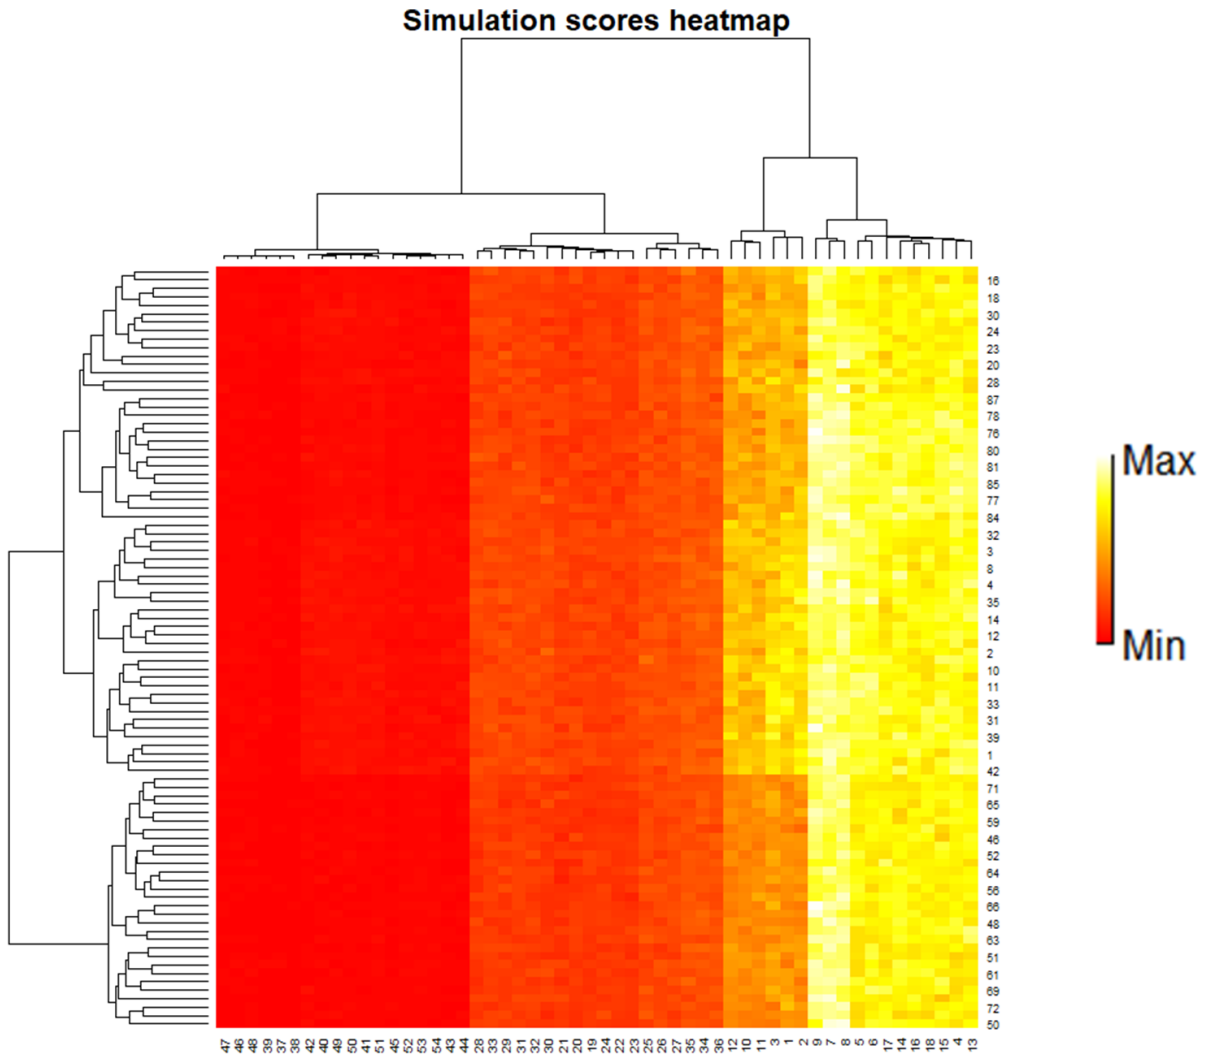


Figure 1.3. Heatmap of mean scores of simulations. Each row is a unique combination of algorithm settings. Each column is a unique combination of problem parameters. Rows and columns are ordered using hierarchical clustering, as shown in their respective dendrograms. Each pixel represents the mean score of 50 simulations with its respective row and column settings. Simulation scores are mean rewards per time-step.

There are three distinct vertical bands in the heatmap of Figure 1.3. These correspond to the three tested values for mean decision thresholds of user agents. The left-most vertical band, shown in deep red, corresponds to mean decision threshold of 0.7. The middle band, which shows a variety of red and orange colors, corresponds to mean decision threshold of 0.5, and the right-most band, with the highest variety of colors from orange to white, corresponds to mean threshold of 0.3. The dendrogram of the columns shows that the scores of the two bands on the left are more similar to each other, while the scores of the right-most band are in a different cluster, which is confirmed by the colors of the heatmap.

The colors of the heatmap of Figure 1.3 are proportional to the values of the pixels of the heatmap, with the lowest value colored red, and the highest value colored white. The visualization in this figure shows that problem parameters have a strong influence in the results. However, our goal is to identify the best solutions to the problems, and from this figure it seems that the differences between the results of various solutions are smaller than the differences between problems. As such, it is not easy to distinguish between different solutions in this figure.

In order to compare different solutions, we prepared Figure 1.4. This figure is the result of column-ranking of the heatmap of Figure 1.3. As such, the values in each column of the heatmap of Figure 1.4 range from 1 to 90, with 1 corresponding to the lowest score and 90 to the highest score in the respective column in the heatmap of Figure 1.3. In this way, the difference between problem settings is eliminated from Figure 1.4 and it is only the difference between row values, that is, algorithm/solution parameters, that causes variations in this heatmap. Each of the 54 columns is a test problem. For each test problem, 90 solutions are given, and they are ranked according to their scores. An ideal solution is one that has high ranks in all or most of the tests. In order to more easily understand the figure, the orders of rows and columns are the same as those of Figure 1.3.

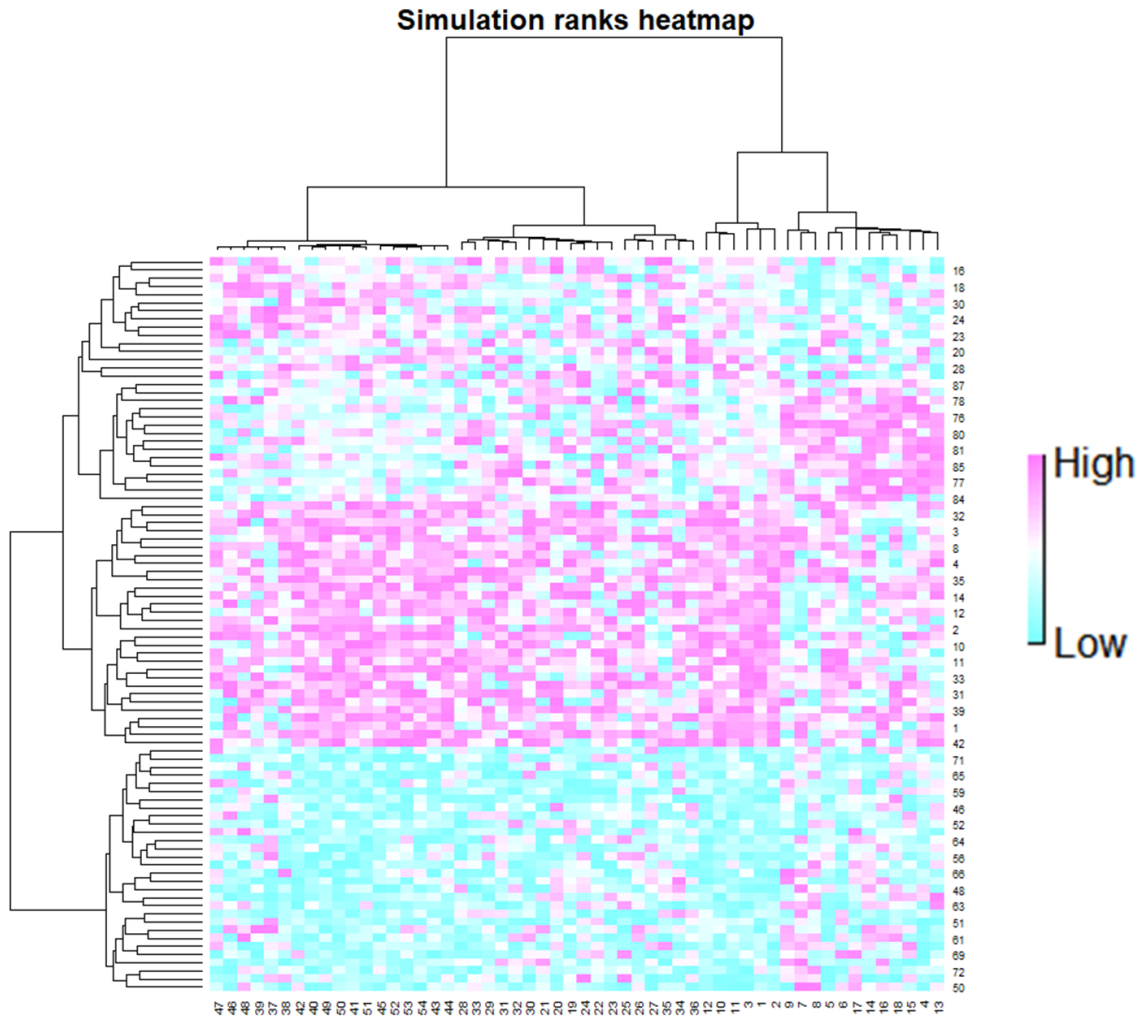


Figure 1.4. Heatmap of ranks of simulations. Each row is a unique combination of algorithm settings. Each column is a unique combination of problem parameters. For each column, the row with the highest mean score of simulations is given the highest rank. The order of rows and columns in the ranks heatmap is the same as that of the mean scores heatmap.

It can be seen that the heatmap of Figure 1.4 is divided into three horizontal zones, with the lowest zone having the lowest ranks, and the middle zone having the highest ranks. The lower zone, with lowest ranks, corresponds to the RL algorithms Q-Learning and Expected SARSA. The middle zone, with the highest ranks, corresponds to the RL algorithms Double SARSA and Double Expected SARSA. The upper zone of the figure, which contains solutions with middle ranks, corresponds to the RL algorithms SARSA and Double Q-Learning. As shown in the dendrograms and confirmed by the colors of the pixels, the scores of the two upper zones are more similar to each other, whereas the lowest zone is in a different cluster.

It is noticeable that within the upper zone of the figure, there is an accumulation of magenta pixels in the bottom-right and in the top-left. We mentioned that the left side of the figure corresponds to problems with mean user agent decision threshold of 0.3, and the right side corresponds to the mean decision threshold of 0.7. We also noted that the latter is a tougher challenge for the RL algorithms because in its space of decisions, rewards are rare. It may seem reasonable to assume that the solutions with higher ranks in the tougher problems are more successful than others. The row numbers of the two groups of solutions show that in this zone, the Double Q-Learning algorithm performs better than the SARSA algorithm.

In all, in Figure 1.4 the Double Learning algorithms showed superior performance. We looked at the row sums of the heatmap in order to identify the best algorithms with their parameters. The highest-ranking algorithms in the 54 problems were: (1) Double Expected SARSA, with an exploration rate of 0.1 and a learning rate of 0.2; (2) Double Expected SARSA, with an exploration rate of 0.2 and a learning rate of 1.0; and (3) Double SARSA, with an exploration rate of 0.01 and a learning rate of 1.0.

Figure 1.5 shows the spread of the rewards obtained in various simulations with the RL algorithms Double SARSA and Double Expected SARSA. Each curve in this figure represents the mean rewards of 50 simulations with similar parameters. For each algorithm, 810 different parameter settings were tested. As seen in the figure, the two algorithms show similar variations in results. There are no areas of the plot that are particularly filled with curves of only one of the two algorithms. In this sense, we cannot visually distinguish between the two algorithms. Appendix 1.C includes flowcharts of these two algorithms.

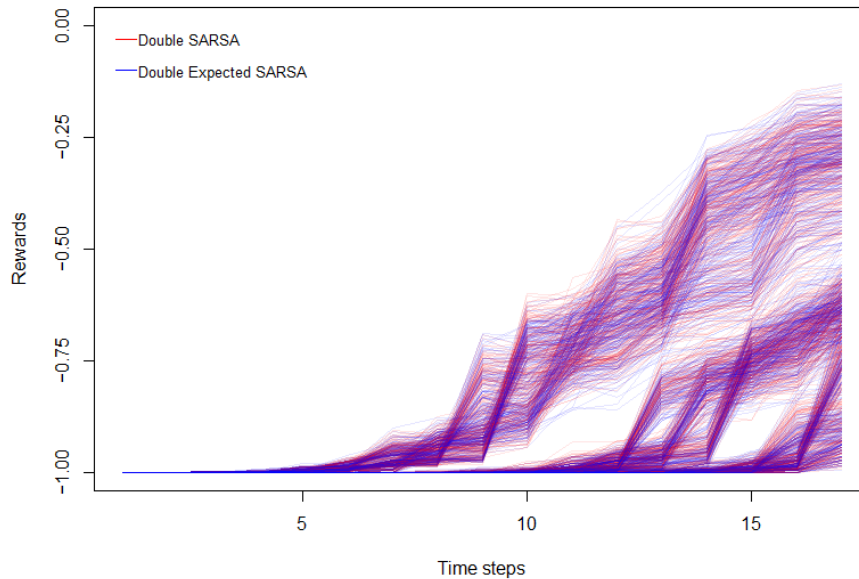


Figure 1.5. Rewards versus time-steps for Double SARSA and Double Expected SARSA algorithms. For each of the two algorithms, 810 curves are shown. Each curve represents a parameter setting for its respective algorithm. For each parameter setting, 50 simulations were run, their mean score was plotted at each time-step, and a line segment was drawn between the score points of consecutive time steps to produce a curve.

Three strands of curves are visible in the plots of rewards of simulations in time steps. These correspond to the three thresholds for decisions of user agents: the lower the thresholds, the higher the rewards. It is noticeable that in simulations with the mean decision threshold of 0.3, higher rewards emerge between the 5th and 10th time steps. Such time of emergence of higher rewards is delayed to between the 10th and 15th time steps in simulations with mean decision threshold of 0.5. The rewards of simulations with mean decision threshold of 0.7 emerge later, after the 15th time step. This shows that as the users agents' decision threshold increases, it takes longer times for the RL algorithms to cause the user agents to cooperate with the governing agent.

In Figure 1.6 we compared the scores of the selected RL algorithms against a baseline. The scores in these histograms are rewards of the 17th time step of 40,500 episodes for each of the RL algorithms and the baseline. In the baseline run, *signals* were produced by the random decisions of the governing agent, and they were given to the user agents. This represents a case

where the governing agent does not have an algorithm. Therefore, this case serves as a basis for comparison against the cases where the governing agent does have an algorithm. Recall that the rewards were defined as unity minus cooperation ratio if cooperation ratio is below 0.5, and zero otherwise. The histograms below show this matter, as they include no rewards between -0.5 and 0 . The histograms show two peaks of frequencies at the highest and lowest ends of score range. Clearly, in comparison with the baseline, the RL algorithms have lower frequencies of low scores and higher frequencies of high scores.

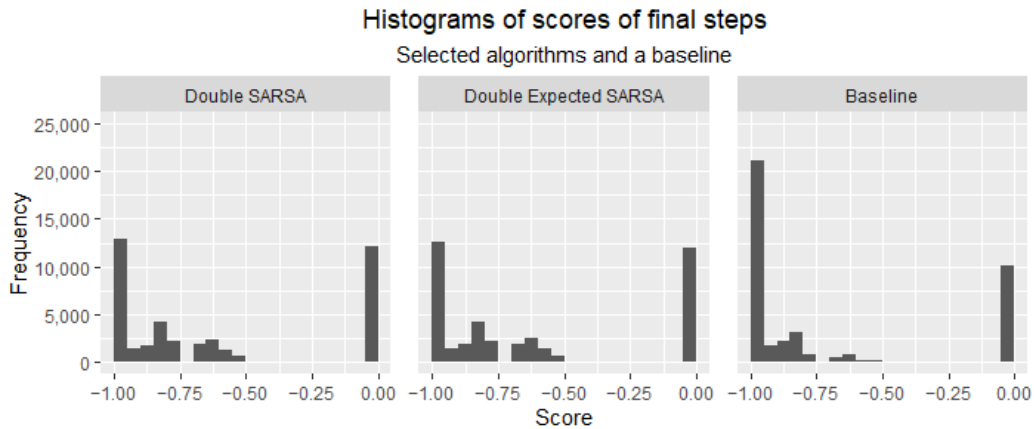


Figure 1.6. Histograms of scores at the final time step for RL algorithms Double SARSA and Double Expected SARSA as well as a random baseline. Each histogram summarizes 40,500 data points.

The choice of rewards of the 17th step as the score in Figure 1.6 was in order to show what happens to the group of user agents at the end of the simulation. It indicates to what extent the target state was achieved throughout simulations. The histograms show that compared to the baseline, the RL algorithms were more successful in encouraging the cooperation of the user agents with the governing agent.

Figure 1.7 shows another comparison of the selected algorithms with the baseline. This figure uses the same simulation and baseline scores as Figure 1.6, but it separates data according to the mean value of user agents' decision threshold. The boxplots of Figure 1.7 show the spread of rewards gained at the final time-step of the runs, for three values of the mean threshold (μ).

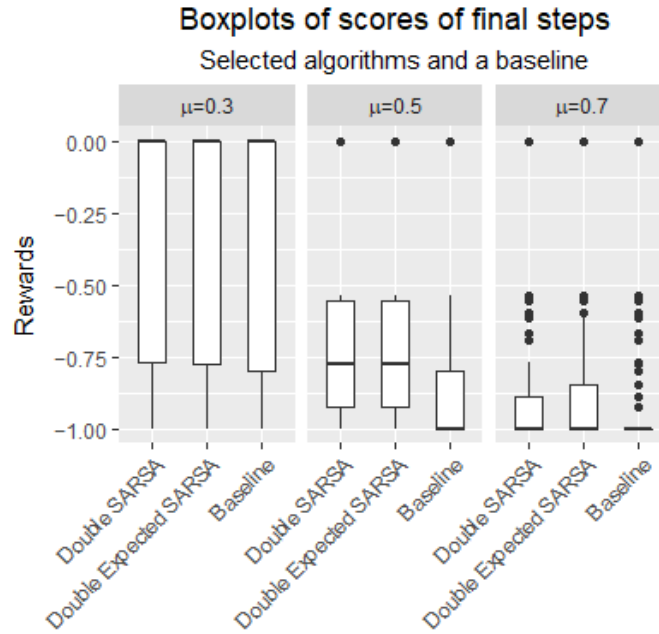


Figure 1.7. Boxplots of scores at the final time step for RL algorithms Double SARSA and Double Expected SARSA as well as a random baseline, classified by mean value of user agents' decision threshold. Each boxplot shows quartile ranges of scores of 13,500 data points.

Figure 1.7 reveals several points. Firstly, it is evident in the figure that the results are dependent on the mean decision threshold of the user agents. Secondly, it is noteworthy that at the lower threshold value ($\mu = 0.3$) half of the runs with RL algorithms, as well as the random baseline, reach the target state and obtain full reward at the final time-step. Lastly, at other threshold values ($\mu = 0.5$ and $\mu = 0.7$) the RL algorithms scored distinctly higher than the baseline. The observation that the median of the baseline dataset is the highest possible score when $\mu = 0.3$ indicates that at the lower threshold, the existence of the mechanism of recognition of 'responsible users' leads to emergence of a norm of behavior in which the user agents cooperate with the governing agent. Conversely, the observation that the median of the baseline dataset is the lowest possible score when $\mu = 0.5$ and $\mu = 0.7$ shows that in these cases it is a challenge for the RL algorithms to find a sequence of decisions for the governing agent to create the desired norm of behavior among user agents. These cases show the superior performance of the RL algorithms in comparison with the baseline.

1.4. Discussion

We started our work with a curiosity about the ability of the governing agent to use reputation as a mechanism for guiding user agents to perform the desired behavior. To that end, we explored the space of possible outcomes of interactions of user agents that consider their reputation. We also equipped the governing agent with a learning algorithm to find a successful policy for its actions. In our abstract study, we consider a policy successful if it leads to the participation of user agents in the desired behavior, such that the percentage of participating users is higher than what could be achieved by random actions. Our criterion for success was inspired by the definition of descriptive norms by Cialdini et al. (Cialdini et al., 1990), which inform us of what others do in the society. This is also in accordance with an interpretation of norms mentioned by Therborn (Therborn, 2002), in which norms inform us about the distribution of things (in our case: behaviors) and indicate what is ‘normal’.

Our simulation of the space of all possible decisions of the governing agent (Figure 1.2) showed that the emergence of the desired norm of behavior among user agents is possible, though it can be rare, depending on the parameters of the problem. We chose the length of the simulation episodes considering computation hardware limits. Substantial computational power and memory are required in the processing and internal verification of this stage, as the space of decisions grows exponentially with the number of time steps. Nevertheless, through the simulations we were able to identify information about the process being studied.

Moreover, we showed that RL algorithms could hasten and facilitate the emergence of norms of group behavior that are costly to the individuals. In particular, comparison of the results of RL algorithms with the baseline (Figure 1.7) showed that in some cases, the algorithms were able to reach results that were rare in their problem settings. On the other hand, the comparison with baseline also showed that in problem settings where the chance of the emergence of the desired behavior is high, a random baseline could reach results comparable with RL algorithms. As such, we can say that if the user agents perceive a low cost for the requested behavior, then having in place a structure in which user agents who performed that behavior are recognized and introduced to the group as ‘responsible users’ can lead to the diffusion of that behavior in the group and emergence of a new behavior norm. If, on the other hand, the perceived cost of performing that behavior is high for the user agents, then the mere existence of the recognition structure is not enough for diffusion of that behavior in the group. In these cases, a governing

agent equipped with an appropriate algorithm may be able to guide the group of user agents towards the desired behavior.

Our initial inspiration for this study comes from our field of work—sustainability—where governments are interested in encouraging individuals to adopt environmentally responsible behavior (Barr, 2003). The research presented in this paper is part of a larger project aimed at understanding the complexities of a system that is composed of social and ecological parts. Such social-ecological systems involve interactions of subsystems that are, in turn, complex (Liu et al., 2007; Ostrom, 2009). In the present study, we were able to select algorithms to use in the construction of a social-ecological model in future. We also identified parameters to use for those algorithms.

In abstraction, our model's governing agent aims to encourage our model's user agents to do something that the user agents perceive is costly for them. This is as if the governing agent was trying to sell something—in our model's case, a 'responsible user' label—to the user agents, where the user agents are not initially convinced that it is worth the price. A field of study that deals with similar problems is marketing. In fact, ABMs are applied in marketing research and are known to be useful because of their cross-scale capabilities: they build individual agents and capture results that emerge at the scale of the society (Rand & Rust, 2011). A similar point has been mentioned in the literature of innovation diffusion (Sebastiano A. Delre, Jager, Bijmolt, & Janssen, 2010; Kiesling et al., 2012). In addition, it has been noted that an individual's decision to purchase a product depends on the quality of the product and the social influence the individual receives from their peers (Sebastiano A. Delre et al., 2010). Similarly, our user agents are influenced by their society. Our model's user agents each perform a cost-benefit analysis. They assess the benefit of performing a task that is costly for them. Such a benefit is social respect. Then, the agents compare that benefit with a threshold, which represents their perception of the cost of the task. Our model's agents, however, do not receive a product in return for the cost that they pay. As such, they compare the cost only with their estimate of the value of the social status that they may gain if they pay the price for it. In a related work, Antinyan et al. (Antinyan, Horváth, & Jia, 2019) built an ABM in which each agent compares its status with the mean status of others in their social network, and decides accordingly to spend a budget to improve its own status. In a similar fashion, our user agents consider the mean status of their

group in their decision to take a costly action for improving their own status. In a different study, Shafiei et al. (Shafiei et al., 2012) built an ABM of market share of electric vehicles and stated that visibility of a new subject can help it become a trend in the society. Similarly, our user agents consider a measure of visibility of the new trend in their group, and the governing agent's actions increase visibility of the 'responsible user' label. In another ABM study about promotional activities in marketing and sales of products, Delre et al. (S. A. Delre, Jager, Bijmolt, & Janssen, 2007) concluded that timing of promotional activities has an important role in the success of a sales campaign, and inappropriate timing may cause the sales of the product to fail. In our study, decisions of the governing agent are indeed about giving free promotional 'responsible user' labels to all user agents. The RL algorithms give the governing agent a policy that prescribes when promotional labels should be given for free. The comparison of the performance of the RL algorithms with the random baseline showed that the timing of promotional offers of the label, which the algorithms prescribed, was influential in achieving results.

Our model is an abstract model that simulates interaction of agents in a hypothetical setting. There is always a concern about such abstract models and whether they are useful, as they are not connected to the real world. Moreover, without connection to the real world, questions arise about the validity of the model and the relevance of its results. Below, we address these issues.

Depending on the model's purpose, ABMs can be classified in two different types: predictive and explanatory (Castle & Crooks, 2006). The aim of predictive models is to extrapolate trends, evaluate scenarios and predict future states, whereas the aim of explanatory models, in terms of Castle and Crooks, is 'to explore theory and generate hypotheses' (Castle & Crooks, 2006). These different purposes justify different approaches. Predictive models try to be detailed enough to make a precise enough replicate of the real world, while explanatory models often involve simplifying assumptions that reduce the real world to abstractions (Castle & Crooks, 2006; Livet, Phan, & Sanders, 2014). There have been many cases where abstract models have led to better understanding of phenomena and theories. For example, Adam Smith developed a theory in which markets emerge as the result of actions of individuals pursuing their own interest (Smith, 1982). Centuries later, Gavin analyzed an abstract ABM based on Smith's

work and put Smith's theory to test with it, to find whether self-interest actions of individuals will result in increased utility overall (Gavin, 2018). For another example, Axelrod developed an abstract ABM of hypothetical agents interacting with each other in a game of repeated Prisoner's Dilemma (Axelrod, 1981) and based on that abstract model he made substantial contributions to the theory of cooperation. Another example is Schelling's segregation model (Schelling, 1971) in which he simulated spatial patterns of distribution of ethnic groups in a hypothetical urban environment. Our model, too, is abstract and aims to provide insights about emergence of certain behaviors in groups of agents. Our model is not intended to represent a real-world system. Rather, it is meant to show whether it is possible that behaviors that are costly to individuals emerge and become a norm in a group, given a mechanism of recognition of agents who perform such a behavior.

In addition to verifying our model at several stages of development, we compared several algorithms with each other in our model assessment effort (Figure 1.4). These comparisons shed light on the simulations and allowed us to distinguish more powerful algorithms and identify some sets of parameters with which the algorithms perform well in various tests. We also assessed our algorithms against a baseline (Figures 1.6 and 1.7) and showed that the identified algorithms make a difference in comparison to a case where those algorithms are not used. Moreover, by constructing the space of outputs of all possible sequences of actions (Figure 1.2) we gained an insight into the results that can be reached, and the rarity of our desired state. Through this integrated approach we put our hypothetical ABM to test, verified it, and learned about its power and its limits.

In two literature reviews, Savarimuthu and Cranefield (Savarimuthu & Cranefield, 2011) and Hollander and Wu (Hollander & Wu, 2011) noted that many authors associate norms with social sanctions and enforcement. Axelrod (Axelrod, 1986) states that in simulations, sanctions facilitate the emergence of norms because an agent's calculation of its utility is affected by the negative score of the sanctions that it might face, if it does not follow the norm. Therefore, it seems that in a setting without sanctions, norms are less likely to emerge than in a similar setting with sanctions. Our study involved a setting without sanctions, and the desired behavior still emerged among the user agents. This indicates two points: firstly, the offer of good reputation is a mechanism that contributes to the emergence of a new norm, even in the absence of sanctions;

and secondly, the governing agent's RL algorithm allows it to effectively use the reputation mechanism and promote the desired behavior. These points address a question that Hollander and Wu (Hollander & Wu, 2011) raise in their literature review, about possible alternatives to social sanctions.

Our governing agent performs the role of centralized leadership (Boman, 1999) in the emergence of a new norm in its society. The new norm is the manifestation of decisions of user agents to cooperate with the governing agent. These decisions are dependent on the user agents' thresholds for assessment of the utility of their choices. In principle, if the governing agent knew the mean value of decision thresholds of user agents, then the governing agent could adjust its actions accordingly and have an efficient policy. The governing agent could do this by repeatedly giving promotional labels to all user agents and increasing their utility, until their utility reached their decision thresholds. Then, the governing agent could ask for the costly behavior, and the user agents would find that the utility of being recognized as a responsible user is worth more than the cost of the requested behavior, so they would cooperate with the governing agent. This is in accordance with Axelrod's (Axelrod, 1986) explanation of how the desire for good reputation can lead to the emergence of a norm. However, the challenge for our governing agent is that it does not know the decision thresholds of user agents. To better understand this challenge, suppose that the governing agent underestimates the mean decision threshold of user agents. In this case, before giving sufficient promotional labels and increasing the utility of the desired behavior in the user agents, the governing agent asks for the costly behavior. As a result, the unprepared user agents do not cooperate with the governing agent. As another result, the promotional activity of the governing agent is delayed by its untimely request, and the emergence of the norm will be postponed. Now suppose another case, where the governing agent overestimates the mean decision threshold of user agents. In this case, the governing agent continues increasing the utility of the desired behavior in the user agents, without realizing when they are ready to cooperate with the governing agent. As such, the governing agent loses time in unnecessary promotional activity and does not ask user agents to perform the desired behavior. This case, as well, results in postponed emergence of the norm. The success of the governing agent, therefore, depends on the proper timing of its activities. Our governing agent's RL algorithm adjusted itself by occasionally making costly requests and checking

the response of the user agents. In this way, the RL algorithm achieved higher user agent cooperation rates than a random baseline, as evident in Figure 1.7.

The scope of this study is the evolution of the simulated society of the governing and user agents, from a situation where no user agents cooperate with the governing agent till a situation where the desired proportion of user agents voluntarily cooperate with the governing agent and perform the costly behavior that the governing agent requests. As such, our work addresses another question raised by Hollander and Wu (Hollander & Wu, 2011) regarding the early stages of norm creation and ideation. What happens afterwards is beyond the scope of this study. Nevertheless, it is worth noting as an implication that when many user agents voluntarily perform the behavior that the governing agent requests, the society of user agents may develop a tendency to take that behavior for granted and sanction those who do not participate in that behavior (Axelrod, 1986). As another implication, the governing agent may introduce new laws to enforce the newly emerged norm (Savarimuthu & Cranefield, 2011). These can be subjects of future works.

As another idea for future work, our model can be used in the study of complex systems that involve our case of Principal–Agent setting in conjunction with another phenomenon. For example, in environmental management there is typically a governing agent or entity with demands from users of an environmental resource. Such social interactions can be simulated in our model. The environmental resource, in turn, is subject to the laws of nature. If there exists a model of natural changes in the environmental resource, then by coupling that model with the model described in the present study, it is possible to simulate the changes in that social-ecological system.

1.5. Conclusions

In this study we developed an abstract ABM of the interactions of a principal and several agents in a hypothetical context where the principal offers the agents recognition and good reputation in return for their cooperation in a behavior that is costly to the agents. Our simulation results showed that in such a setting, the emergence of the desired behavior as a norm is possible. If the agents perceive that the cost of the behavior is low, then emergence of that norm is possible even without guidance of the agents by the principal. If the perceived cost of the behavior is not low, then cases of emergence of that norm are rare. However, we demonstrated

through comparison with a random baseline that RL algorithms can effectively guide the agents towards adopting the said behavior. Among the six TD RL algorithms that we tried—namely, SARSA, Q-Learning, Expected SARSA, Double SARSA, Double Q-Learning, and Double Expected SARSA—we noted that Double Learning algorithms obtained better results in the setting of this study. We conclude that with a proper learning algorithm it is possible to create norms of costly behaviors, by using recognition as a reward for participation in the behavior, even in the absence of social sanction and enforcement.

Author Contributions: Conceptualization, S.H., L.P. and R.M.-H.; methodology, S.H., L.P. and R.M.-H.; software, S.H.; validation, S.H.; formal analysis, S.H., L.P. and R.M.-H.; investigation, S.H., L.P. and R.M.-H.; resources, L.P.; data curation, S.H.; writing—original draft preparation, S.H., L.P. and R.M.-H.; writing—review and editing, S.H., L.P. and R.M.-H.; visualization, S.H.; supervision, L.P. and R.M.-H.; project administration, L.P.; funding acquisition, L.P. and R.M.-H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada through the Discovery Grant number RGPIN/05396-2016 awarded to L.P. R.M.-H. received financial support from the European Union’s Seventh Framework Programme through NEWFOREST programme number PIRSES-GA-2013-612645.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Code of the model is available at <https://github.com/s-harati/model-Cooperation> (accessed on September 9, 2021). Datasets of results of simulations as well as flowcharts of the RL algorithms used are available at <https://osf.io/jyqu7/> (accessed on September 9, 2021).

Acknowledgments: We are thankful to two anonymous reviewers who read the manuscript and provided helpful feedback towards its improvement. We are also thankful to the Université de Montréal’s International Affairs Office (IAO) for their financial support through

the International Partnership Development program, which allowed the collaboration between researchers from UdeM and CREAM.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix 1.A. Model Parameters

Table S1.1. Model parameters.

Parameter Group	Parameter	Symbol	Values
Problem parameters	Number of user agents	n	5, 9, 13
	User agents mean threshold	μ	0.3, 0.5, 0.7
	User agents' threshold standard deviation	σ	0.06, 0.08
	Future discounting rate	γ	0.1, 0.5, 0.9
Algorithm parameters	Exploration rate	ϵ	0.01, 0.1, 0.2
	Learning rate	α	0.2, 0.4, 0.6, 0.8, 1

Appendix 1.B. Parameter Combination Codes

Table S1.2. Codes and combinations of problem parameters. These codes correspond to column numbers of the heatmap figures in the text. μ and σ are the mean and standard deviation of User agent thresholds, respectively. n is the number of user agents. γ is the future discount rate.

Code	μ	σ	n	γ	Code	μ	σ	n	γ	Code	μ	σ	n	γ
1	0.3	0.06	5	0.1	19	0.5	0.06	5	0.1	37	0.7	0.06	5	0.1
2	0.3	0.06	5	0.5	20	0.5	0.06	5	0.5	38	0.7	0.06	5	0.5
3	0.3	0.06	5	0.9	21	0.5	0.06	5	0.9	39	0.7	0.06	5	0.9
4	0.3	0.06	9	0.1	22	0.5	0.06	9	0.1	40	0.7	0.06	9	0.1
5	0.3	0.06	9	0.5	23	0.5	0.06	9	0.5	41	0.7	0.06	9	0.5
6	0.3	0.06	9	0.9	24	0.5	0.06	9	0.9	42	0.7	0.06	9	0.9
7	0.3	0.06	13	0.1	25	0.5	0.06	13	0.1	43	0.7	0.06	13	0.1
8	0.3	0.06	13	0.5	26	0.5	0.06	13	0.5	44	0.7	0.06	13	0.5
9	0.3	0.06	13	0.9	27	0.5	0.06	13	0.9	45	0.7	0.06	13	0.9
10	0.3	0.08	5	0.1	28	0.5	0.08	5	0.1	46	0.7	0.08	5	0.1
11	0.3	0.08	5	0.5	29	0.5	0.08	5	0.5	47	0.7	0.08	5	0.5
12	0.3	0.08	5	0.9	30	0.5	0.08	5	0.9	48	0.7	0.08	5	0.9
13	0.3	0.08	9	0.1	31	0.5	0.08	9	0.1	49	0.7	0.08	9	0.1
14	0.3	0.08	9	0.5	32	0.5	0.08	9	0.5	50	0.7	0.08	9	0.5
15	0.3	0.08	9	0.9	33	0.5	0.08	9	0.9	51	0.7	0.08	9	0.9
16	0.3	0.08	13	0.1	34	0.5	0.08	13	0.1	52	0.7	0.08	13	0.1
17	0.3	0.08	13	0.5	35	0.5	0.08	13	0.5	53	0.7	0.08	13	0.5
18	0.3	0.08	13	0.9	36	0.5	0.08	13	0.9	54	0.7	0.08	13	0.9

Table S1.3. Codes and combinations of algorithm settings. These codes correspond to row numbers of the heatmap figures in the text. Algorithms are Double Expected SARSA (DXS), Double Q-Learning (DQ), Double SARSA (DS), Expected SARSA (ES), Q-Learning (Q), and SARSA (S). ϵ is the exploration rate. α is the learning rate.

Code	Algorithm	ϵ	α	Code	Algorithm	ϵ	α	Code	Algorithm	ϵ	α
1	DXS	0.01	0.2	31	DS	0.01	0.2	61	Q	0.01	0.2
2	DXS	0.01	0.4	32	DS	0.01	0.4	62	Q	0.01	0.4
3	DXS	0.01	0.6	33	DS	0.01	0.6	63	Q	0.01	0.6
4	DXS	0.01	0.8	34	DS	0.01	0.8	64	Q	0.01	0.8
5	DXS	0.01	1.0	35	DS	0.01	1.0	65	Q	0.01	1.0
6	DXS	0.10	0.2	36	DS	0.10	0.2	66	Q	0.10	0.2
7	DXS	0.10	0.4	37	DS	0.10	0.4	67	Q	0.10	0.4
8	DXS	0.10	0.6	38	DS	0.10	0.6	68	Q	0.10	0.6
9	DXS	0.10	0.8	39	DS	0.10	0.8	69	Q	0.10	0.8
10	DXS	0.10	1.0	40	DS	0.10	1.0	70	Q	0.10	1.0
11	DXS	0.20	0.2	41	DS	0.20	0.2	71	Q	0.20	0.2
12	DXS	0.20	0.4	42	DS	0.20	0.4	72	Q	0.20	0.4
13	DXS	0.20	0.6	43	DS	0.20	0.6	73	Q	0.20	0.6
14	DXS	0.20	0.8	44	DS	0.20	0.8	74	Q	0.20	0.8
15	DXS	0.20	1.0	45	DS	0.20	1.0	75	Q	0.20	1.0
16	DQ	0.01	0.2	46	ES	0.01	0.2	76	S	0.01	0.2
17	DQ	0.01	0.4	47	ES	0.01	0.4	77	S	0.01	0.4
18	DQ	0.01	0.6	48	ES	0.01	0.6	78	S	0.01	0.6
19	DQ	0.01	0.8	49	ES	0.01	0.8	79	S	0.01	0.8
20	DQ	0.01	1.0	50	ES	0.01	1.0	80	S	0.01	1.0
21	DQ	0.10	0.2	51	ES	0.10	0.2	81	S	0.10	0.2
22	DQ	0.10	0.4	52	ES	0.10	0.4	82	S	0.10	0.4
23	DQ	0.10	0.6	53	ES	0.10	0.6	83	S	0.10	0.6
24	DQ	0.10	0.8	54	ES	0.10	0.8	84	S	0.10	0.8
25	DQ	0.10	1.0	55	ES	0.10	1.0	85	S	0.10	1.0
26	DQ	0.20	0.2	56	ES	0.20	0.2	86	S	0.20	0.2
27	DQ	0.20	0.4	57	ES	0.20	0.4	87	S	0.20	0.4
28	DQ	0.20	0.6	58	ES	0.20	0.6	88	S	0.20	0.6
29	DQ	0.20	0.8	59	ES	0.20	0.8	89	S	0.20	0.8
30	DQ	0.20	1.0	60	ES	0.20	1.0	90	S	0.20	1.0

Appendix 1.C. Flowcharts of Selected Algorithms

Among the algorithms used in this study, Double SARSA and Double Expected SARSA performed better than the others. Their flowcharts are shown in Figures A1 and A2, respectively. For larger images of these flowcharts and the flowcharts of other algorithms used in the study, see the Data Availability section.

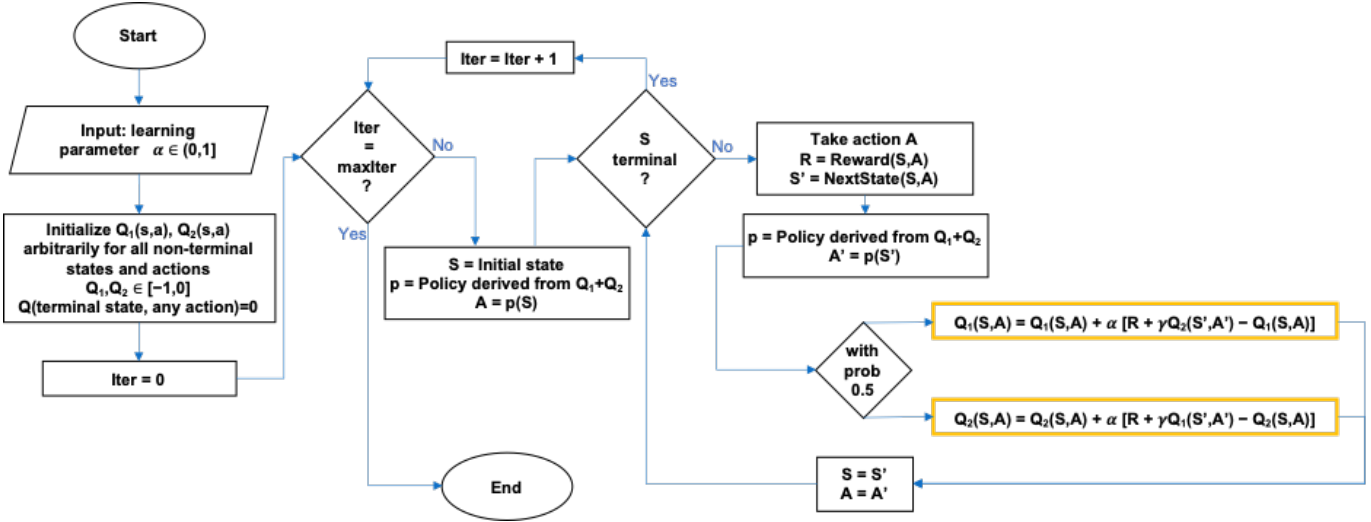


Figure S1.1. Double SARSA flowchart.

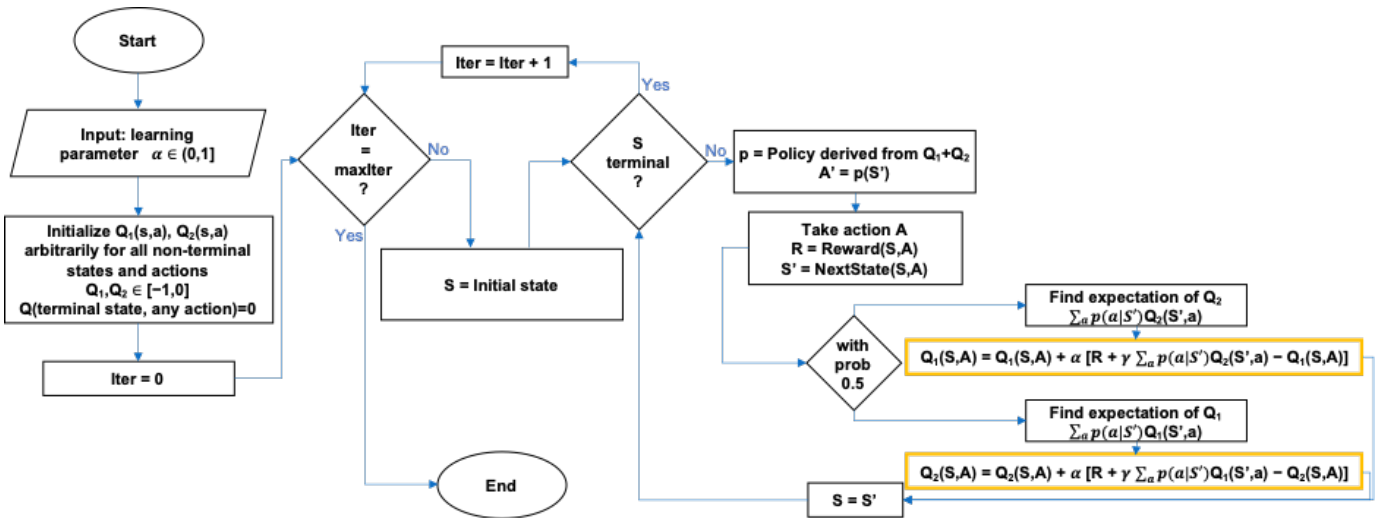


Figure S1.2. Double Expected SARSA flowchart.

Appendix 1.D. Post-publication remarks

This appendix has been added in the present thesis chapter after the publication of the paper. This appendix includes additional information on the contents of the chapter based on remarks and comments received between publication in journal and final submission of this thesis. Where necessary, references have been cited in this appendix and added in the chapter's bibliography.

Reinforcement Learning algorithms comprise the core of the intelligence of the governing agent introduced in this chapter. It is noteworthy to emphasize that the training of the algorithms is done iteratively in episodes. In the particular example of this study, and for computational reasons, 4000 training episodes were given to each model. Episodes are composed of time steps. Again, in the particular example of this study, 17 time steps were considered for each episode, and that number was chosen due to computational reasons. As explained in the text, building the reference dataset with 17 time steps would require 2^{16} runs for each configuration of model parameters.

It is also noteworthy that the algorithms may end an episode before the 17th time step if the target state occurs. The latter, in turn, has been defined as a state where the majority of user agents opt for the label and cooperate with the governing agent. A schematic representation of what happens in an episode in the governing agent's algorithm is given in the example flow charts of Appendix 1.C. In those flow charts, the right-most loop – which only runs if state S is not terminal – corresponds to the time steps within an episode.

Note that the RL algorithms of this study are all based on calculating the value of being in a certain state and then taking a certain action. For this reason, the *(State, Action)* pairs appear in various equations presented in this chapter. Values calculated for such pairs are stored in Q. In fact, retrieval the value stored for a specific *(State, Action)* pair is done by calling $Q(\text{State}, \text{Action})$. Related equations should be read with the consideration that the first component of such pairs refers to a state, and the second component refers to an action.

In RL problems typically involve an agent interacting with an entity and desiring to maximize the rewards that it receives through those interactions. In that regard, one strategy is to choose the policy that has lead to the highest rewards in previous efforts. This is called exploitation, because it involves using the knowledge gained thus far. The problem with exploitation is that it may keep the agent in a local maximum point without any knowledge of other possibilities in the state space. To overcome this problem, an alternative strategy is used: exploration. This alternative typically involves random behavior, so that the agent finds itself in a new part of the state space. In practice, learning involves a balance between exploration and exploitation, and this is included in RL algorithms with a parameter that indicates the rate of exploration vs. exploitation (Sutton & Barto, 2018).

As regards RL, one of the special features of this work is its definition of states. In the governing agent's RL algorithm, state is defined as a two-dimensional categorical variable. One of the dimensions of this variable is related to the response of user agents to the governing agent. As defined in the model, if more than half of the user agents cooperate with the governing agent, then the governing agent is in the target state. In non-target states, the proportion of cooperating user agents is less than half. It was decided to have an odd number of user agents, so as to avoid cases with exactly half of the user agents receiving the label. Among non-target states, it is ideal to distinguish states with higher proportions of cooperating user agents from those with lower cooperation. This means there should be several levels (or modalities) in that dimension of the state, in order to make such ideal distinction possible. In the simplest case, there will be 2 such levels (as no distinction is possible with 1 level). Then, the smallest number of user agents to make such distinction possible is 5. In this example, 3 or more user agents receiving the label constitutes the target state; and 2 and 1 user agents receiving the label constitute the two non-target modalities. In a similar fashion, in order to have 3 non-target modalities there should be at least 9 user agents; and in order to have 4 non-target modalities there should be at least 13 user agents. This was the reasoning behind the choice of number of user agents in the study.

The two selected algorithms, Double SARSA and Double Expected SARSA, had very similar performance in the tests. This was evident in Figure 1.5. In addition to this visualization, quantitative measures, too, confirm the similarity of their performance. For example, mean and standard deviation of the final time step's scores were -0.603 and 0.241 for Double SARSA, and -0.604 and 0.242 for Double Expected SARSA, respectively.

Finally, in order to help better understand the ranks heatmap, the following figure is added. This figure shows the zones associated with the algorithms in the heatmap.

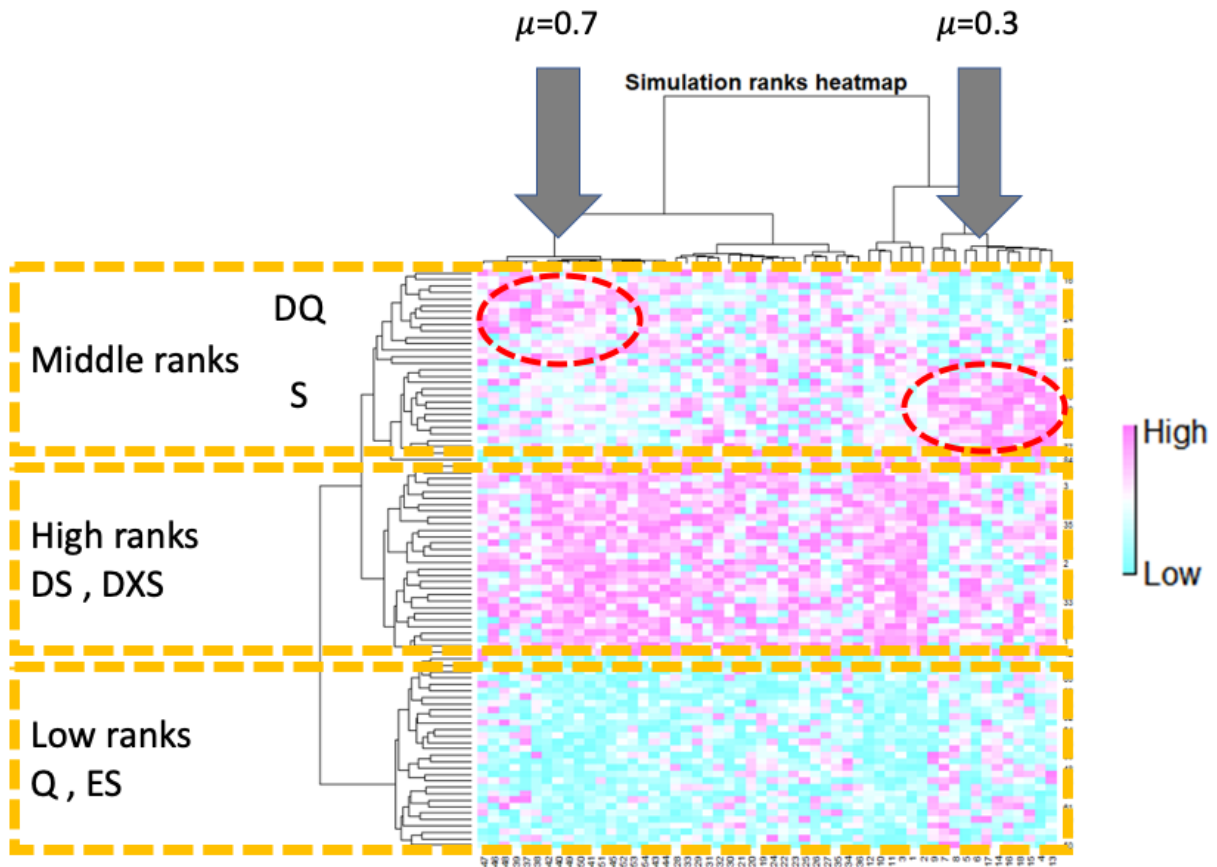


Figure S1.3. Simulation ranks heatmap with zones associated to the learning algorithms.

References

- Afshar Sedigh, A. H., Purvis, M. K., Savarimuthu, B. T. R., Frantz, C. K., & Purvis, M. A. (2021). Impact of Different Belief Facets on Agents' Decision---A Refined Cognitive Architecture to Model the Interaction Between Organisations' Institutional Characteristics and Agents' Behaviour. In A. Aler Tubella, S. Cranefield, C. Frantz, F. Meneguzzi, & W. Vasconcelos (Eds.), *Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XIII* (pp. 133–155). Cham: Springer International Publishing.
- Afshar Sedigh, A. H., Purvis, M. K., Savarimuthu, B. T. R., Purvis, M. A., & Frantz, C. K. (2021). Impact of Meta-roles on the Evolution of Organisational Institutions. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation* (pp. 66–80). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-030-66888-4_6
- Alpaydin, E. (2014). Reinforcement Learning. In *Introduction to Machine Learning* (2nd ed., pp. 517–545). Cambridge, MA, USA: MIT Press.
- Amato, F., Moscato, F., Moscato, V., Pascale, F., & Picariello, A. (2020). An agent-based approach for recommending cultural tours. *Pattern Recognition Letters*, *131*, 341–347. <https://doi.org/10.1016/j.patrec.2020.01.005>

- An, L. (2012). Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecological Modelling*, 229, 25–36. <https://doi.org/10.1016/j.ecolmodel.2011.07.010>
- Anderson, C., Hildreth, J. A. D., & Howland, L. (2015). Is the desire for status a fundamental human motive? A review of the empirical literature. *Psychological Bulletin*, 141(3), 574–601. <https://doi.org/10.1037/a0038781>
- Anderson, T., Leung, A., Dragicevic, S., & Perez, L. (2021). Modeling the geospatial dynamics of residential segregation in three Canadian cities: An agent-based approach. *Transactions in GIS*, 25(2), 948–967. <https://doi.org/10.1111/tgis.12712>
- Anderson, T., Leung, A., Perez, L., & Dragičević, S. (2021). Investigating the Effects of Panethnicity in Geospatial Models of Segregation. *Applied Spatial Analysis and Policy*, 14(2), 273–295. <https://doi.org/10.1007/s12061-020-09355-2>
- Andrighetto, G., Campenni, M., Cecconi, F., & Conte, R. (2010). The Complex Loop of Norm Emergence: A Simulation Model. In *Simulating Interacting Agents and Social Phenomena* (pp. 19–35). Tokyo: Springer Japan. https://doi.org/10.1007/978-4-431-99781-8_2
- Angourakis, A., Santos, J. I., Galán, J. M., & Balbo, A. L. (2015). Food for all: An agent-based model to explore the emergence and implications of cooperation for food storage. *Environmental Archaeology*, 20(4), 349–363. <https://doi.org/10.1179/1749631414Y.0000000041>
- Antinyan, A., Horváth, G., & Jia, M. (2019). Social status competition and the impact of income inequality in evolving social networks: An agent-based model. *Journal of Behavioral and Experimental Economics*, 79(March 2018), 53–69. <https://doi.org/10.1016/j.socec.2018.12.008>
- Arnoux Hebert, G., Perez, L., & Harati, S. (2018). An Agent-Based Model to Identify Migration Pathways of Refugees: The Case of Syria. In L. Perez, E.-K. Kim, & R. Sengupta (Eds.), *Agent-Based Models and Complexity Science in the Age of Geospatial Big Data, Advances in Geographic Information Science* (1st ed., pp. 45–58). Basel, Switzerland: Springer International Publishing. <https://doi.org/10.1007/978-3-319-65993-0>
- Axelrod, R. (1981). The Emergence of Cooperation among Egoists. *American Political Science Review*, 75(02), 306–318. <https://doi.org/10.2307/1961366>
- Axelrod, R. (1986). An evolutionary approach to norms. *The American Political Science Review*, 80(4), 1095–1111.
- Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Vol. 3). Princeton, NJ, USA: Princeton University Press.
- Barr, S. (2003). Strategies for sustainability: Citizens and responsible environmental behaviour. *Area*, 35(3), 227–240. <https://doi.org/10.1111/1475-4762.00172>
- Batty, M., & Torrens, P. M. (2005). Modelling and prediction in a complex world. *Futures*, 37(7), 745–766. <https://doi.org/10.1016/j.futures.2004.11.003>
- Blundell, A. G., & Gullison, R. E. (2003). Poor regulatory capacity limits the ability of science to influence the management of mahogany. *Forest Policy and Economics*, 5(4), 395–405. [https://doi.org/10.1016/S1389-9341\(03\)00038-8](https://doi.org/10.1016/S1389-9341(03)00038-8)
- Bohensky, E. (2014). Learning Dilemmas in a Social-Ecological System: An Agent-Based Modeling Exploration. *Journal of Artificial Societies and Social Simulation*, 17(1), 1–2. <https://doi.org/10.18564/jasss.2448>
- Boman, M. (1999). Norms in artificial decision making. *Artificial Intelligence and Law*, 7(1), 17–35. <https://doi.org/10.1023/A:1008311429414>

- Bone, C., & Dragičević, S. (2010). Simulation and validation of a reinforcement learning agent-based model for multi-stakeholder forest management. *Computers, Environment and Urban Systems*, 34(2), 162–174. <https://doi.org/10.1016/j.compenvurbsys.2009.10.001>
- Bratman, M. (1987). *Intention, plans, and practical reason*. Cambridge, Massachusetts: Harvard University Press.
- Braun, D., & Guston, D. H. (2003). Principal-agent theory and research policy: An introduction. *Science and Public Policy*, 30(5), 302–308. <https://doi.org/10.3152/147154303781780290>
- Bristow, M., Fang, L., & Hipel, K. W. (2014). Agent-Based Modeling of Competitive and Cooperative Behavior Under Conflict. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 44(7), 834–850. <https://doi.org/10.1109/TSMC.2013.2282314>
- Canese, L., Cardarilli, G. C., Di Nunzio, L., Fazzolari, R., Giardino, D., Re, M., & Spanò, S. (2021). Multi-Agent Reinforcement Learning: A Review of Challenges and Applications. *Applied Sciences*, 11(11), 4948. <https://doi.org/10.3390/app11114948>
- Castle, C., & Crooks, A. (2006). *Principles and concepts of agent-based modelling for developing geospatial simulations*. London, UK: Centre for Advanced Spatial Analysis (UCL).
- Chan, C. K., & Steiglitz, K. (2008). *An agent-based model of a minimal economy*. Princeton, NJ, USA: Princeton University.
- Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology*, 58(6), 1015–1026. <https://doi.org/10.1037/0022-3514.58.6.1015>
- Crawford, S. E. ., & Ostrom, E. (1995). A grammar of institutions. *American Political Science Review*, 89(3), 582–600.
- Crooks, A., Castle, C., & Batty, M. (2008). Key challenges in agent-based modelling for geospatial simulation. *Computers, Environment and Urban Systems*, 32(6), 417–430. <https://doi.org/10.1016/j.compenvurbsys.2008.09.004>
- Delre, S. A., Jager, W., Bijmolt, T. H. A., & Janssen, M. A. (2007). Targeting and timing promotional activities: An agent-based model for the takeoff of new products. *Journal of Business Research*, 60(8), 826–835. <https://doi.org/10.1016/j.jbusres.2007.02.002>
- Delre, Sebastiano A., Jager, W., Bijmolt, T. H. A., & Janssen, M. A. (2010). Will it spread or not? the effects of social influences and network topology on innovation diffusion. *Journal of Product Innovation Management*, 27(2), 267–282. <https://doi.org/10.1111/j.1540-5885.2010.00714.x>
- Feeny, D., Berkes, F., McCay, B. J., & Acheson, J. M. (1990). The Tragedy of the Commons: Twenty-Two Years Later. *Human Ecology*, 18(1), 1–19.
- Gaudreau, J., Perez, L., & Harati, S. (2018). Towards Modelling Future Trends of Quebec’s Boreal Birds’ Species Distribution under Climate Change. *ISPRS International Journal of Geo-Information*, 7(9), 335. <https://doi.org/10.3390/ijgi7090335>
- Gavin, M. (2018). An agent-based computational approach to “the Adam Smith problem.” *Historical Social Research*, 43(1), 308–336. <https://doi.org/10.12759/hsr.43.2018.1.308-336>
- Goldstein, J. (1999). Emergence as a Construct: History and Issues. *Emergence*, 1, 49–72. https://doi.org/10.1207/s15327000em0101_4
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., ... DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126. <https://doi.org/10.1016/j.ecolmodel.2006.04.023>

- Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, 221(23), 2760–2768. <https://doi.org/10.1016/j.ecolmodel.2010.08.019>
- Grimm, V., Railsback, S. F., Vincenot, C. E., Berger, U., Gallagher, C., DeAngelis, D. L., ... Ayllón, D. (2020). The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism. *Journal of Artificial Societies and Social Simulation*, 23(2), 1–7. <https://doi.org/10.18564/jasss.4259>
- Hales, D. (2002). Group reputation supports beneficent norms. *Journal of Artificial Societies and Social Simulation*, 5(4), 1–4.
- Hollander, C. D., & Wu, A. S. (2011). The Current State of Normative Agent-Based Systems. *Journal of Artificial Societies and Social Simulation*, 14(2). <https://doi.org/10.18564/jasss.1750>
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 237–285. <https://doi.org/10.1613/jair.301>
- Kiesling, E., Günther, M., Stummer, C., & Wakolbinger, L. M. (2012). Agent-based simulation of innovation diffusion: A review. *Central European Journal of Operations Research*, 20(2), 183–230. <https://doi.org/10.1007/s10100-011-0210-y>
- Lazaric, N., Le Guel, F., Belin, J., Oltra, V., Lavaud, S., & Douai, A. (2020). Determinants of sustainable consumption in France: the importance of social influence and environmental values. *Journal of Evolutionary Economics*, 30(5), 1337–1366. <https://doi.org/10.1007/s00191-019-00654-7>
- Li, X., & Liu, X. (2008). Embedding sustainable development strategies in agent-based models for use as a planning tool. *International Journal of Geographical Information Science*, 22(1), 21–45. <https://doi.org/10.1080/13658810701228686>
- Lindström, B., Jangard, S., Selbing, I., & Olsson, A. (2018). The role of a “common is moral” heuristic in the stability and change of moral norms. *Journal of Experimental Psychology: General*, 147(2), 228–242. <https://doi.org/10.1037/xge0000365>
- Liu, J., Dietz, T., Carpenter, S. R., Alberti, M., Folke, C., Moran, E., ... Taylor, W. W. (2007). Complexity of Coupled Human and Natural Systems. *Science*, 317(5844), 1513–1516. <https://doi.org/10.1126/science.1144004>
- Livet, P., Phan, D., & Sanders, L. (2014). Diversité et complémentarité des modèles multi-agents en sciences sociales. *Revue Française de Sociologie*, 55(4), 689. <https://doi.org/10.3917/rfs.554.0689>
- Maggi, E., & Vallino, E. (2021). Price-based and motivation-based policies for sustainable urban commuting: An agent-based model. *Research in Transportation Business & Management*, 39, 100588. <https://doi.org/10.1016/j.rtbm.2020.100588>
- Manson, S. M. (2001). Simplifying complexity: A review of complexity theory. *Geoforum*, 32(3), 405–414. [https://doi.org/10.1016/S0016-7185\(00\)00035-X](https://doi.org/10.1016/S0016-7185(00)00035-X)
- Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., ... Schwarz, N. (2013). Describing human decisions in agent-based models – ODD + D, an extension of the ODD protocol. *Environmental Modelling & Software*, 48, 37–48. <https://doi.org/10.1016/j.envsoft.2013.06.003>
- Nolan, J. M., Schultz, P. W., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2008). Normative Social Influence is Underdetected. *Personality and Social Psychology Bulletin*, 34(7), 913–923. <https://doi.org/10.1177/0146167208316691>
- North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge,

- UK: Cambridge University Press. <https://doi.org/10.1017/CBO9780511808678>
- North, M. J., Collier, N. T., Ozik, J., Tataru, E. R., Macal, C. M., Bragen, M., & Sydelko, P. (2013). Complex adaptive systems modeling with Repast Symphony. *Complex Adaptive Systems Modeling*, 1, 3. <https://doi.org/10.1186/2194-3206-1-3>
- Okdinawati, L., Simatupang, T. M., & Sunitiyoso, Y. (2017). Multi-agent Reinforcement Learning for Collaborative Transportation Management (CTM). In *Agent-Based Approaches in Economics and Social Complex Systems IX* (pp. 123–136). Singapore: Springer Singapore. https://doi.org/10.1007/978-981-10-3662-0_10
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge, UK: Cambridge University Press.
- Ostrom, E. (2009). A General Framework for Analyzing Sustainability of Social-Ecological Systems. *Science*, 325(July), 419–422. <https://doi.org/10.1126/science.1172133>
- Perez, L., & Dragicevic, S. (2009). An agent-based approach for modeling dynamics of contagious disease spread. *International Journal of Health Geographics*, 8(1), 50. <https://doi.org/10.1186/1476-072X-8-50>
- Perez, L., & Dragicevic, S. (2010). Modeling mountain pine beetle infestation with an agent-based approach at two spatial scales. *Environmental Modelling & Software*, 25(2), 223–236. <https://doi.org/10.1016/j.envsoft.2009.08.004>
- Perez, L., Dragicevic, S., & Gaudreau, J. (2019). A geospatial agent-based model of the spatial urban dynamics of immigrant population: A study of the island of Montreal, Canada. *PLOS ONE*, 14(7), e0219188. <https://doi.org/10.1371/journal.pone.0219188>
- R Core Team. (2019). R: A Language and Environment for Statistical Computing. Vienna, Austria, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org>
- Rand, W., & Rust, R. T. (2011). Agent-based modeling in marketing: Guidelines for rigor. *International Journal of Research in Marketing*, 28(3), 181–193. <https://doi.org/10.1016/j.ijresmar.2011.04.002>
- Rao, A., & Georgeff, M. (n.d.). BDI Agents: From Theory to Practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*. San Francisco, CA, USA, 12-14 June 1995.
- Rasch, S., Heckeley, T., Oomen, R., & Naumann, C. (2016). Cooperation and collapse in a communal livestock production SES model – A case from South Africa. *Environmental Modelling & Software*, 75, 402–413. <https://doi.org/10.1016/j.envsoft.2014.12.008>
- Ross, H. L. (1973). *Perspectives on social order*. New York: McGraw-Hill.
- Sanchez-Anguix, V., Tunalı, O., Aydoğan, R., & Julian, V. (2021). Can Social Agents Efficiently Perform in Automated Negotiation? *Applied Sciences*, 11(13), 6022. <https://doi.org/10.3390/app11136022>
- Savarimuthu, B. T. R., & Cranefield, S. (2011). Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, 7(1), 21–54. <https://doi.org/10.3233/MGS-2011-0167>
- Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–186. <https://doi.org/10.1080/0022250X.1971.9989794>
- Shafiei, E., Thorkelsson, H., Ásgeirsson, E. I., Davidsdottir, B., Raberto, M., & Stefansson, H. (2012). An agent-based modeling approach to predict the evolution of market share of electric vehicles: A case study from Iceland. *Technological Forecasting and Social Change*, 79(9), 1638–1653. <https://doi.org/10.1016/j.techfore.2012.05.011>

- Shoham, Y., & Tennenholtz, M. (n.d.). Emergent conventions in multi-agent systems: Initial experimental results and observations. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR)* (pp. 225–231). San Mateo, CA, USA, 25-29 October 1992: Morgan Kaufmann.
- Smith, A. (1982). *An Inquiry Into the Nature and Causes of the Wealth of Nations*. (R. H. Campbell & A. S. Skinner, Eds.). Indianapolis, IN, USA (first published 1776): Liberty Fund.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). Cambridge, Massachusetts: MIT Press.
- Tascioglu, M., Eastman, J. K., & Iyer, R. (2017). The impact of the motivation for status on consumers' perceptions of retailer sustainability: the moderating impact of collectivism and materialism. *Journal of Consumer Marketing*, 34(4), 292–305. <https://doi.org/10.1108/JCM-03-2015-1351>
- Therborn, G. (2002). Back to Norms! on the Scope and Dynamics of Norms and Normative Action. *Current Sociology*, 50(6), 863–880. <https://doi.org/10.1177/0011392102050006006>
- Wagner, W. (2004). Commons Ignorance: The Failure of Environmental Law to Produce Needed Information on Health and the Environment. *Duke Law Journal*, 53(6), 1619–1746.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. New York, NY, USA: Springer-Verlag. Retrieved from <https://ggplot2.tidyverse.org>
- Wickham, H., & Seidel, D. (2020). R Package “scales”: Scale Functions for Visualization. Retrieved from <https://cran.r-project.org/package=scales>
- Wittemyer, G., Daballen, D., & Douglas-Hamilton, I. (2011). Rising ivory prices threaten elephants. *Nature*, 476(7360), 282–283. <https://doi.org/10.1038/476282c>
- Wolfe, B. E. (2020). R Package “signs”: Insert Proper Minus Signs. Retrieved from <https://cran.r-project.org/package=signs>
- Yazan, D. M., & Fraccascia, L. (2020). Sustainable operations of industrial symbiosis: an enterprise input-output model integrated by agent-based simulation. *International Journal of Production Research*, 58(2), 392–414. <https://doi.org/10.1080/00207543.2019.1590660>

Chapter 2

Presentation of the article

The subject of the second chapter of the thesis is a model that simulates an ecological change – forest insect infestations. This model was built with several simplifying assumptions. Specifically, it may be argued that the model did not include many potentially important explanatory factors, such as age and species composition of forest stands as well as wind, precipitation and other climatic variables. Moreover, the state of infestation of grid cells was only expressed in binary form, that is, instead of reading and predicting percentages of infestation of cells, the model simply considered each cell as either infested or not infested. Although these simplifications may be regarded as limits for the model, the simulations arguably demonstrated spatiotemporal complexities of the studied dispersion phenomenon.

This chapter has been published in the peer-reviewed journal “Forests” in 2020. My coauthors in this publication were my supervisors, Dr. Liliana Perez and Dr. Roberto Molowny-Horas. The chapter as it appears in this thesis involves modifications in the layout and style of the published paper, and slight modifications in figure and table numbers. Other than those, this chapter includes no changes in the content of the published paper.

Citation:

Harati, S., Perez, L., Molowny-Horas, R. Integrating Neighborhood Effect and Supervised Machine Learning Techniques to Model and Simulate Forest Insect Outbreaks in British Columbia, Canada. *Forests*. **2020**, *11*, 1215. <https://doi.org/10.3390/f11111215>

Integrating Neighborhood Effect and Supervised Machine Learning Techniques to Model and Simulate Forest Insect Outbreaks in British Columbia, Canada

Saeed Harati ^{1,*}, Liliana Perez ¹ and Roberto Molowny-Horas ²

¹ Laboratory of Environmental Geosimulation (LEDGE), Department of Geography, Université de Montréal, 1375 Ave. Thérèse-Lavoie-Roux, Montreal, H2V 0B3 QC, Canada; l.perez@umontreal.ca

² CREAM, Centre de Recerca Ecològica i Aplicacions Forestals, E08193 Bellaterra (Cerdanyola de Vallès), Catalonia, Spain; roberto@creaf.uab.es

* Correspondence: saeed.harati.asl@umontreal.ca

Harati, S., Perez, L., Molowny-Horas, R. Integrating Neighborhood Effect and Supervised Machine Learning Techniques to Model and Simulate Forest Insect Outbreaks in British Columbia, Canada. *Forests*. **2020**, *11*, 1215. <https://doi.org/10.3390/f11111215>

Abstract

Background and Objectives: Modelling and simulation of forest land cover change due to epidemic insect outbreaks are powerful tools that can be used in planning and preparing strategies for forest management. In this study, we propose an integrative approach to model land cover changes at a provincial level, using as a study case the simulation of the spatiotemporal dynamics of mountain pine beetle (MPB) infestation over the lodgepole pine forest of British Columbia (BC), Canada. This paper aims to simulate land cover change by applying supervised machine learning techniques to maps of MPB-driven deforestation. *Materials and Methods:* We used a 16-year series (1999–2014) of spatial information on annual mortality of pine trees due to MPB attacks, provided by the BC Ministry of Forests. We used elevation, aspect, slope, ruggedness, and weighted neighborhood of infestation as predictors. We implemented a) generalized linear regression (GLM), and b) random forest (RF) algorithms to simulate forestland cover changes due to MPB between 2005 and 2014. To optimize the ability of our models to predict MPB infestation in 2020, a cross-validation procedure was implemented.

Results: Simulating infestations from 2008 to 2014, RF algorithms produced less error than GLM. Our simulations for the year 2020 confirmed the predictions from the BC Ministry of Forest by forecasting a slower rate of spread in future MPB infestations in the province.

Conclusions: Integrating neighborhood effects as variables in model calibration allows spatiotemporal complexities to be simulated.

Keywords: land cover change; complex systems; model calibration; random forest; insect outbreaks; regression; machine learning

2.1. Introduction

Disturbances are a critical component of forests dynamics, which shape and substantially affect these key ecosystems (Bourbonnais, Nelson, & Wulder, 2014; McCullough, Werner, & Neumann, 1998). Particularly of interest, forest insect epidemics can exert severe impacts on ecosystem dynamics due to mortality or growth reduction of millions of trees over widespread areas (Axelson, Alfaro, & Hawkes, 2010; MacLean, 2016; Pelz & Smith, 2012; Robbins, 2008). Both indigenous and invasive species can disturb natural and managed forests habitats (Herms & McCullough, 2014; Pelz & Smith, 2012; Sturtevant et al., 2013). Additional to ecological impacts, insect epidemics may have devastating effects associated with economic (e.g., losses within the forestry industry) (Chang, Lantz, Hennigar, & MacLean, 2012; Corbett, Withey, Lantz, & Ochuodho, 2016; Patriquin, Wellstead, & White, 2007) and social (e.g., unemployment due the sawmill closures) development or stability (Flint, McFarlane, & Müller, 2009; Petersen & Stuart, 2014).

In the province of British Columbia (BC), Canada, an unprecedented insect outbreak of mountain pine beetle (MPB; *Dendroctonus ponderosae* Hopkins) began in the early 1990s and reached a peak between 2005 and 2006, which facilitated a massive migration of beetles into the province of Alberta (AB) (Patriquin et al., 2007; Petersen & Stuart, 2014; Strohm, Reid, & Tyson, 2016). The cumulative forest area in BC alone that has been attacked by the MPB, since the ongoing outbreak that began in 1990, is estimated to be over 25 million hectares (Bone & Nelson, 2019; James & Huber, 2019). An indigenous insect to western North American forests, the MPB is a bark beetle that feeds mainly on lodgepole pine (*Pinus contorta* var. *latifolia* Engelm.), but also feeds on sugar pine (*Pinus lambertiana* Douglas), western white pine (*Pinus*

monticola Douglas ex D. Don), and ponderosa pine (*Pinus ponderosa* var. *scopulorum* Engelm.) (Bentz, Boone, & Raffa, 2015; L. Safranyik & Carroll, 2007). Until the end of the last century, the historical range of MPB was limited to the west of the continent (Logan & Powell, 2001; L. Safranyik et al., 2010). However, today the MPB is present outside of its historical range, extending into northern BC and eastward into the boreal forest of north-central AB, where approximately 1.43 million trees have been infested (Hodge, Cooke, & McIntosh, 2017). One of the greatest future threats from the current expansion of the range of MPB is that the beetle is no longer limited to attacking lodgepole pine, but is also reproducing in jack pine (*Pinus banksiana* Lambert), one of the dominant pine species of the boreal forest (L. Safranyik et al., 2010).

In the light of the severity of the current insect disturbance faced by BC and AB, it has become a pressing matter to continue monitoring, assessing, modelling, and simulating future changes of forest cover due to MPB outbreaks to assist environmental policy development and forestry resource management (Bentz & Jönsson, 2015; Cooke & Carroll, 2017; Ferretti, 1997; Hall, Castilla, White, Cooke, & Skakun, 2016). To assess the impact of the MPB epidemic on forest ecological systems, different methodologies have been proposed to date. Remote sensing (Coops, Wulder, & Waring, 2012; Liang, Hawbaker, Chen, Zhu, & Gong, 2014; M. A. Wulder, Ortlepp, White, Nelson, & Coops, 2010), equation-based (Lewis, Nelson, & Xu, 2010; Powell, Logan, & Bentz, 1996; L. Safranyik, Barclay, Thomson, & Riel, 1999), Geographical Information System (GIS)-based (C. Bone, Dragicevic, & Roberts, 2005; Christopher Bone, Wulder, White, Robertson, & Nelson, 2013; Liang, Li, Huang, Qin, & Huang, 2017; Macias Fauria & Johnson, 2009), and complex systems theory approaches (Christopher Bone & Altaweel, 2014; Liang et al., 2017; Perez & Dragicevic, 2012; Pérez & Dragičević, 2011) are some of the methodologies that have been most frequently applied to detect, model, and predict the spatial dispersal of the MPB population and attack patterns. Although the aforementioned simulation efforts have been successful at modelling the spread of MPB infestation, they have done so at a very detailed scale that ranges from tree to stand level. By comparison, the studies that claim to have modeled the infestations at a landscape scale have only gone as far as a county level for the United States, for example. The MPB spread appears to depend basically on topographic conditions and the state of neighboring areas (Perez & Dragicevic, 2012; L. Safranyik & Carroll, 2007; L. Safranyik et al., 2010; L. Safranyik et al., 1999). These drivers may in turn be mediated by local or regional climatic conditions (e.g., milder winters), although these

conditions may affect winter survival of the beetle, and not spatial spread per se. With the aim of carrying out spatio-temporal modelling and simulation of the MPB infestation at a province scale in BC, Canada, we set out to apply supervised machine learning techniques to maps of MPB-caused deforestation. This research study proposes an integrative approach to model land cover changes at a provincial level, using as a study case the simulation of the spatiotemporal dynamics of MPB infestation of the lodgepole pine forest of BC, Canada. The main objective of this work is three-fold, namely:

1. Compare the performance of two methodologies, namely binomial regression and random forests, to model the MPB spread between 1999 and 2014.
2. Evaluate the usefulness of a set of predictor variables, describing the influence of local topography and the state (i.e., infested/non-infested) of neighboring localities, to determine the extent and speed of the MPB infestation.
3. Simulate possible land cover changes in 2020, due to MPB infestation.

2.2. Materials and Methods

2.2.1. Study Area

The study was conducted in the Canadian province of British Columbia (BC), and covers an area of 944,735 km², extending from 59°59'27"N 138°54'19"W to 48°59'53"N 114°2'37"W (Figure 2.1). BC is known for its highly diverse mountainous landscape subject to a diversity of disturbance regimes (Axelson, Alfaro, & Hawkes, 2009; Haughian, Burton, Taylor, & Curry, 2012; Klenner, Walton, Arsenault, & Kremsater, 2008). The climatic conditions in the province are largely controlled by the Pacific Ocean to the west, continental air masses in the interior plateaus, and the Rocky Mountains to the east (Lemmen, Warren, Lacroix, & Bush, 2008).



Figure 2.1. The Province of British Columbia in Canada is mainly dominated by a Lodgepole pine (*Pinus contorta*) forest, which by 2014 had been decimated by almost 50% due to mountain pine beetle (MPB) infestation.

Seventy percent of the total area is covered by forest, whereas only two percent of the total area is used by humans to live, cultivate, etc. Forest in central BC, where lodgepole pines are the main tree species, have been experiencing an epidemic infestation of MPB, due to factors including fire suppression and changing climate (Natural Resources Canada, 2016).

2.2.2. Pine Mortality Dataset

The original source of data for the project is a collection of 16 maps indicating cumulative lodgepole pine mortality caused by MPB attacks on the forests of the province as observed in the period between 1999 and 2014 (L. Safranyik et al., 2010). Observations were acquired by the BC Ministry of Forest from aerial photographs and LANDSAT satellite images, wherein infested areas are identifiable based on their spectral response and by calculating a Normalized Difference Moisture Index (NDMI), contrasting the near-infrared (NIR) band 4, which is sensitive to the reflectance of leaf chlorophyll content, to the mid-infrared (MIR) band 5, which is sensitive to the absorbance of leaf moisture. These maps are in a raster format with an

Albers equal area projection and a cell size of 400 m. The cell values equal 10 times the percentage of infestation in each cell, hence ranging from 0 to 1000; the Ministry of Forests, Lands and Natural Resource Operations of the Canadian province of BC have made this dataset publicly accessible (BC Ministry of Forests, 2000). Existing literature and reports emphasize the importance of infestation levels above which the risk of MPB attack should be considered seriously by forest managers for further investigation (BC Ministry of Forests, 2000; Carroll et al., 2006; Shore & Safranyik, 1992).

For the purposes of this study we applied a threshold to the cumulative infestation maps in order to transform their continuous percentage scale into a binary scale (i.e., infested = 1/not-infested = 0). This is a simplified assumption that enabled us to apply well-known statistical methods to our datasets. The procedure to calculate that threshold value is described in detail in Appendix 2.A.

2.2.3. Predictor Variables

For modelling purposes, we assumed that the infestation pattern in the near future depended directly on the status of the infestation during past years. For the sake of notation, if we denote by t_2 the future year for which a prediction is sought, then t_1 represents the starting date from which the simulated map of MPB infestation is projected and t_{1p} designates a previous year, for which further explanatory variables, used by the model, must be determined. Throughout the study, $t_{1p} < t_1 < t_2$ and $t_1 - t_{1p} = 1$ year, although $t_2 - t_1 = 3$ years.

We also assumed that the probability of beetle infestation in our thresholded maps depended entirely, for a given pixel, both on the local topography and on the state of the infestation in adjacent pixels. Arguably, local topography may either enhance or stall MPB outbreaks by, e.g., boosting or blocking beetle flights, respectively. It may also determine local climatic and environmental conditions in forests, making them a more or less suitable habitat for beetles to settle and attack. In turn, the infestation status of nearby areas should arguably have a direct influence on the number of beetles that affect a given location. These are distance-dependent variables that must be determined for every individual simulation. That dependence, however, is not known a priori and must be approximated. Bearing these ideas in mind, we set out to select a set of variables that could serve as valid drivers for the MPB infestation.

Our choices for explanatory variables represented the drivers that we fed the MPB numeric infestation model. These variables are listed in Table 2.1.

Table 2.1. List of predictor variables used for modeling. The Acronym column identifies the corresponding variable in Equations (2.2a) and (2.2b). The Time column indicates whether the variable is calculated at t_{1p} or t_1 (see text for an explanation).

Predictor description	Acronym	Units	Time
Elevation	e	m	-
Aspect	a	Arbitrary	-
Slope	s	Radians	-
Ruggedness	r	Arbitrary	-
Identity	$Z_{iden,t_{1p}}$	Arbitrary	t_{1p}
Linear weight	$Z_{lin,t_{1p}}$	Arbitrary	t_{1p}
Inverse-distance weight	$Z_{inv,t_{1p}}$	Arbitrary	t_{1p}
Square-inverse-distance weight	$Z_{squ,t_{1p}}$	Arbitrary	t_{1p}
Identity	Z_{iden,t_1}	Arbitrary	t_1
Linear weight	Z_{lin,t_1}	Arbitrary	t_1
Inverse-distance weight	Z_{inv,t_1}	Arbitrary	t_1
Square-inverse-distance weight	Z_{squ,t_1}	Arbitrary	t_1

Topography-based explanatory variables were calculated only once at the beginning of the calibration because we assumed that local topographic conditions did not change during the time intervals spanned by the simulations:

1. Elevation: MPB infestation has been observed to take place mostly at low or medium heights (Michael A. Wulder et al., 2006). Elevation is defined as height above sea level per pixel. We used a Digital Elevation Map provided by GeoBC. The original pixel size of 500 was changed to 400 to match the resolution of the MPB infestation map.
2. Slope: Steeper areas may affect, for example, distances between tree canopies on a hillside (Cooke & Carroll, 2017; McIntire, 2004), which, in turn, may make it easier or harder for beetles to fly from one tree to another. Slope was calculated from the elevation map with the “terrain” function of the “raster” R package.
3. Aspect: The spread of the MPB infestation may benefit from milder temperatures (Milne & Lewis, 2011; Perez & Dragicevic, 2012; L. Safranyik et al., 2010) on south-oriented slopes. For that reason, aspect was calculated from the elevation map as the compass direction of the pixel slope face. We employed the “terrain” function of the

“raster” R package. Next, it was sine-transformed to avoid the discontinuity at point $0-2\pi$ radians ($0^\circ-360^\circ$). Sine and cosine functions were used to avoid ambiguity at 0 radians.

4. Ruggedness: Adult beetle flight may be faster and/or longer over open ground (Reid, 1962). To account for this effect, we implemented the Terrain Ruggedness Index (TRI) (Riley, DeGloria, & Elliot, 1999). The TRI index was calculated from the elevation map with the “terrain” function of the “raster” R package using an 8-pixel window.

In contrast, adjacency-type predictors had to be computed at every temporal step of the simulation. As a measure of the dependence of infestation rate on the state of the neighboring pixels, we computed a weighted sum of the surrounding pixels (containing 0s or 1s) at each location. The basic equation to account for the adjacency effect can be written as:

$$z_j = \sum_{\forall i, d_i \leq d_{max}} p_{ij} \cdot w_i$$

where z_j stands for the generic adjacency effect at location j . Index i runs such that the distance from location j at which pixel value p_{ij} is summed, i.e., d_i , is smaller or equal than the maximum size d_{max} of the weighting window, w_i . To account for the unknown true dependency, we used four different weighted sum expressions:

1. No-weighting (z_{iden}):

$$w_i = 1$$

2. Linear weighting (z_{lin}): weights decrease linearly until $d_i = d_{max}$

$$w_i = d_{max} - d_i$$

3. Inverse-distance weighting (z_{inv}): weights decrease as a function of the inverse of distance until $d_i = d_{max}$

$$w_i = \frac{1}{d_i}$$

4. Squared-inverse-distance weighting (z_{squ}): weights decrease as a function of the inverse of the squared distance until $d_i = d_{max}$

$$w_i = \frac{1}{d_i^2}$$

In parentheses, we have included the corresponding variable name in Equation (2.2) and Table 2.1. Appendix 2.C demonstrates the aforementioned four weighting functions in a neighborhood of radius 5.

All of these predictor variables were determined at every infested and non-infested pixel in the t_{1p} and t_1 maps. Regarding adjacency-type predictors, we included the weighted total number of surrounding infested pixels both at t_{1p} and at t_1 as adjacency-type predictors. Because we ignored the exact dependence of the infestation rate on these predictors, we chose several weighting procedures to account for this unknown dependency and included all of them in the calibration procedure as independent variables.

We carried out a preliminary exploration of the relationship between infestation probability, represented by the thresholded infestation map specified above, and the set of predictor variables described in the previous paragraphs. Appendix 2.A shows the log-transformed average infestation as a function of the binned predictors. In general, infestation rate appears to change linearly with predictors. The mean response showed a parabolic response vs. elevation, indicating that the expected infestation rate is proportional to a quadratic function of the elevation. In addition, it displayed an approximately linear response vs. slope, aspect, ruggedness, and the log-transformed adjacency measures.

2.2.4. Approaches to Model and Simulate Land Cover Changes

To model and simulate the changes in the lodgepole pine forest cover within the province of BC during the sixteen-year period of the data set on recorded MPB epidemics, we implemented two algorithms: (1) generalized linear regression (GLM), and (2) random forest (RF).

2.2.4.1. Generalized Linear Regression (GLM)

Generalized linear regression is a maximum-likelihood regression methodology that computes the parameters of a linear model that maximize an appropriate log-likelihood. It can be applied to cases where the error distribution of the response variable is not Gaussian. The linear

model and the dependent variable are related directly via a continuous link function, which relates the expected value of the response variable with a linear combination of the predictors. In turn, the error distribution of the dependent variable determines our choice for the distributional model from which a log-likelihood can then be derived. There are several common options for log-likelihood and link functions, and the modeler must make his/her own choice by considering the observed relationships between response and predictor variables. The link function imposes a final expression for the log-likelihood that is often non-linear in the parameters. As a consequence, maximization of the log-likelihood is carried out iteratively with an appropriate optimization scheme (see e.g., McCullagh and Nelder, 1989).

To account for the binary dependent variable representing an MPB affected/non-affected pixel, we chose a Bernoulli log-likelihood, which corresponds to a binomial log-likelihood with number of trials equal to one. In turn, we determined the most appropriate dependence between expected binary response and the explanatory variables by carrying out an exploratory exercise (see Appendix 2.B). Based on those observed relationships, we proposed the following linear logit link function:

$$\text{logit}(\mu_j) = \beta_0 + \beta_1 e_j + \beta_3 s_j + \beta_4 a_j + \beta_5 r_j + \beta_6 z_{iden,t_1} + \beta_7 z_{lin,t_1} + \beta_8 z_{inv,t_1} + \beta_9 z_{squ,t_1} + \beta_{10} z_{iden,t_{1p}} + \beta_{11} z_{lin,t_{1p}} + \beta_{12} z_{inv,t_{1p}} + \beta_{13} z_{squ,t_{1p}}$$

$$\text{logit}(\mu_j) = \beta_0 + \beta_1 e_j + \beta_2 e_j^2 + \beta_3 s_j + \beta_4 a_j + \beta_5 r_j + \beta_6 z_{iden,t_1} + \beta_7 z_{lin,t_1} + \beta_8 z_{inv,t_1} + \beta_9 z_{squ,t_1} + \beta_{10} z_{iden,t_{1p}} + \beta_{11} z_{lin,t_{1p}} + \beta_{12} z_{inv,t_{1p}} + \beta_{13} z_{squ,t_{1p}}$$

For the sake of clarity, sub index j ($i \in j \dots n$, where n is the number of observations) was omitted in the neighborhood covariates z . In these equations, μ_j indicates the logit-transformed expected value of the binary response variable at location j , the sub-indexed β s are the unknown parameters to be calculated, and the predictor variables are shown in Latin letters, following the notation illustrated in Table 2.1. In our calculations we used the “glm” function of the built-in stats package of the R software (R Core Team, 2019).

The generalized linear regression scheme described above produced a continuous function, μ , bounded between 0 and 1 such that it represented the probability of infestation. Therefore, at each location j , that probability is given by:

$$\mu_j = \frac{1}{1 + e^{-(\beta_0 + \beta_1 e_j + \dots)}}$$

where the ellipsis in the exponent indicates all linear and their interactions, depending on whether we use Equation (2.2a) or (2.2b).

We then used μ to elaborate a predictive binary map pinpointing the location of newly infested pixels. This last step entailed the selection of a valid cut-off value to map the continuous μ variable onto a binary scale. We computed a suitable cut-off for μ by maximizing the fuzzy kappa index between observed and predicted infestation maps.

2.2.4.2. *Random Forests (RF)*

Random forests (RF) belong to the group of ensemble learning methods. Used for both classification and regression, RF are sets of decision trees each including a random subset of the data, that work based on the principle of highest voted decision, or the average of decided scores, depending on the type of the model. Throughout the calculations, we used the RF algorithms implemented in the “ranger” package of the R software. As predictor variables we used the same variables shown in Equation 2.2a. Preliminary results suggested that it was better to regress than to classify with RF, so we used the “ranger” function of that package in regression mode.

2.2.5. **Model Calibration and Validation**

To optimize the ability of our models to predict infestation in non-infested pixels we set up a cross-validation strategy by dividing the datasets into training and test subsets. The training dataset was used to calculate the unknown parameters of the model, whether binomial or RF, as shown above. The test dataset, in turn, was used to gauge the predictive ability of the models and to adjust a relevant parameter. We selected a training set containing 75% of the original dataset, whereas the test set included the remaining 25%.

Figure 2.2 shows the method of analysis in a flow chart. The flow of the algorithm is downwards and generally to the right. Raw inputs include geographical variables of elevation,

slope, and aspect, in addition to three infestation maps. The time interval between the first two maps at t_{1p} and t_1 is one year, and the time step between the last two maps at t_1 and t_2 is 3 years. The goal of the algorithm is to predict the binary spread of infestations in the next time step. After reading inputs, the algorithm is divided into two main parts. In the first part, the inputs are analyzed, and a model is developed and parameterized to simulate the change from the first two images (t_{1p} and t_1 maps) to the third (t_2). In the second part, the parameterized model is used to predict the change from the last two given images (t_{2p} and t_2 maps) to the next time step in the future (t_3).

In this model, parameters were determined and adjusted in three stages. These stages are shown with gray flowchart shapes in Figure 2.2. First, the initial threshold for conversion of numeric data into binary (infested/not infested) was identified. The input images were converted to binary assuming a threshold of 0.1586553 (see Appendix 2.A), i.e., infestation proportions of above 0.158 are taken as presence and lower values as non-presence. The second stage was the adjustment of the regression model parameters. To do so, the compiled dataset was divided into training and test datasets. Using the training dataset, regression parameters were estimated. This preliminary model was then applied to the test dataset in order to assess its ability to predict the changes in input data. Such assessment involved the third stage of model parameterization—the selection of the final threshold for creation of the binary output. This was done by applying and testing the goodness-of-fit of 1000 images created by different cut-off values. The test involved giving test data to the model and comparing its output with the most recent input image. The cut-off thresholds giving the highest kappa and the highest Youden's J statistic was selected for application in the prediction phase.

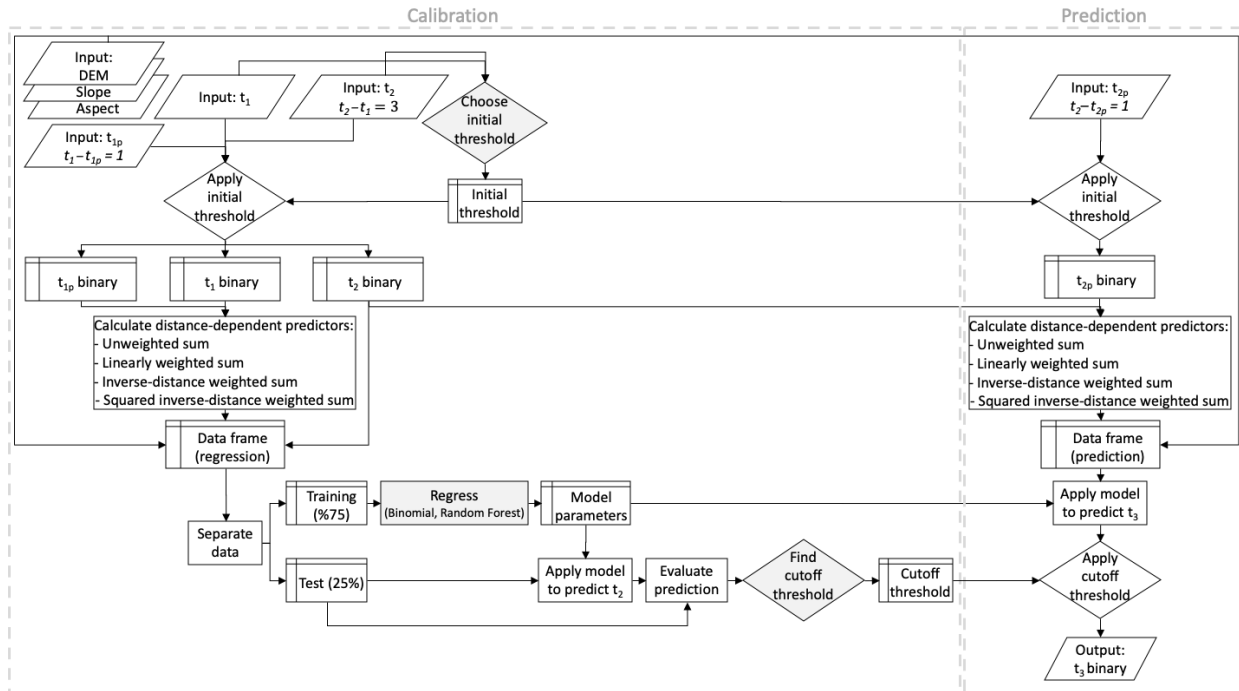


Figure 2.2. Schematic representation of the integrative approach to parametrize and calibrate our model to simulate forest insect disturbances.

Model output was compared with the reference data for validation. It should be noted that the similarity of simulation and reference maps of the final year does not provide sufficient information for validating the model (Pontius, Huffaker, & Denman, 2004; Rutherford, Guisan, & Zimmermann, 2007; White & Engelen, 1993). In fact, if the overall change in the landscape is small, even an erroneous prediction can still be highly similar to the reference. Therefore, model validation should account for the simulated and observed change. This requires considering three maps: a reference map of the beginning time, and reference and simulation maps at the ending time. In this study, the process of change is infestation, which is irreversible. As such, change may only occur in areas that were initially not infested. We analyzed these areas when validating the model.

2.2.6. Software

Calibration and simulation algorithms were computed using R 3.6.0 (R Core Team, 2019) and cartography was produced using ArcGIS Pro 2.4 (ESRI, 2019).

2.3. Results

Validation testing of the model involved simulation of changes in the study area from 2008 to 2014 and comparison with reference observations. Prior to testing, the model was parameterized using data of changes from 2005 to 2008. Then, in two rounds, it simulated the change from 2008 to 2011 and from the predicted 2011 to 2014. These predictions were made with 3-year time steps. Figure 2.3 shows model results for three algorithms—binomial regression (Equation (2.2a); hereafter GLM1), binomial regression with parabolic elevation (Equation (2.2b); hereafter GLM2), and random forest (RF) with maximum kappa final cut-off threshold—in addition to the observed change.

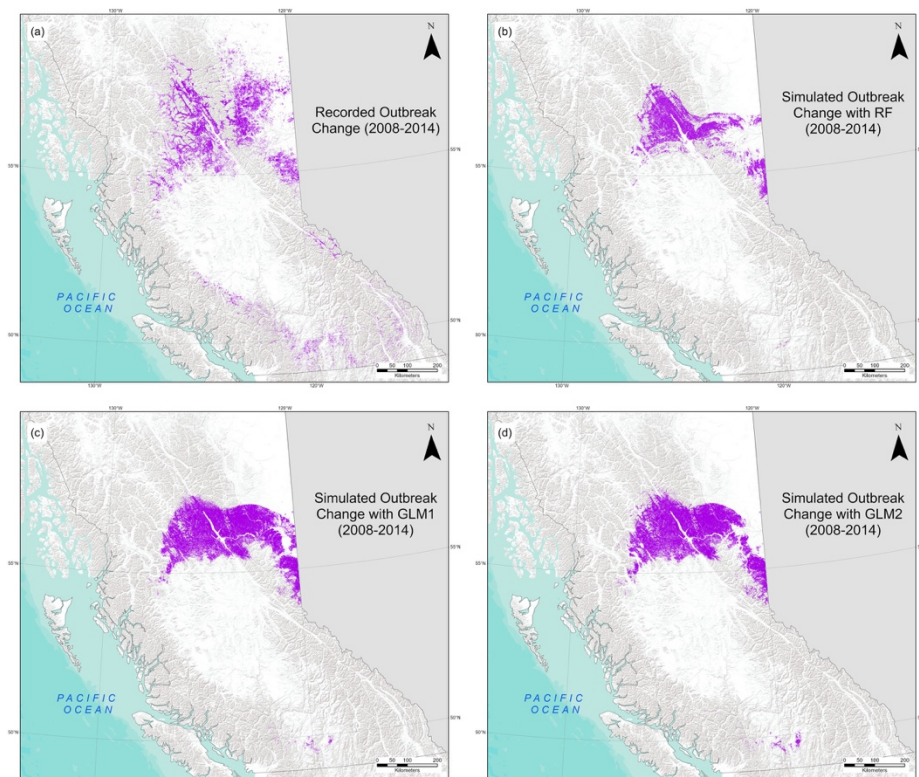


Figure 2.3. Comparison of model simulations with reference observations of changes in the study area from 2008 to 2014. (a) Mountain pine beetle outbreak change in British Columbia (BC) recorded by satellite imagery and aerial surveys; mountain pine beetle outbreak change simulated with algorithms: (b) random forests (RF), (c) binomial regression (GLM1), (d) binomial regression with parabolic elevation (GLM2).

Outputs of validation analysis for the three algorithms with maximum kappa final cut-off threshold are presented in Figure 2.4. Each map of this figure corresponds to one simulation algorithm and demonstrates a comparison between three images: observed infestations in 2008, simulated infestations in 2014, and observed infestations in 2014. In this figure, correctly simulated changes are identified as “hits”; observed changes that are missing from simulations are identified as “misses”; simulated changes that are not observed are identified as “false alarms”; and correct simulations of no change are identified as “correct rejections”. Because infestation is a one-way process that is not reversible, it is evident that zones of “prior infestation”, that is, zones that were already infested at the beginning of the study period, are not susceptible to future change and they should be excluded from the assessment of model performance.

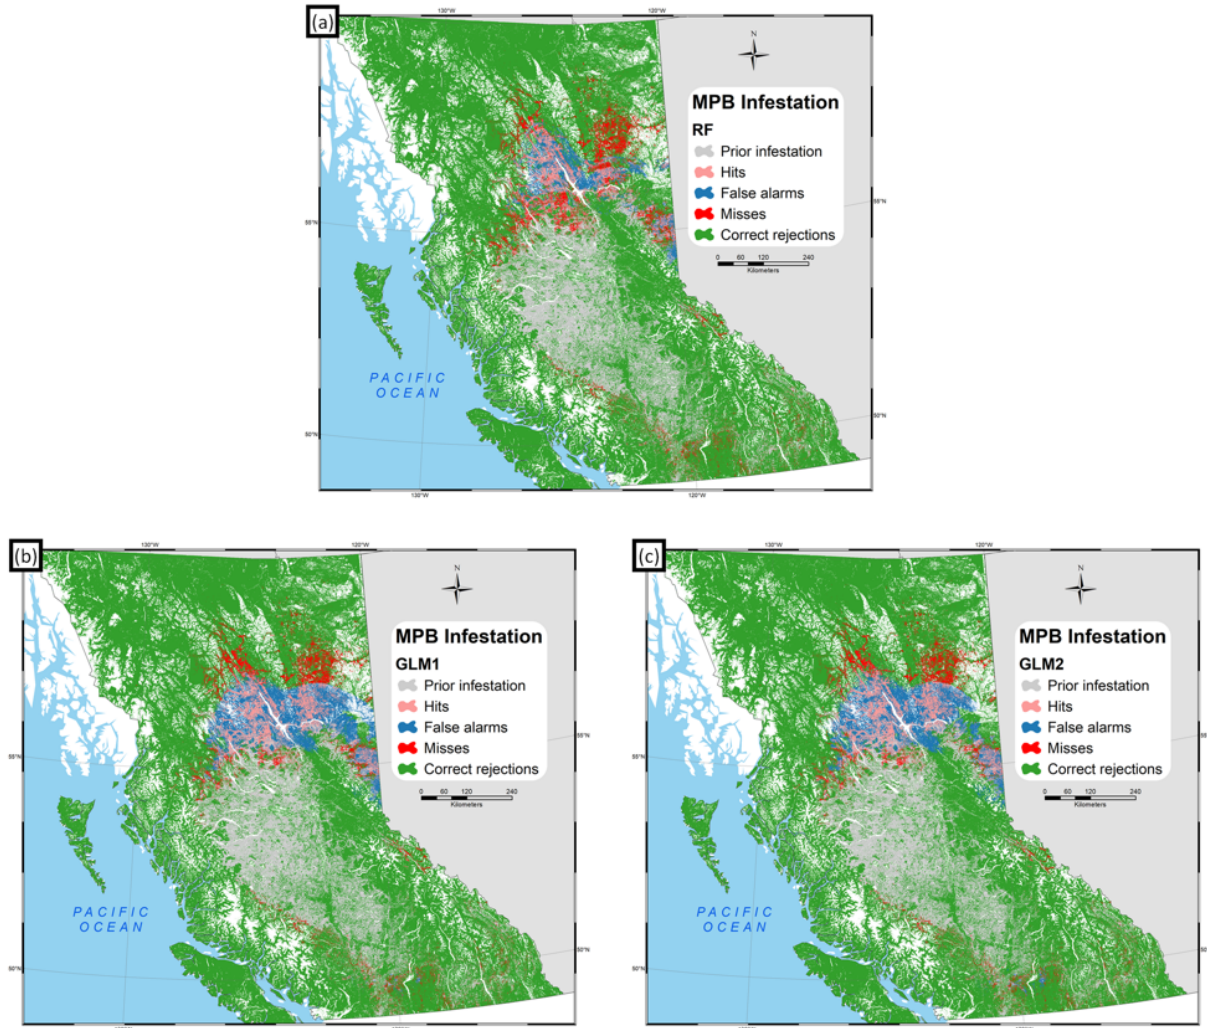


Figure 2.4. Validation analysis by comparison of simulated and observed changes in study area from 2008 to 2014. White color indicates no data. Algorithms: (a) RF, (b) GLM1, (c) GLM2. Hits are correct simulations of change. False alarm errors are persistence simulated as change. Miss errors are change simulated as persistence. Correct rejections are correct simulations of persistence. Infestations prior to the study are excluded from the study area.

Table 2.2 shows hits, misses and false alarms in numbers of pixels for each of the simulation algorithms. In addition, this table gives the figure of merit, which is the percentage of hits in the sum of hits, misses, and false alarms. The table also shows the overall accuracy of each model, which is the percentage of correct predictions in the study area.

Table 2.2. Comparison of hits, misses, and false alarms for different algorithms of simulation of change in study area from 2008 to 2014. GLM stands for Generalized Linear regression Model.

Algorithm	Cutoff	Hits (Pixels)	Misses (Pixels)	False Alarms (Pixels)	Figure of Merit (%)	Overall Accuracy (%)
-----------	--------	------------------	--------------------	--------------------------	------------------------	-------------------------

Binomial (GLM1)	Kappa	115,511	149,877	227,472	23.4	91.0
Binomial (GLM1)	Youden's J	171,359	94,029	433,560	24.5	87.4
Binomial-parabolic elevation (GLM2)	Kappa	120,653	144,735	214,695	25.1	91.5
Binomial-parabolic elevation (GLM2)	Youden's J	182,475	82,913	488,038	24.2	86.4
Random forest (RF)	Kappa	74,883	190,505	108,450	20	92.9
Random forest (RF)	Youden's J	90,631	174,757	144,737	22.1	92.4

Using the model results, which indicate locations of simulated infestations and the Ministry's map of pine volume density (BC Ministry of Forests, 2015a) in the study area, cumulative estimates of pine volume killed were calculated for each algorithm. Table 2.3 shows these estimates. In comparison, the real percentage of cumulative change based on the Ministry's map of observed infestations for year 2014 is 49% for the entire province (excluding the Tree Farm License Zone).

Table 2.3. Estimates of cumulative percentages of pine volume killed in MPB attacks by 2014 for three algorithms.

Algorithm	Cutoff	Cumulative Volume of Pine Killed (%)
Binomial (GLM1)	Kappa	57.5
Binomial (GLM1)	Youden's J	62.7
Binomial-parabolic elevation (GLM2)	Kappa	57.8
Binomial-parabolic elevation (GLM2)	Youden's J	63.3
Random forest (RF)	Kappa	54
Random forest (RF)	Youden's J	55

Figure 2.5 shows simulations of changes from 2014 to 2020, where the model was parameterized with observed changes from 2011 to 2014. The time step for all these simulations was 3 years. This is a demonstration of a possible application of the model for future predictions. Model outputs are projections of the infestations in the future, for which no reference data is available yet. As such, it is not possible to calculate the accuracy of these predictions. Rather, based on the results of validation tests described above and summarized in Table 2.2, we assume that the model is able to produce valid predictions. Aggregate percentages of cumulative pine volume killed are estimated for each simulation algorithm using the Ministry's pine volume density map, and are presented in Table 2.4.

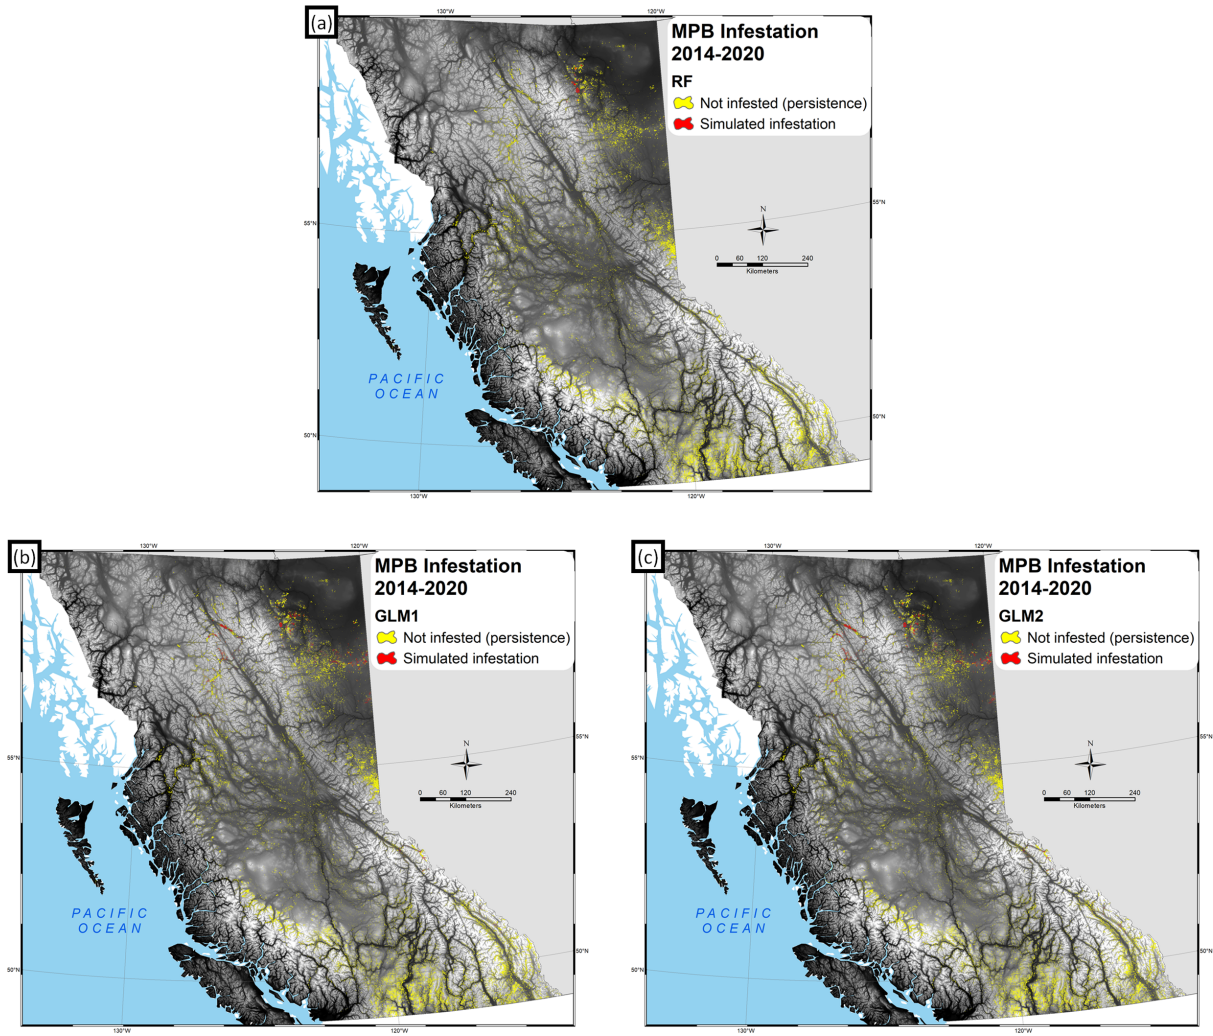


Figure 2.5. Predictions of changes in study area from 2014 to 2020. Algorithms: (a) random forest (RF); (b) binomial regression (GLM1); (c) binomial regression with parabolic elevation (GLM2).

Table 2.4. Estimates of cumulative percentages of pine volume killed in MPB attacks by 2020 for three algorithms.

Algorithm	Cutoff	Cumulative Volume of Pine Killed (%)
Binomial (GLM1)	Kappa	64.1
Binomial (GLM1)	Youden's J	70.5
Binomial-parabolic elevation (GLM2)	Kappa	64.2
Binomial-parabolic elevation (GLM2)	Youden's J	69.9
Random forest (RF)	Kappa	64
Random forest (RF)	Youden's J	64

2.4. Discussion

2.4.1. Model Validation

All of the simulation and observation images show a noticeable area of new infestations north of the center. A closer examination also reveals scattered infestations in other regions. Regarding the larger infested area, we know based on prior information that it is the result of the spread of previous infestations from the center of the province towards the north. A visual inspection suggests that, in the two GLM simulations, the southern and northern envelopes of the large infested areas are somewhat similar, as if the models have offset the prior infestations to the north. Considering that the model calculates the adjacency-type predictors in windows of a predefined size, it appears from Figure 2.3 that the GLM algorithms predicted a large number of infestations in a limited distance.

Comparing the images with respect to the large zone of new infestations north of the center, it appears that the GLM algorithms predicted a larger area for the spread of the insect. Moreover, the two GLM algorithms appear to have predicted a more relatively uniform spread inside the zone of new infestations. In contrast, the reference observation image shows a less homogeneous pattern in the same area. Regarding the RF algorithm, its result is somewhat non-homogeneous in that area. However, identifying correspondence between this image and the reference observation requires further analysis.

Figure 2.4 clearly shows that the larger parts of observed and simulated changes occurred in the area north of the center. It is noticeable that the two GLM algorithms predicted larger numbers of infestations in this area. Some of these predictions of infestation matched reference observation (hits), but many of them did not (false alarms). False alarms in the results of these two algorithms are visually distinctive. In contrast, the RF algorithm predicted a smaller spread of infestations and missed some infestations that actually occurred.

Table 2.2 reveals useful information about the performance of the models. Considering that hits indicate correct predictions of change, and that misses and false alarms indicate errors in prediction of change, the figure of merit is the proportion of correct predictions in all predictions involving change. This table confirms the above visual analyses of simulation outputs. The GLM algorithms predicted more change (hits + false alarms) than RF. Some of these excess

predictions were correct, and these algorithms had more hits than RF. However, many of them were incorrect. The GLM algorithms, especially with Youden's J final cutoff threshold, produced a large number of false alarms. In fact, although they have more hits and therefore higher figures of merit, this advantage is overshadowed by their error in estimating the quantity of change. The GLM algorithms, particularly with Youden's J final cutoff threshold, produced more false alarms than misses. This means that they overestimated the quantity of change. In contrast, the table shows that the RF algorithm underestimated the quantity of change, producing more misses than false alarms. With the kappa final cutoff threshold, the number of these errors in the estimation of quantity in RF is comparable with that of the GLM algorithms. With Youden's J final cutoff threshold, however, the underestimation of the quantity of change in the RF algorithm is relatively small.

Prior to producing prediction maps, each algorithm creates a probability map. Then it compares that map with the final cutoff threshold; that is, it classifies pixel values above the final cutoff threshold as 'Infested', and others as 'Not infested'. The Youden's J threshold was calculated to be lower value than that of the Kappa threshold for each algorithm. Therefore, for the same probability map, Youden's J cutoff predicts more infestations than the kappa cutoff. On the other hand, the table shows that regardless of the final cutoff threshold, the GLM algorithms tend to predict more change. This tendency, combined with the lower Youden's J final cutoff threshold, resulted in the large overestimation of quantity seen in the 2nd and 4th rows of the table. One avenue for improvement of the model in future works can be to reduce its error of estimated quantity of change.

As seen in Table 2.3, the random forest algorithm produces a closer estimate of the aggregate quantity of change in the study area. In comparison, the other two algorithms (GLM1 and GLM2) appear to overestimate this quantity. This is also confirmed in Table 2.2 and Figure 2.4; it may be noted that these algorithms produce more false alarms than misses. Having seen the model's performance in the validation test, it is easier to understand the model's prediction of future changes.

2.4.2. Model Predictions

Predictions of future infestations are summarized in Table 2.4. These calculations are based on the Ministry's previous inventory and do not include the effect of harvesting or

management decisions. Different simulations estimate that by 2020 between 64 and 70 percent of the pine volume in BC forests will be killed in MPB attacks. These assessments have been made for the entire forests—or the Whole Land Base (WLB)—of the province. The Ministry has made estimates of the spread of infestation in the Timber Harvesting Land Base (THLB), which is a smaller subset of the WLB. According to the Ministry the cumulative percentage of pine volume killed in the THLB could reach 55% by 2020 (BC Ministry of Forests, 2015b). Taking this difference in land base or study area into account, the results of our model can be considered as a complement to the Ministry’s calculations, particularly for zones outside the THLB. In both studies, it is predicted that in future the spread of infestations in the province will subside.

2.4.3. Limitations of the Study

Applying a threshold to transform the initial continuous infestation maps into binary maps may entail a loss of information because we discard valid knowledge about the shape of the infestation curve, i.e., whether it grows faster or slower until it reaches saturation. However, we have shown (Appendix 2.A) that, in general, infestation curves in all pixels seem to follow a similar sigmoidal pattern. Therefore, our choice of a single threshold to compute all binary maps appears to be a reasonable compromise between simplifying the analysis and retaining as much information relating to the infestation process as possible.

It is obvious that, for a process that spreads spatially, the state of the surroundings is important in determining the expansion rate. Presumably, definitions of adjacency other than those we used in the present study may lead to a better characterization of the infestation. Proximity laws based on insect phenology (e.g., flight time and strength, sensibility to extreme heat or sunlight) may yield better results than the generic adjacency functions used in this work. Future work may shed some light on this crucial topic.

2.5. Conclusions

In this work, we built a spatial change model by integrating neighborhood selection and transition rule identification in one process. Although we did not have detailed information about distance and neighborhood effects of MPB infestations at the beginning of this study, we compensated for this lack of knowledge by calculating several mathematically independent functions of distance as predictors in the model, and allowed the rule identification algorithm to find the transition rule that best described the input data. We validated our model by comparison

of its output with the observed data. Validation based on components of figure of merit gave us better insight into the performance of different algorithms applied in the model. In validation analyses, we noted that although the GLM simulations resulted in higher figures of merit, the RF algorithm was better able to estimate the quantity of change. We then used the validated model to predict the upcoming changes in the study area. By integrating neighborhood effects as variables in the parameterization of the model, we were able to simulate spatiotemporal complexities of a forest land change process.

Author Contributions: Conceptualization, L.P., R.M.H. and S.H.; methodology, L.P., R.M.H. and S.H.; software, R.M.H. and S.H.; validation, L.P. and S.H.; formal analysis, L.P., R.M.H. and S.H.; investigation, L.P., R.M.H. and S.H.; resources, L.P., R.M.H. and S.H.; data curation, S.H.; writing—original draft preparation, L.P., R.M.H. and S.H.; writing—review and editing, L.P., R.M.H. and S.H.; visualization, L.P. and S.H.; supervision, L.P. and R.M.H.; project administration, L.P.; funding acquisition, L.P. and R.M.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada through the Discovery Grant number RGPIN/05396–2016 awarded to LP. RMH received financial support from the European Union’s Seventh Framework Programme through NEWFOREST programme number PIRSES-GA-2013–612645.

Acknowledgments: We are thankful to two anonymous reviewers who helped improve this paper. We are thankful to the Université de Montréal’s International Affairs Office (IAO) for their financial support through the International Partnership Development program, which allowed the collaboration between researchers from UdeM and CREAM. We are also thankful to Mr. Graham Hawkins from the Ministry of Forests, Lands and Natural Resource Operations of British Columbia for his support in providing data, and to the Netherlands Environmental Assessment Agency (MNP) as the owner and RIKS BV as the developer of the Map Comparison Kit for providing access to the software. The analyses were made possible thanks to provision of high-performance computing by Calcul Quebec.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix 2.A. Calculation of Threshold Value on Cumulative *Pinus contorta* Mortality Data

To calculate the initial threshold value (i.e., the cutoff value that we must apply to the original infestation map), we first convert the original 0–1000 scale of the cumulative data into a 0%–100% scale. At every spatial pixel the infestation starts at a value close to 0% starting in the year 1999, which then increases quickly and finally arrives at saturation level (that is, there are no more trees left in that pixel to attack) near 100%. When sample curves of the cumulative infestation values are plotted as a function of time (see examples in Figure S2.1, left panel) we see that they approximately display a sigmoidal behavior, although the years the infestation starts are different.

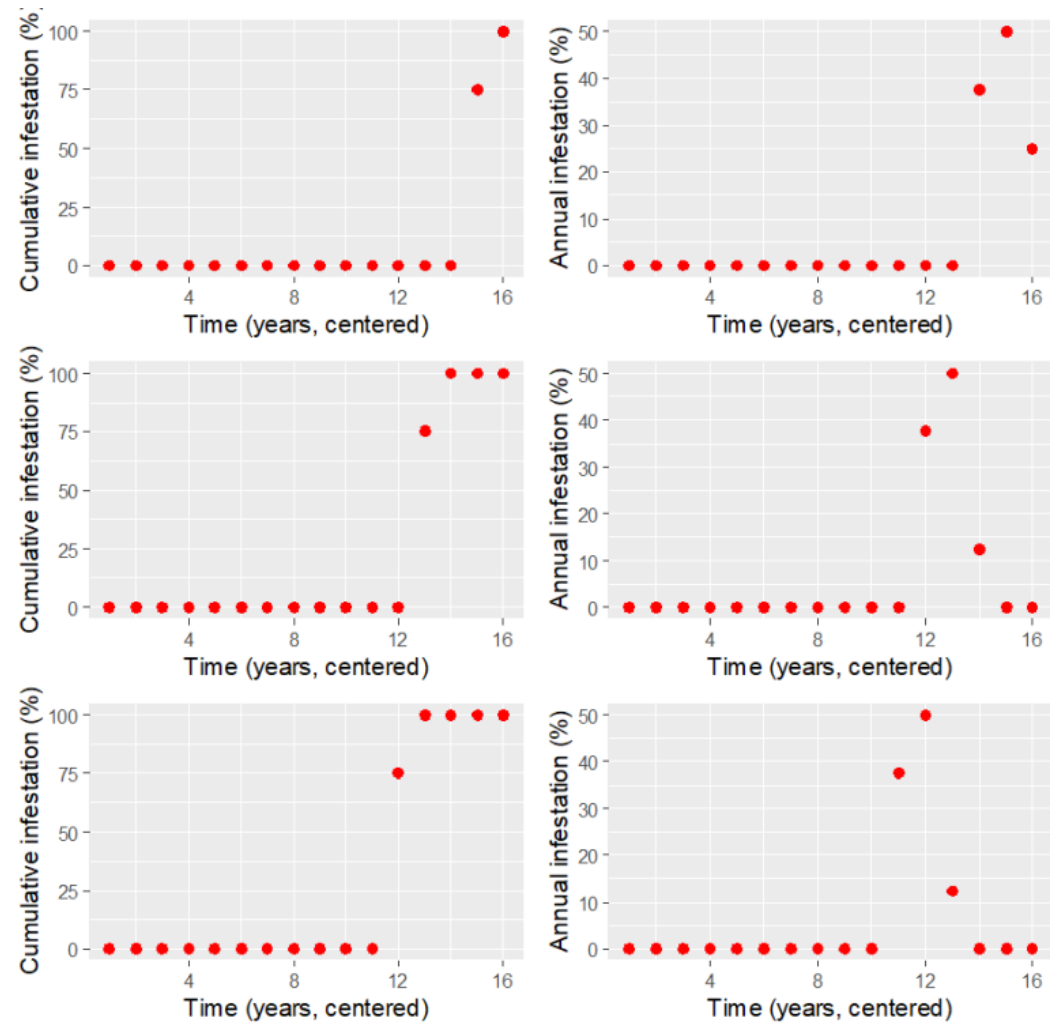


Figure S2.1. Three sample curves of cumulative (left panels) and annual (right panels) infestation percentages.

Next, we determine the annual infestation percentage by computing the numerical derivative of the cumulative curves. The resulting curves (Figure S2.1, right panel) now show a coarse bell-like shape with a given center and width.

If we conjecture that the curves represent unimodal and symmetric distributions, we can calculate, at each pixel i , (a) the centroid c_i as the first statistical moment, and (b) a proxy of the width w_i as the square-root of second moment (i.e., variance), as follows:

$$c_i = \frac{\sum_{j=1}^{16} p_{ij} \cdot t_j}{\sum_{j=1}^{16} p_{ij}}$$

and:

$$w_i = \sqrt{\frac{\sum_{j=1}^{16} p_{ij} \cdot (t_j - c_i)^2}{\sum_{j=1}^{16} p_{ij}}}$$

where p_{ij} stands for the infestation percentage for pixel i at year t_j ($t_1 = 1999, \dots, t_{16} = 2014$). We then use centroids c_i to re-center the offset of the original cumulative infestation curves. A random selection of 50,000 offset-corrected cumulative curves is shown in Figure S2.2. This figure shows the resulting re-centered data points, where it is now easier to make out an approximate sigmoidal shape for the cumulative infestation process. Notice that no correction for different widths has been carried out.

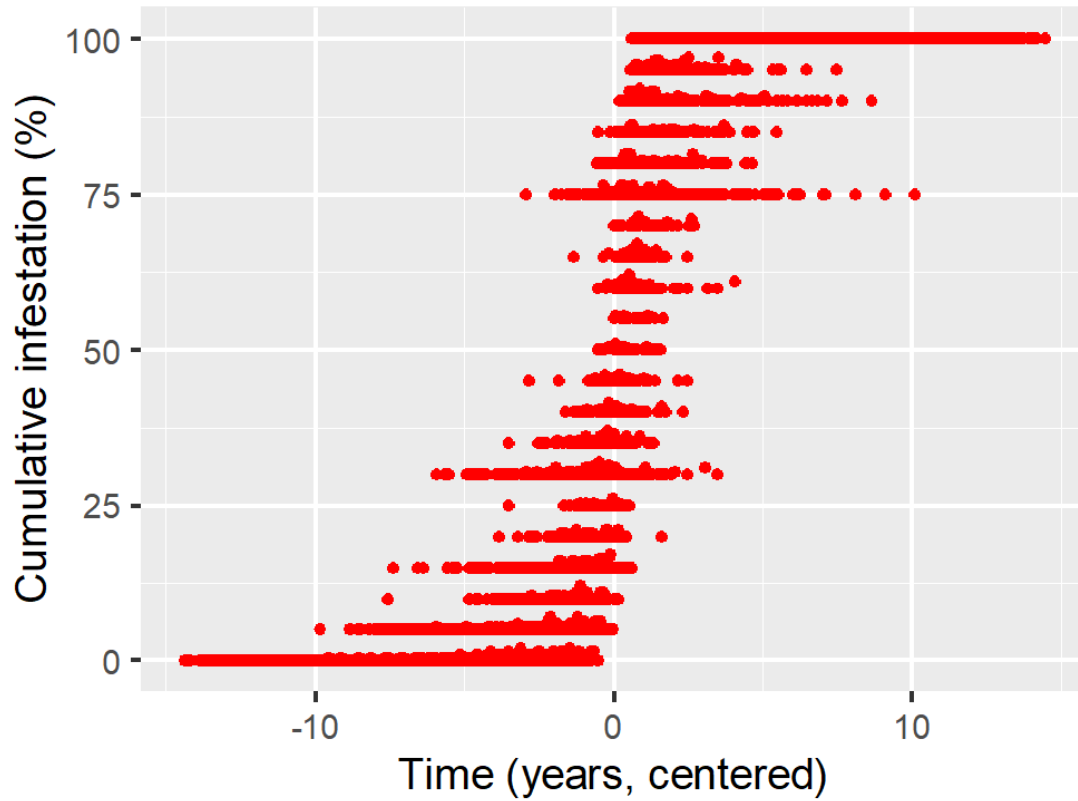


Figure S2.2. Offset-corrected cumulative infestation data points.

Next, we compute the mean cumulative infestation in small intervals along the temporal axis. When those mean infestation data points are plotted as a function of time (Figure S2.3, red dots), the sigmoidal-like shape of the infestation process becomes more conspicuous. Finally, we plot a normal cumulative distribution function (CDF, hereafter $I(t)$) with standard deviation $\sigma = \bar{w}$ (Figure S2.3, blue curve), where:

$$\bar{w} = \frac{1}{N} \cdot \sum_{i=1}^N w_i$$

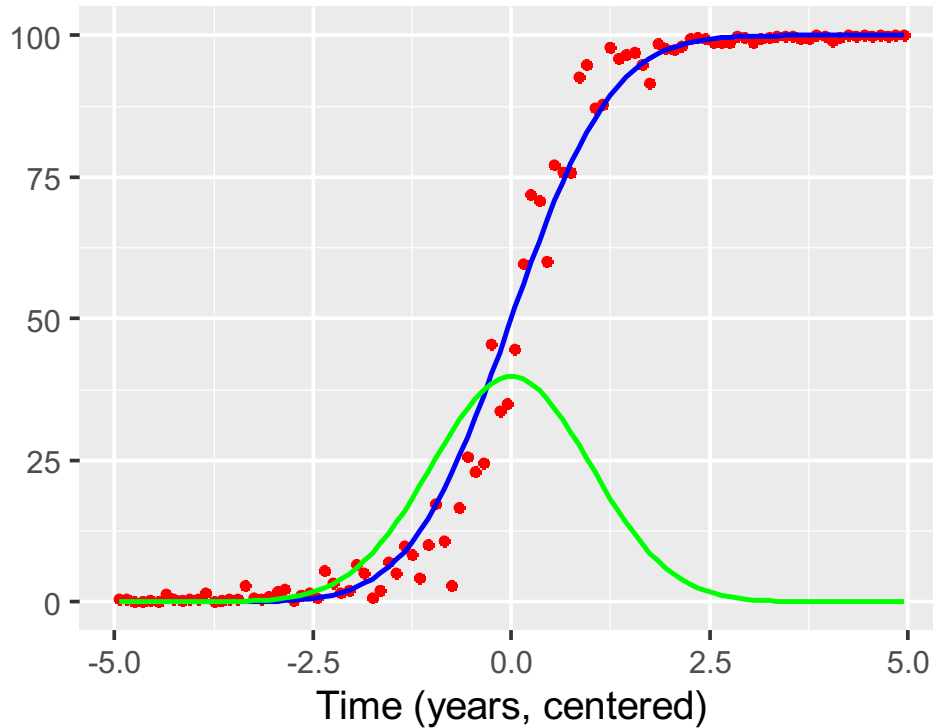


Figure S2.3. Mean observed and offset-corrected cumulative infestation (red dots), normal cumulative distribution function $I(t)$ (CDF, in blue, scaled to 0–100), whose standard deviation is explained in text, and its derivative $I'(t)$, i.e., the normal probability density function (in green).

For reference, we have also plotted (in green) the probability density function of the corresponding normal distribution, $I'(t)$, in Figure S2.3. Clearly, $I(t)$ serves as a good approximation to the observed cumulative infestation curve. In curves $I(t)$ and $I'(t)$ we can distinguish five phases as we move from left to right:

1. initial negligible or very low infestation that spreads slowly ($I(t)$ and $I'(t)$ are close to zero);
2. a transitional phase in which the infestation starts picking up speed ($I(t)$ still low but $I'(t)$ increases);
3. fast but steady infestation that increases constantly ($I(t)$ increases but $I'(t)$ reaches a maximum);

4. transitional phase during which the infestation slows down ($I(t)$ increases further, $I'(t)$ decreases);
5. saturation level ($I(t)$ highest, $I'(t)$ very close to zero).

We assume in the present study that a critical moment in the spread of the infestation occurs when, in phase 2 above, the infestation starts picking up speed, therefore accelerating its spread. Consequently, we will use the infestation level at which that acceleration is maximum as our threshold value I_{th} , which will enable us to categorize the continuous cumulative infestation maps into binary maps I_b such that:

$$I_b = \begin{cases} 0 \text{ (not infested)} & I \leq I_{th} \\ 1 \text{ (infested)} & I > I_{th} \end{cases}$$

The so-called acceleration can be calculated as the second derivative of $I(t)$. Therefore, to find its maximum we must compute in turn the third derivative and solve $I'''(t) = 0$. Assuming that $I(t)$ is well approximated by a normal CDF, we know that, for a centered normal distribution:

$$I'''(t) = \frac{(t^2 - \sigma^2) \cdot e^{-\frac{t^2}{2\sigma^2}}}{\sigma^5 \cdot \sqrt{2\pi}}$$

Equating $I'''(t) = 0$ we arrive at $t = \pm\sigma$, excluding the trivial solutions $t = \pm\infty$. Because $t = +\sigma$ corresponds to phase 4 above, i.e., the deaccelerating phase, we take $t_{th} = -\sigma$ as the temporal location of the sought-after threshold. Inserting t_{th} into $I(t)$ yields, finally, $I(-\sigma) = 0.1586553$ (or 15.86553 in our percentage scale). Figure S2.4 depicts the procedure.

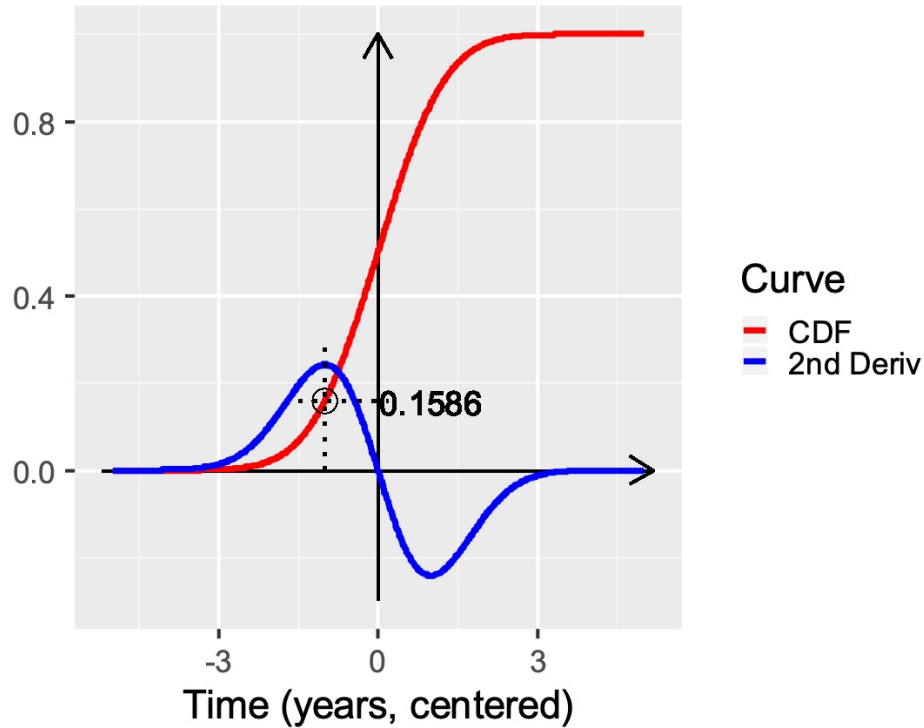


Figure S2.4. Cumulative distribution function (red) of the normal distribution and its 2nd derivative (blue). The location of the maximum of the 2nd derivative and the corresponding ordinate is indicated with a dashed line.

Appendix 2.B. Plots of Average Mortality vs. Predictors

Each predictor variable in Table 2.1 (see main text) was binned at consecutive intervals and at each interval we averaged over the corresponding 0 and 1 pixels in the map. This yielded the expected infestation rates, which were then logit-transformed and plotted vs. the center of the intervals, as shown below.

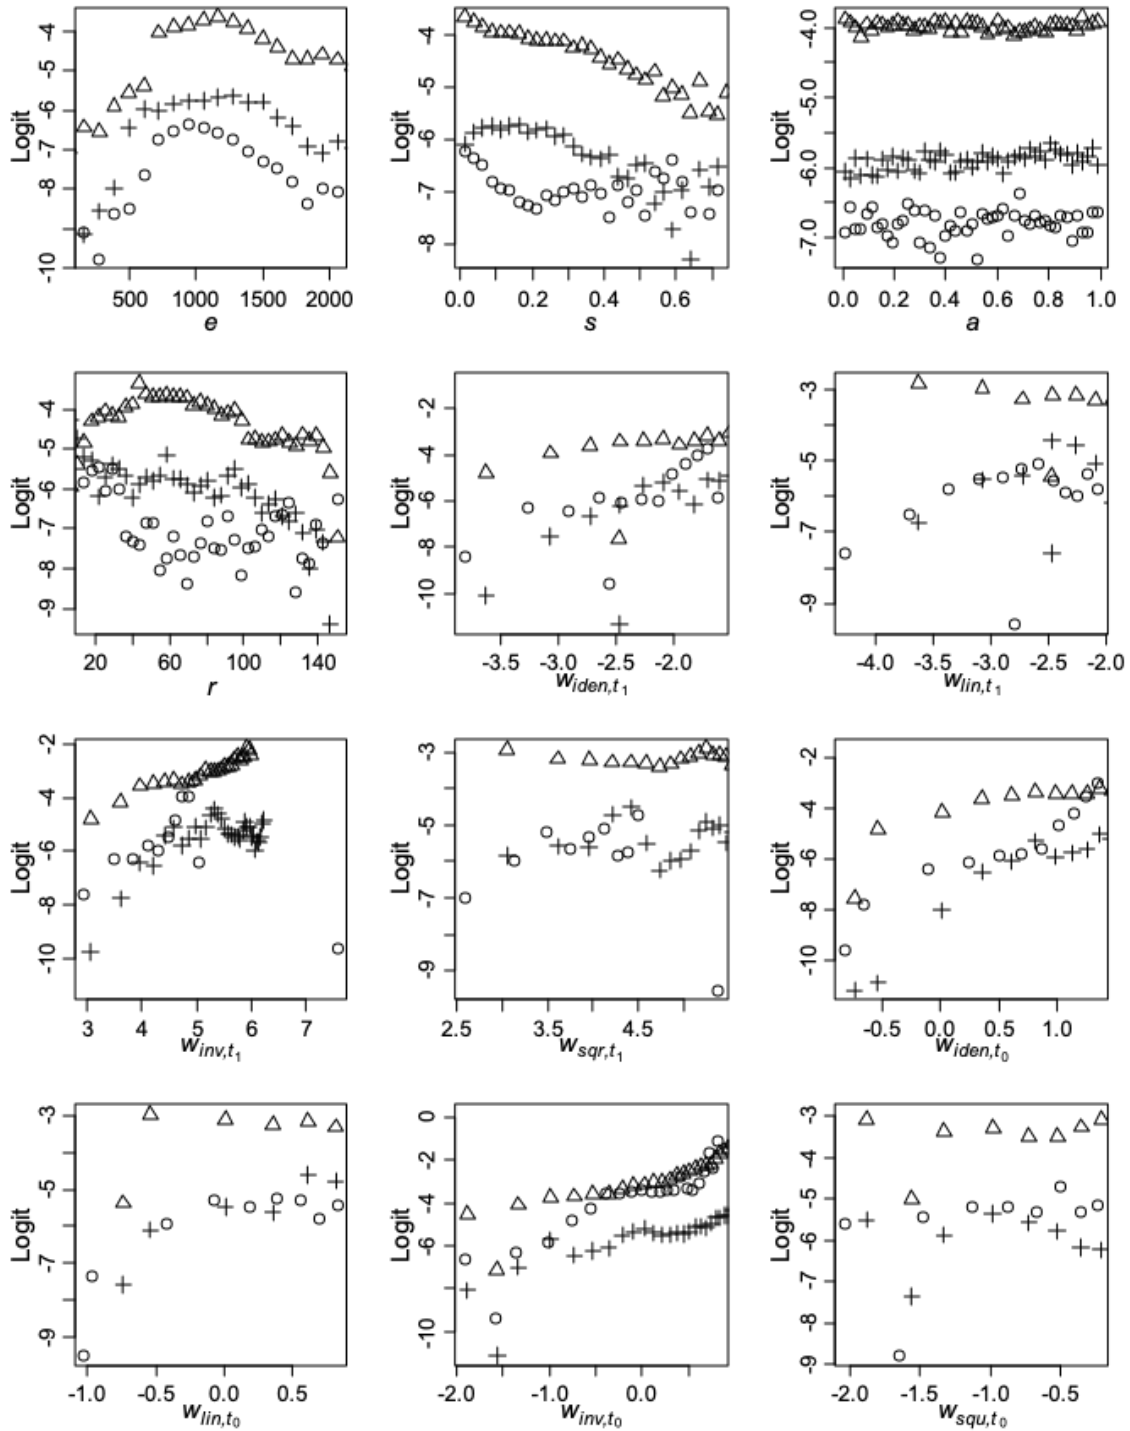


Figure S2.5. Plots of average mortality vs. predictors in three transition periods: 2001–2002 (circles), 2007–2008 (triangles), and 2013–2014 (plus signs).

Appendix 2.C. Graphic Representation of the Four Different Neighborhood Types Implemented




Figure S2.6. The four neighborhood types implemented in this model are (a) no-weight: neighborhood effect is uniformly distributed; (b) linear: neighborhood effect decreases linearly with increasing distance; (c) inverse-distance: neighborhood effect decreases inversely proportional to distance; and (d) squared-inverse-distance: neighborhood effect decreases inversely proportional to distance squared.

Appendix 2.D. Model Parameterization

The results of parameterization of the GLM1, GLM2, and RF models are summarized in this appendix.

The GLM1 algorithm resulted in a Receiver Operating Characteristic Area Under Curve (AUC) of 0.7802905. The maximum values of kappa and Youden's J obtained in model parameterization were 0.3489783 and 0.4540825, respectively. Statistics for this model's variables are presented in Table S2.1.

Table S2.1. Parameterization statistics for the GLM1 model.

Variable	Acronym	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-	$-2.518927 \times 10^{+00}$	7.437363×10^{-03}	-338.685488	$0.000000 \times 10^{+00}$
elevation	<i>e</i>	2.415789×10^{-04}	5.497418×10^{-06}	43.944072	$0.000000 \times 10^{+00}$
ruggedness	<i>r</i>	3.418483×10^{-04}	3.661161×10^{-05}	9.337157	9.895502×10^{-21}

aspect.sin (sine)	<i>a</i>	$-1.249106 \times 10^{-02}$	2.968032×10^{-03}	-4.208534	2.570333×10^{-05}
aspect.cos (cosine)	<i>a</i>	$-8.139534 \times 10^{-03}$	2.982952×10^{-03}	-2.728684	6.358761×10^{-03}
slope	<i>s</i>	$-2.682364 \times 10^{+00}$	1.821996×10^{-02}	-147.221146	$0.000000 \times 10^{+00}$
identity.1	$Z_{iden,t_{1p}}$	$-2.990441 \times 10^{+01}$	3.971936×10^{-01}	-75.289257	$0.000000 \times 10^{+00}$
linear.1	$Z_{lin,t_{1p}}$	6.979188×10^{-06}	1.514023×10^{-07}	46.096967	$0.000000 \times 10^{+00}$
inverse.1	$Z_{inv,t_{1p}}$	$-3.289507 \times 10^{-02}$	1.136399×10^{-03}	-28.946755	$3.083004 \times 10^{-184}$
squared.1	$Z_{squ,t_{1p}}$	$-9.227856 \times 10^{-02}$	4.722606×10^{-03}	-19.539752	5.042652×10^{-85}
identity.2	Z_{iden,t_1}	$3.895161 \times 10^{+01}$	3.035022×10^{-01}	128.340446	$0.000000 \times 10^{+00}$
linear.2	Z_{lin,t_1}	$-6.558742 \times 10^{-06}$	1.114139×10^{-07}	-58.868233	$0.000000 \times 10^{+00}$
inverse.2	Z_{inv,t_1}	1.187859×10^{-02}	8.019164×10^{-04}	14.812753	1.211739×10^{-49}
squared.2	Z_{squ,t_1}	1.341532×10^{-01}	3.245616×10^{-03}	41.333659	$0.000000 \times 10^{+00}$

The GLM2 algorithm resulted in an AUC of 0.7966494. The maximum values of kappa and Youden’s J obtained in model parameterization were 0.3603257 and 0.4610181, respectively. Statistics for this model’s variables are presented in Table S2.2.

Table S2.2. Parameterization statistics for the GLM2 model.

Variable	Acronym	Estimate	Std. Error	Z Value	Pr(> z)
(Intercept)	-	$-5.792189 \times 10^{+00}$	2.340479×10^{-02}	-247.478780	$0.000000 \times 10^{+00}$
elevation	<i>e</i>	6.035065×10^{-03}	3.824616×10^{-05}	157.795349	$0.000000 \times 10^{+00}$
ruggedness	<i>r</i>	9.025530×10^{-04}	3.765714×10^{-05}	23.967649	$6.049586 \times 10^{-127}$
aspect.sin (sine)	<i>a</i>	$-1.940276 \times 10^{-02}$	2.988868×10^{-03}	-6.491677	8.488584×10^{-11}
aspect.cos (cosine)	<i>a</i>	$-8.848104 \times 10^{-03}$	2.996480×10^{-03}	-2.952833	3.148722×10^{-03}
slope	<i>s</i>	$-2.494952 \times 10^{+00}$	1.827487×10^{-02}	-136.523626	$0.000000 \times 10^{+00}$
identity.1	$Z_{iden,t_{1p}}$	$-3.198807 \times 10^{+01}$	3.989017×10^{-01}	-80.190364	$0.000000 \times 10^{+00}$
linear.1	$Z_{lin,t_{1p}}$	6.664911×10^{-06}	1.514503×10^{-07}	44.007261	$0.000000 \times 10^{+00}$
inverse.1	$Z_{inv,t_{1p}}$	$-2.847800 \times 10^{-02}$	1.135279×10^{-03}	-25.084582	$7.326996 \times 10^{-139}$
squared.1	$Z_{squ,t_{1p}}$	$-9.944317 \times 10^{-02}$	4.717755×10^{-03}	-21.078495	1.253054×10^{-98}
identity.2	Z_{iden,t_1}	$3.926965 \times 10^{+01}$	3.044064×10^{-01}	129.004015	$0.000000 \times 10^{+00}$
linear.2	Z_{lin,t_1}	$-6.281823 \times 10^{-06}$	1.115482×10^{-07}	-56.314888	$0.000000 \times 10^{+00}$
inverse.2	Z_{inv,t_1}	1.006898×10^{-02}	8.026541×10^{-04}	12.544607	4.255184×10^{-36}
squared.2	Z_{squ,t_1}	1.190556×10^{-01}	3.241775×10^{-03}	36.725455	$2.866455 \times 10^{-295}$
I(elevation^2)	e^2	$-2.332311 \times 10^{-06}$	1.521917×10^{-08}	-153.248243	$0.000000 \times 10^{+00}$

The RF algorithm resulted in an AUC of 0.9999999. The maximum values of kappa and Youden’s J obtained in model parameterization were 0.9996198 and 0.9997796, respectively. The RF algorithm produces relative importance ranks for model variables. These ranks are presented in Table S2.3, with the most important variable given the score of 100.

Table S2.3. RF model variable ranks and relative importance scores rounded to two decimal digits.

Variable	Acronym	Relative Importance Score
inverse.2	Z_{inv,t_1}	100.00
identity.2	Z_{iden,t_1}	99.96
linear.1	Z_{lin,t_1p}	90.85
inverse.1	Z_{inv,t_1p}	90.23
linear.2	Z_{lin,t_1}	80.11
identity.1	Z_{iden,t_1p}	75.00
squared.1	Z_{squ,t_1p}	74.04
squared.2	Z_{squ,t_1}	61.32
elevation	e	22.43
slope	s	11.16
aspect.cos (cosine)	a	6.95
aspect.sin (sine)	a	6.76
ruggedness	r	3.83

Appendix 2.E. Code and Data Availability

Model code is available at <https://github.com/s-harati/model-MPB>. Data used in the study is accessible at <https://doi.org/10.17605/OSF.IO/V7ATJ>.

Appendix 2.F. Post-publication remarks

This appendix has been added in the present thesis chapter after the publication of the paper. This appendix includes additional information on the contents of the chapter based on remarks and comments received between publication in journal and final submission of this thesis. Where necessary, references have been cited in this appendix and added in the chapter's bibliography.

As noted in the introduction, the beetles excavate bores in the bark of the tree and lay eggs there. The larvae also excavate in the bark. Both the beetles and the larvae damage the tree's phloem. Moreover, MPB carries the blue stain fungi, which further damage the nutrient distribution system as well as the defence mechanism of the tree. The beetle and the fungi together kill the host tree in this manner. The flight of the insect occurs in the summer and the larval stage of its life is passed inside the tree in the cold season. Cold winter temperatures can substantially reduce survival rate of the larvae. Conversely, mild winter temperatures may contribute to substantial increase in the population of the insect. MPB attacks can be identified in aerial photos as they cause change in foliage color and cover. However, this insect does not feed on foliage; rather, it is a bark beetle (Natural Resources Canada, 2016).

In this work we focused on the recent outbreaks of the Mountain Pine Beetle in the forests of British Columbia. We assumed that some of the attacked trees become infested, and the infested trees die and remain dead for the duration of the simulation. In this sense, infestation was considered a one-way process in this study. The model of this study is constructed on grid cells of the map of the study area, and not on individual trees. As a consequence of assumption of infestation as a one-way change in the duration of the study, the infested cells are excluded from the calculation of future targets of the outbreak.

The scope of this study was simulation of spread of infestations at provincial scale. A previous work (Perez & Dragicevic, 2012) noted that Agent Based Models are powerful at simulating interactions between trees and insects at small geographical scales, but for larger scales they become computationally intensive. In such larger scales as the province of BC, spatial models were preferred (idem). Our choice of modelling approach was also informed by the type of data available to us, which was matrix maps of infestation in the province. These maps were composed of grids of cells, where the value of each cell represented the amount of trees that were damaged by the MPB in that cell. Given such data, it was reasonable to work with spatial models that produce matrix maps and can be calibrated and validated by comparison of their output with respective reference maps.

As a side note, the search for infestation data resulted in finding datasets of the years 1999-2014. These datasets were explored for gaining insight about the infestations and the extent and speed of their spread. They were also used in our preliminary works that were reviewed and modified several times before evolving into the present work. However, not all of those datasets were used in the development and testing of the model described in this chapter. Rather, as mentioned in the text, in this particular study, a model is built and trained with reference data on spread of infestations from 2005 to 2008 (hence a 3-year interval), and subsequently it was used to produce simulations of the spread of infestations from 2008 to 2011 and then to 2014. The end result, that is, the simulation of 2014, was compared with reference data for validation.

One of the most important factors in the simulation of spread of infestations was the weighting of potential locations of future infestations based on their distance from locations of previous infestations. This was done in the present study by introducing various distance functions. Indeed, this was an innovation in this work, as the distance functions were weighted

automatically through the calibration process, in which the model was trained and then tested against reference data. The definition of the distance functions of this study includes applying various distance-decaying functions (as well as a uniform function) in a circular area of radius d_{max} , which in the example case of this study was 150 pixels or 60 kilometers. This is the maximum range that was considered for the spread of infestations in one time-step.

A noteworthy point about this work is the large size of the study area. During the time interval of the study, spread of infestations comprises a relatively small part of the overall study area. That is why a large part of the study area remains not infested. As a result, all simulations show relatively high values of overall accuracy. The figure of merit focuses on areas where change occurred or was simulated, and shows very different values from the overall accuracy. However, the figure of merit, alone, is not descriptive enough of the details of the performance of the models. That is why instead of relying on a single combined dimensionless metric, several measures of agreement and disagreement of simulated change with reference change – namely, hits, misses and false alarms – are used in this chapter. The choice of these measures was also in part influenced by a desire to avoid eliminating information in model assessment. Specifically, a preference for quantities with dimensions over dimensionless quantities was considered in this choice. Such preference, especially with regard to showing areas of components of agreement and disagreement between simulated and reference change has been mentioned in model assessment literature as well (Pontius Jr & Si, 2014).

References

- Axelsson, J. N., Alfaro, R. I., & Hawkes, B. C. (2009). Influence of fire and mountain pine beetle on the dynamics of lodgepole pine stands in British Columbia, Canada. *Forest Ecology and Management*, 257(9), 1874–1882. <https://doi.org/10.1016/j.foreco.2009.01.047>
- Axelsson, J. N., Alfaro, R. I., & Hawkes, B. C. (2010). Changes in stand structure in uneven-aged lodgepole pine stands impacted by mountain pine beetle epidemics and fires in central British Columbia. *The Forestry Chronicle*, 86(1), 87–99. <https://doi.org/10.5558/tfc86087-1>
- BC Ministry of Forests. (2000). *Forest health aerial overview survey standards for British Columbia: the B.C. Ministry of Forests adaptation of the Canadian Forest Service's FHN Report 97-1 "Overview Aerial Survey Standards for British Columbia and the Yukon."* Retrieved from <http://www.for.gov.bc.ca/hts/risc/pubs/teveg/foresthealth/assets/aerial.pdf>
- BC Ministry of Forests. (2015a). [dataset] BC MPB Observed Cumulative Kill - vol.12.

- BC Ministry of Forests. (2015b). [dataset] BCMPB Summary of Volume Killed. Version 12.
- Bentz, B. J., Boone, C., & Raffa, K. F. (2015). Tree response and mountain pine beetle attack preference, reproduction and emergence timing in mixed whitebark and lodgepole pine stands. *Agricultural and Forest Entomology*, *17*(4), 421–432. <https://doi.org/10.1111/afe.12124>
- Bentz, B. J., & Jönsson, A. M. (2015). Modeling Bark Beetle Responses to Climate Change. *Bark Beetles*, 533–553. <https://doi.org/10.1016/B978-0-12-417156-5.00013-7>
- Bone, C., Dragicevic, S., & Roberts, A. (2005). Integrating high resolution remote sensing, GIS and fuzzy set theory for identifying susceptibility areas of forest insect infestations. *International Journal of Remote Sensing*, *26*(21), 4809–4828. <https://doi.org/10.1080/01431160500239180>
- Bone, Christopher, & Altaweel, M. (2014). Modeling micro-scale ecological processes and emergent patterns of mountain pine beetle epidemics. *Ecological Modelling*, *289*, 45–58. <https://doi.org/10.1016/j.ecolmodel.2014.06.018>
- Bone, Christopher, Wulder, M. A., White, J. C., Robertson, C., & Nelson, T. A. (2013). A GIS-based risk rating of forest insect outbreaks using aerial overview surveys and the local Moran's I statistic. *Applied Geography*, *40*, 161–170. <https://doi.org/10.1016/j.apgeog.2013.02.011>
- Bone, & Nelson. (2019). Improving Mountain Pine Beetle Survival Predictions Using Multi-Year Temperatures Across the Western USA. *Forests*, *10*(10), 866. <https://doi.org/10.3390/f10100866>
- Bourbonnais, M. L., Nelson, T. A., & Wulder, M. A. (2014). Geographic analysis of the impacts of mountain pine beetle infestation on forest fire ignition. *The Canadian Geographer / Le Géographe Canadien*, *58*(2), 188–202. <https://doi.org/10.1111/j.1541-0064.2013.12057.x>
- Carroll, A. L., Aukema, B. H., Raffa, K. F., Linton, D. A., Smith, G. D., & Lindgren, B. S. (2006). *Mountain Pine Beetle Outbreak Development : the Endemic - Incipient Epidemic Transition*.
- Chang, W.-Y., Lantz, V. A., Hennigar, C. R., & MacLean, D. A. (2012). Economic impacts of forest pests: a case study of spruce budworm outbreaks and control in New Brunswick, Canada. *Canadian Journal of Forest Research*, *42*(3), 490–505. <https://doi.org/10.1139/x11-190>
- Cooke, B. J., & Carroll, A. L. (2017). Predicting the risk of mountain pine beetle spread to eastern pine forests: Considering uncertainty in uncertain times. *Forest Ecology and Management*, *396*, 11–25. <https://doi.org/10.1016/j.foreco.2017.04.008>
- Coops, N. C., Wulder, M. A., & Waring, R. H. (2012). Modeling lodgepole and jack pine vulnerability to mountain pine beetle expansion into the western Canadian boreal forest. *Forest Ecology and Management*, *274*, 161–171.

<https://doi.org/10.1016/j.foreco.2012.02.011>

- Corbett, L. J., Withey, P., Lantz, V. A., & Ochuodho, T. O. (2016). The economic impact of the mountain pine beetle infestation in British Columbia: provincial estimates from a CGE analysis. *Forestry*, *89*(1), 100–105. <https://doi.org/10.1093/forestry/cpv042>
- ESRI. (2019). ArcGIS Pro 2.4.
- Ferretti, M. (1997). Forest health assessment and monitoring - issues for consideration. *Environmental Monitoring and Assessment*, *48*(1), 45–72.
- Flint, C. G., McFarlane, B., & Müller, M. (2009). Human dimensions of forest disturbance by insects: An international synthesis. *Environmental Management*, *43*(6), 1174–1186. <https://doi.org/10.1007/s00267-008-9193-4>
- Hall, R. J., Castilla, G., White, J. C., Cooke, B. J., & Skakun, R. S. (2016). Remote sensing of forest pest damage: a review and lessons learned from a Canadian perspective. *The Canadian Entomologist*, *148*(S1), S296–S356. <https://doi.org/10.4039/tce.2016.11>
- Haughian, S. R., Burton, P. J., Taylor, S. W., & Curry, C. (2012). Expected effects of climate change on forest disturbance regimes in British Columbia. *BC Journal of Ecosystems and Management*, *13*(1). Retrieved from <http://www.jem.forrex.org/index.php/jem/article/viewArticle/152>
- Hermes, D. a, & McCullough, D. G. (2014). Emerald ash borer invasion of North America: history, biology, ecology, impacts, and management. *Annual Review of Entomology*, *59*, 13–30. <https://doi.org/10.1146/annurev-ento-011613-162051>
- Hodge, J., Cooke, B., & McIntosh, R. (2017). *A strategic approach to slow the spread of mountain pine beetle across Canada*.
- James, P. M. A., & Huber, D. P. W. (2019). TRIA-Net: 10 years of collaborative research on turning risk into action for the mountain pine beetle epidemic. *Canadian Journal of Forest Research*, *49*(12), iii–v. <https://doi.org/10.1139/cjfr-2019-0384>
- Klenner, W., Walton, R., Arsenault, A., & Kremsater, L. (2008). Dry forests in the Southern Interior of British Columbia: Historic disturbances and implications for restoration and management. *Forest Ecology and Management*, *256*(10), 1711–1722. <https://doi.org/10.1016/j.foreco.2008.02.047>
- Lemmen, D. S., Warren, F. J., Lacroix, J., & Bush, E. (2008). *From impacts to adaptation : Canada in a changing climate 2007*. Ottawa: Government of Canada. Retrieved from <http://adaptation2007.nrcan.gc.ca>
- Lewis, M. A., Nelson, W., & Xu, C. (2010). A Structured Threshold Model for Mountain Pine Beetle Outbreak. *Bulletin of Mathematical Biology*, *72*(3), 565–589. <https://doi.org/10.1007/s11538-009-9461-3>

- Liang, L., Hawbaker, T. J., Chen, Y., Zhu, Z., & Gong, P. (2014). Characterizing recent and projecting future potential patterns of mountain pine beetle outbreaks in the Southern Rocky Mountains. *Applied Geography*, *55*, 165–175. <https://doi.org/10.1016/j.apgeog.2014.09.012>
- Liang, L., Li, X., Huang, Y., Qin, Y., & Huang, H. (2017). Integrating remote sensing, GIS and dynamic models for landscape-level simulation of forest insect disturbance. *Ecological Modelling*, *354*, 1–10. <https://doi.org/10.1016/j.ecolmodel.2017.03.007>
- Logan, J. A., & Powell, J. A. (2001). Ghost forests, global warming, and the mountain pine beetle (Coleoptera : Scolytidae). *American Entomologist*, *47*(3), 160–173.
- Macias Fauria, M., & Johnson, E. A. (2009). Large-scale climatic patterns and area affected by mountain pine beetle in British Columbia, Canada. *Journal of Geophysical Research*, *114*(G1), G01012. <https://doi.org/10.1029/2008JG000760>
- MacLean, D. A. (2016). Impacts of insect outbreaks on tree mortality, productivity, and stand development. *The Canadian Entomologist*, *148*(S1), S138–S159. <https://doi.org/10.4039/tce.2015.24>
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). London: Chapman and Hall.
- McCullough, D. G., Werner, R. A., & Neumann, D. (1998). Fire and Insects in Northern and Boreal Forest Ecosystems of North America. *Annual Review of Entomology*, *43*(1), 107–127. <https://doi.org/10.1146/annurev.ento.43.1.107>
- McIntire, E. J. B. (2004). UNDERSTANDING NATURAL DISTURBANCE BOUNDARY FORMATION USING SPATIAL DATA AND PATH ANALYSIS. *Ecology*, *85*(7), 1933–1943. <https://doi.org/10.1890/03-0579>
- Milne, M., & Lewis, D. (2011). Considerations for rehabilitating naturally disturbed stands : Part 1 – Watershed hydrology. *BC Journal of Ecosystems and Management*, *11*(3), 55–65.
- Natural Resources Canada. (2016). Mountain Pine Beetle. Retrieved from <http://www.nrcan.gc.ca/forests/fire-insects-disturbances/top-insects/13381>
- Patriquin, M. N., Wellstead, A. M., & White, W. A. (2007). Beetles, trees, and people: Regional economic impact sensitivity and policy considerations related to the mountain pine beetle infestation in British Columbia, Canada. *Forest Policy and Economics*, *9*(8), 938–946. <https://doi.org/10.1016/J.FORPOL.2006.08.002>
- Pelz, K. A., & Smith, F. W. (2012). Thirty year change in lodgepole and lodgepole/mixed conifer forest structure following 1980s mountain pine beetle outbreak in western Colorado, USA. *Forest Ecology and Management*, *280*, 93–102. <https://doi.org/10.1016/J.FORECO.2012.05.032>
- Perez, L., & Dragicevic, S. (2012). Landscape-level simulation of forest insect disturbance: Coupling swarm intelligent agents with GIS-based cellular automata model. *Ecological*

Modelling, 231, 53–64. <https://doi.org/10.1016/j.ecolmodel.2012.01.020>

- Pérez, L., & Dragičević, S. (2011). ForestSimMPB: A swarming intelligence and agent-based modeling approach for mountain pine beetle outbreaks. *Ecological Informatics*, 6(1), 62–72. <https://doi.org/10.1016/j.ecoinf.2010.09.003>
- Petersen, B., & Stuart, D. (2014). Explanations of a changing landscape: A critical examination of the British Columbia bark beetle epidemic. *Environment and Planning A*, 46(3), 598–613. <https://doi.org/10.1068/a4672>
- Pontius Jr, R. G., & Si, K. (2014). The total operating characteristic to measure diagnostic ability for multiple thresholds. *International Journal of Geographical Information Science*, 28(3), 570–583. <https://doi.org/10.1080/13658816.2013.862623>
- Pontius, R. G., Huffaker, D., & Denman, K. (2004). Useful techniques of validation for spatially explicit land-change models. *Ecological Modelling*, 179(4), 445–461. <https://doi.org/10.1016/j.ecolmodel.2004.05.010>
- Powell, J. A., Logan, J. A., & Bentz, B. J. (1996). Local projections for a global model of mountain pine beetle attacks. *Journal of Theoretical Biology*, 179(3), 243–260. <https://doi.org/10.1006/jtbi.1996.0064>
- R Core Team. (2019). R: A Language and Environment for Statistical Computing. Vienna, Austria, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org>
- Reid, R. W. (1962). Biology of the Mountain Pine Beetle, *Dendroctonus monticolae* Hopkins, in the East Kootenay Region of British Columbia II. Behaviour in the Host, Fecundity, and Internal Changes in the Female. *The Canadian Entomologist*, 94(06), 605–613. <https://doi.org/10.4039/Ent94605-6>
- Riley, S. J., DeGloria, S. D., & Elliot, S. D. (1999). A Terrain Ruggedness Index that Quantifies Topographic Heterogeneity. *Intermountain Journal of Sciences*, 5(1–4), 23–27.
- Robbins, J. (2008). Bark Beetles Kill Millions of Acres of Trees in West - NYTimes.com. *New York Times*, p. D3.
- Rutherford, G. N., Guisan, a., & Zimmermann, N. E. (2007). Evaluating sampling strategies and logistic regression methods for modelling complex land cover changes. *Journal of Applied Ecology*, 44(2), 414–424. <https://doi.org/10.1111/j.1365-2664.2007.01281.x>
- Safranyik, L., & Carroll, A. . (2007). The biology and epidemiology of the mountain pine beetle in lodgepole pine forests. In L. Safranyik & W. R. Wilson (Eds.), *The mountain pine beetle: a synthesis of biology, management, and impacts on lodgepole pine* (pp. 3–66). Victoria, British Columbia.
- Safranyik, L., Carroll, A. L., Régnière, J., Langor, D. W., Riel, W. G., Shore, T. L., ... Taylor, S. W. (2010). Potential for range expansion of mountain pine beetle into the boreal forest of

- North America. *The Canadian Entomologist*, 142(05), 415–442.
<https://doi.org/10.4039/n08-CPA01>
- Safranyik, L., Barclay, H., Thomson, A., & Riel, W. G. (1999). *A population dynamics model for the mountain pine beetle, Dendroctonus ponderosae Hopk. (Coleoptera : Scolytidae)*. Victoria, British Columbia.
- Shore, T., & Safranyik, L. (1992). Susceptibility and risk rating systems for the mountain pine beetle in lodgepole pine stands. Victoria, B.C: Forestry Canada.
- Strohm, S., Reid, M. L., & Tyson, R. C. (2016). Impacts of management on Mountain Pine Beetle spread and damage: A process-rich model. *Ecological Modelling*, 337, 241–252.
<https://doi.org/10.1016/J.ECOLMODEL.2016.07.010>
- Sturtevant, B. R., Achtemeier, G. L., Charney, J. J., Anderson, D. P., Cooke, B. J., & Townsend, P. a. (2013). Long-distance dispersal of spruce budworm (*Choristoneura fumiferana* Clemens) in Minnesota (USA) and Ontario (Canada) via the atmospheric pathway. *Agricultural and Forest Meteorology*, 168, 186–200.
<https://doi.org/10.1016/j.agrformet.2012.09.008>
- White, R., & Engelen, G. (1993). Cellular automata and fractal urban form: a cellular modelling approach to the evolution of urban land-use patterns. *Environment and Planning A*, 25(8), 1175–1199. <https://doi.org/10.1068/a251175>
- Wulder, M. A., Ortlepp, S. M., White, J. C., Nelson, T., & Coops, N. C. (2010). A provincial and regional assessment of the mountain pine beetle epidemic in British Columbia: 1999-2008. *Journal of Environmental Informatics*, 15(1), 1–13. <https://doi.org/10.3808/jei.201000161>
- Wulder, Michael A., Dymond, C. C., White, J. C., Erickson, B., Safranyik, L., & Wilson, B. (2006). Detection, mapping, and monitoring of the mountain pine beetle. In Les Safranyik & B. Wilson (Eds.), *The Mountain Pine Beetle: A Synthesis of Biology, Management, and Impacts on Lodgepole Pine* (pp. 123–154). Victoria, BC: Natural Resources Canada, Canadian Forest Service, Pacific Forestry Centre.

Chapter 3

Presentation of the article

The previous chapter presented a land change model for simulation of spread of forest insect infestations. In this doctoral work, the land change model described in Chapter 2 is used as a virtual laboratory for testing several management scenarios in Chapter 4. However, before proceeding with such tests, it is necessary for the spatial model to be validated with empirical reference data. To that end, I used existing methods for validation of land change models. My efforts to compare simulation results with reference data as well as with results of simple baseline models gave rise to new questions that could not be answered by existing methods. Particularly, it seems that existing methods focus only on the *count* and *type* of simulation errors, and do not take note of *where* errors occur. This means that existing methods lose some information in data aggregation. My endeavor to understand the strengths and weaknesses of the land change model of Chapter 2 resulted in new and improved methods of map comparison, which constitute a humble contribution to land change science. Chapter 3 describes these methods. These methods also allowed to gain a better insight about the studied process of forest infestation from available data. This insight was later used in Chapter 4 in defining an infestation control measure.

This chapter has been published in the peer-reviewed journal “Landscape Ecology” in 2021. My coauthors in this publication were my supervisors, Dr. Liliana Perez, Dr. Roberto Molowny-Horas, as well as Dr. Robert Gilmore Pontius Jr, who has developed the land change model assessment methods that I tried to improve. The chapter as it appears in this thesis involves modifications in the layout and style of the published paper, and slight modifications in figure and table numbers. Other than those, this chapter includes no changes in the content of the published paper.

Citation:

Harati, S., Perez, L., Molowny-Horas, R., Pontius Jr, R. G. Validating models of one-way land change: an example case of forest insect disturbance. *Landscape Ecol* **36**, 2919–2935 (2021). <https://doi.org/10.1007/s10980-021-01272-0>

Validating models of one-way land change: an example case of forest insect disturbance

Saeed Harati^{1,*}, Liliana Perez¹, Roberto Molowny-Horas², Robert Gilmore Pontius Jr³

¹ Laboratory of Environmental Geosimulation (LEDGE), Department of Geography, Université de Montréal, 1375 Ave. Thérèse-Lavoie-Roux, Montreal, H2V 0B3 QC, Canada

² CREAM, Centre de Recerca Ecològica i Aplicacions Forestals, E08193 Bellaterra (Cerdanyola de Vallès), Catalonia, Spain

³ School of Geography, Clark University, Worcester, MA 01610, USA

* Correspondence: saeed.harati.asl@umontreal.ca

Harati, S., Perez, L., Molowny-Horas, R., Pontius Jr, R. G. Validating models of one-way land change: an example case of forest insect disturbance. *Landscape Ecol* **36**, 2919–2935 (2021). <https://doi.org/10.1007/s10980-021-01272-0>

Abstract

Context: Validation of models of Land Use and Cover Change often involves comparing maps of simulated and reference change. The interpretation of differences between simulated and reference change depends on the characteristics of the process being studied. Our paper focuses on validation of models of one-way land change processes that spread in space. *Objectives:* Our objective is to develop a method for validation of one-way land change models, such that the method provides objective information about the spatial distribution of errors. *Methods:* Using distance analysis on reference data, we build a baseline model for comparison with simulations. We then simultaneously compare the four maps of reference at initial time, reference at final time, simulation at final time, and baseline at final time. We also use Total Operating Characteristic curves and multiple-resolution map comparison. We illustrate the methods with a simulation of forest insect infestations. *Results:* The methods give insights concerning the reference data, as well as to information concerning the spatial distribution of misses, hits, and

false alarms with respect to initial points of infestations. The new methods reveal that the simulations underestimated change near initial points of spread. *Conclusions:* The spatial distribution of errors is a topic of land change models that deserves attention. For models of one-way, geographically-spreading processes, we recommend that validation should distinguish between near and far allocation errors with respect to initial points of spread.

Keywords: Model validation, one-way change, Total Operating Characteristic, multiple resolution, distance analysis, area partition

3.1. Introduction

3.1.1 Validation of land change models

Modelling is a major theme of the Land Use and Land Cover Change (LUCC) science (de Sousa-Neto, Gomes, Nascimento, Pacheco, & Ometto, 2018; Lambin, Geist, & Rindfuss, 2006). Models are useful tools in science and policy applications, but before they can be used, their credibility should be evaluated (Pérez, Dragičević, & White, 2013; Robert Gilmore Pontius Jr, Boersma, et al., 2008; Robert Gilmore Pontius Jr, Huffaker, & Denman, 2004). Validation is an important step in the process of evaluation of any model. According to Rykiel (1996), “validation is a demonstration that a model within its domain of applicability possesses a satisfactory range of accuracy consistent with the intended application of the model”. Measuring a model’s accuracy and deciding a satisfactory level have been subjects of debate and disagreement in the scientific community (Verburg, Kok, Pontius Jr, & Veldkamp, 2006). Examples of intended uses of a model are to project future pathways of change and to run tests to understand processes of change and quantify our knowledge about them (Brown, Walker, Manson, & Seto, 2012; Lambin et al., 2006). It is essential to consider model objectives when validating the model (Batty & Torrens, 2005; Brown, Verburg, Pontius Jr, & Lange, 2013; Brown et al., 2012; National Research Council, 2014; van Vliet et al., 2016). LUCC literature recommends development of four classes of model testing methods: sensitivity analysis to test variation of model results with changing model parameters, uncertainty analysis to test variation of stochastic model results without changing model parameters, structural validation for models built with the aim of understanding change processes, and pattern validation for models aiming at prediction of future changes (Brown et al., 2013; National Research Council, 2014). The latter category usually employs techniques of map comparison (Brown et al., 2013; Foody, 2004;

Hagen-Zanker, 2006; National Research Council, 2014; Robert Gilmore Pontius Jr, 2000; White, 2006). This category of validation methods is further classified into two groups: tests of erroneous composition of land classes i.e. quantity disagreement, and tests of erroneous configuration of land classes i.e. allocation disagreement (Robert Gilmore Pontius Jr & Millones, 2011; van Vliet et al., 2016). A review of 114 models in land change literature found that the most common validation approach among them was location accuracy (van Vliet et al., 2016).

3.1.1.1 Definitions

Multiple disciplines apply techniques of map comparison and model validation – some examples are urban growth (Pijanowski et al. 2005; Chen and Pontius Jr 2010), landscape ecology (Cushman, Macdonald, Landguth, Malhi, & Macdonald, 2017; Paudel & Yuan, 2012), forestry (Rollins, Keane, & Parsons, 2004), agriculture (Li, Huffman, Zhang, Zhou, & McConkey, 2012), conservation (Hermoso, Morán-Ordóñez, & Brotons, 2018), and remote sensing (Foody, 2004) – with each discipline using its own terminology. In order to avoid confusion about meanings of the terms used throughout the text, we define them in this section.

The basis of the analyses described in this paper is the contingency table, also known as the confusion matrix or the error matrix. This table summarises the comparison of two binary maps of change. Generally, in model validation, one of the maps shows the change simulated by the model, and the other shows the reference change for comparison. The professional convention is that the rows of the contingency table indicate the simulation change, and the columns indicate the reference change. The value of each cell of the table is the count of units, e.g., pixels, with that cell's row category in simulation and its column category in reference. The sum of values in a row of the contingency table equals the quantity of the respective category in the simulation. Likewise, the sum of values in a column of the contingency table equals the quantity of the respective category in the reference. Finally, the sum of all cells of the contingency table equals the total count of units, e.g. pixels, in the study area. It is also possible to divide all values in the table by this total count, which is what Figure 3.1 shows. Then, the values will be proportions of the study area, and their total will be 1.

		Reference		
		Change	Persistence	Sum
Simulation	Change	<i>Hits</i>	<i>False alarms</i>	<i>Simulation change</i>
	Persistence	<i>Misses</i>	<i>Correct rejections</i>	$1 - \textit{Simulation change}$
	Sum	<i>Prevalence</i>	$1 - \textit{Prevalence}$	1

Figure 3.1. Outline of a contingency table

Figure 3.1 shows the layout of a contingency table with two categories: *Change* and *Persistence*. The terms *Hits*, *Misses*, *False Alarms* and *Correct Rejections* are used in this text. For clarification, they are defined in Table 3.1.

Table 3.1. Definitions

Term	Definition
<i>Hits</i>	Cases where change is simulated correctly. Also known as true positives.
<i>Misses</i>	Cases where reference change is simulated as persistence. Also known as false negatives.
<i>False Alarms</i>	Cases where reference persistence is simulated as change. Also known as false positives.
<i>Correct Rejections</i>	Cases where persistence is simulated correctly. Also known as true negatives.

The above terms are defined as proportions, and sum to 1. In this case, the proportion of area of reference change (hits+misses) is called *Prevalence*. However, it is equally justifiable to define the terms of Table 3.1 as size. In that case, their sum is the total size of the study area, and the size of area of reference change (hits+misses) is referred to as *Abundance*. In this study, we use various tools of analysis. Our main and concluding analysis is better interpreted when the components of map comparison are described as proportions of the study area. Still, one of our other analyses uses a tool that calculates respective areas of those components. Describing each analysis, we clarify how it interprets components of map comparison. More detailed explanation about the contingency table can be found in several references, such as Pontius Jr and Parmentier (2014) in LUCC modelling literature, and Congalton (2004) in remote sensing literature.

3.1.2 Scope and objectives

The scope of this paper is the validation of models of land change with a single transition and spatiotemporal dependency. Within this scope, we present a method for assessment of strengths and weaknesses of models in comparison with a baseline. In this section, we first

describe the above-said scope. Then, we explain the reason for using baselines in model assessment. Next, we highlight challenges and problems of using baselines. Finally, we define the objective of the paper with respect to the above-said challenges and problems.

We begin by describing the scope. A single transition is a one-way change between two land classes. This means that the subset of land change models within the scope of this study includes binary models of processes of change that are one-way during the study time interval. Beyond the scope of this study, in models that involve gain of several classes, or in models that involve simultaneous gain and loss of the same class in different parts of the study area, correct prediction of change is more challenging (Pontius Jr et al. 2018), and model assessment is accordingly more complicated. Figure 3.2 demonstrates a hypothetical example, which shows an important difference between models within the scope of this study and models beyond the scope of this study. Figure 3.2 shows that for models beyond the scope of this study, even if a place appears unchanged, it may have undergone change several times. As such, assessment of models beyond the scope of this study involves an additional challenge in interpretation of maps of change. In contrast, the definition of scope of this study involves a simplifying assumption regarding land classes and transitions. Note that an implication of the assumption of one-way change is that the study area will be the region that was unchanged at the beginning of the time interval of validation of the model, as this region is the only candidate for change during the simulation.

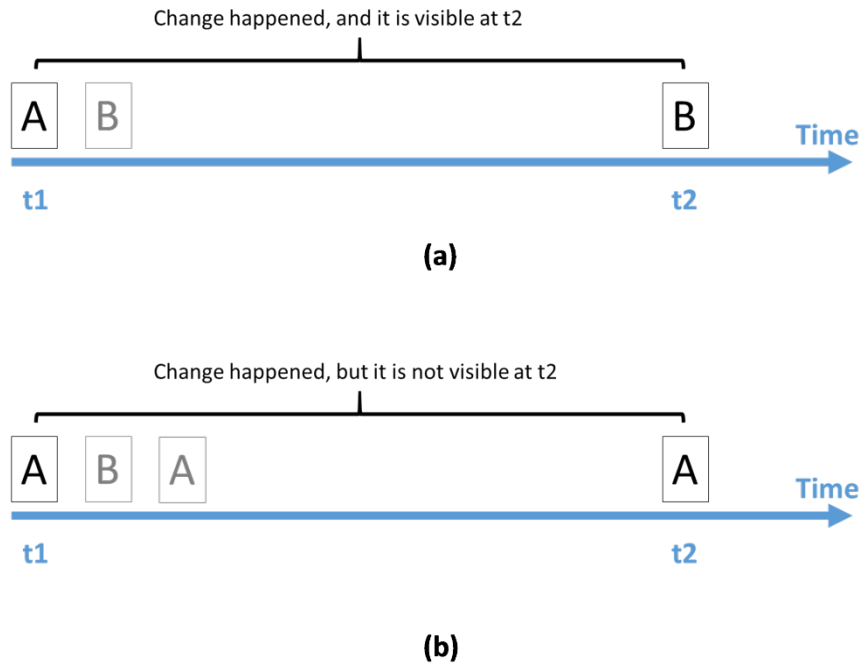


Figure 3.2. Implications of reversibility and irreversibility in interpretation of assessment data: (a) in a one-way process, a change of state during the study period can be identified by comparing reference data at the beginning and end of the period; (b) in a reversible process, change of state may happen during the study period without a trace in the reference data at the beginning and end of the period

Another assumption that we make in defining the scope of this study is spatiotemporal dependency. Spatiotemporal dependency occurs when places influence and are influenced by their surrounding neighborhoods. This is an implication of the first law of geography (Tobler, 1970), which states that things that are nearer to each other tend to be in a stronger relationship with each other. Some examples of applications within the scope of this study are forest fires (Gaudreau, Perez, & Drapeau, 2016), forest insect infestations (Perez, Molowny-Horas, & Harati, 2016), deforestation (Robert Gilmore Pontius Jr, 2018), urban expansion (Pijanowski et al., 2005), and spread of contaminations (Di Gregorio, Serra, & Villani, 1997). These examples, like many other phenomena in geography and landscape ecology, involve spatiotemporal dependency. Moreover, the process of change in each of these examples is one-way.

The combination of spatiotemporal dependency and one-way change has an implication about the above applications: it implies that these applications involve neighborhoods where change spreads. In the above applications, there are places which, at the beginning of the study time interval, can transmit the phenomenon to their surroundings. In other words, the

phenomenon is present in those places at the beginning of the study. Throughout this paper, we refer to these places as the initial points of spread. For example, in applications such as forest fires or forest insect infestations, initial points of spread are places that underwent change immediately before the beginning of the study, such that at the start of the study, fire or insects were present in those places. In other applications such as deforestation, urban expansion, or soil contamination, initial points of spread are on the border between the unchanged and changed zones. In any case, the modeler can define the initial points of spread based on the nature of the studied phenomenon. We use the idea of initial points of spread in our methods.

Having described the scope of the study, we now explain the rationale and issues regarding the use of baselines in model assessment. It is desirable to identify strengths and weaknesses of models. However, strength and weakness are relative terms, and they make sense when they are expressed in a comparison of two or more things. It is a known approach of model assessment to compare the results of the model with another model as a baseline, also known as a benchmark (van Vliet et al., 2016). This comparison reveals whether the performance of the model is better or worse than the baseline. The use of baselines in assessment of LUCC models is not rare; in fact, in their review of 114 modeling works in LUCC literature, van Vliet et al. (2016) reported that 30% of the studied works involved baselines or benchmarks. Two questions need to be addressed before a baseline model can be used for assessment: (1) how to define the baseline model, and (2) how to compare the baseline model with the simulation. Regarding the first question, the modeler defines the baseline according to common sense and the particularities of the subject of study. For models within the scope of this paper, we define the baseline using the idea of neighborhoods around initial points of spread. We discuss this matter in detail in the methods section. Regarding the second question, some works in land change literature assess the baseline model in the same way that the simulation is assessed, and then compare the results of the two assessments (Robert Gilmore Pontius Jr, 2018; Robert Gilmore Pontius Jr et al., 2007). This approach reveals useful information. For example, it answers questions such as: *Does the simulation make more Hits than the baseline? Which model makes more False Alarms? Which model makes more Misses?*

Although the above-said approach provides useful information for model assessment, it has a shortcoming, as it cannot answer more detailed questions such as: *How many of the Hits of*

the simulation are also Hits of the baseline? How many of the Hits of the simulation are missed by the baseline? How many of the Hits of the baseline are missed by the simulation? Answering these questions reveals useful details about strengths and weaknesses of the model with respect to the baseline. For example, Hits of the simulation that are missed by the baseline show relative strength of the simulation with respect to the baseline. Moreover, for models within the scope of this paper, we want objective information about where the errors occur with respect to initial points of spread. To that end, we are interested in finding answers to questions such as: *How much of the model's correct predictions of change are near initial points of spread? What is the dominant type of error far from initial points of spread? Does the model underestimate or overestimate change near or far from initial points of spread?* These questions are not answered by existing methods of model assessment, and we address them in this paper.

Considering the above, the objectives of this study are: (1) to present a method for building a baseline model for applications within the scope of this study; and (2) to present a method for comparing a simulation with a baseline, in such detail that the comparison of the simulation and baseline provides answers to the questions raised in the previous paragraph. We demonstrate our methods by applying them to an example case of study, and we discuss the results that our methods reveal about the example case. We also apply existing methods in the validation of our example case, in order to compare the existing methods with our methods.

3.1.3 Case study

The motivation for this work comes from previous efforts to evaluate a model of forest insect infestations (Harati, Perez, & Molowny-Horas, 2020). The Mountain Pine Beetle (MPB) is a native wood-boring insect infesting forests of western North America. During the past two decades, MPB outbreaks have become epidemic and damaged over half of the commercial pine volume in the province of British Columbia (BC), Canada (Natural Resources Canada, 2019). Each summer, the insects fly in search of new hosts to infest. In the earlier years of the outbreaks, the eastward spread of infestations was limited by the Rocky Mountains, but eventually MPB crossed this natural barrier and arrived in the neighboring province of Alberta (Natural Resources Canada, 2019). Contrary to infestation, which kills trees, are processes of forest regrowth and succession. However, these processes are much slower than the spread of insect infestations. As such, in our case of modeling MPB outbreaks in a 6-year period, we

assume that the processes of succession did not have enough time to make any change in infested forests. For this reason, we consider infestation a one-way process of land change.

In a previous study, we developed a model to simulate the spread of the insect in BC (Harati et al., 2020). The predictor variables of the MPB model were elevation, aspect, slope, surface ruggedness, and sums of surrounding infestations weighted by four different distance functions, calculated for infestation data of the year of start of simulation and its preceding year. We used various algorithms and calibration settings to train the model, and used it to predict the spread of infestations from 2008 to 2014, hence producing one simulation for each algorithm. The scope of the present paper is assessment and validation of two simulations produced by the MPB model, and their comparison with a third model that is defined as a baseline. Therefore, development of simulation algorithms and adjustment of their parameters are not within the scope of this paper. Rather, we discuss how we can compare simulation outputs with reference datasets in order to obtain information that is useful in model assessment.

3.2 Methods

Our model assessment approach in this study is to compare the simulation with a baseline. As suggested by Pontius Jr et al. (2007), we select our baseline as a model that is easy to understand. For models with a single one-way transition and spatiotemporal dependency, we define the baseline as a neighborhood in proximity of initial points of spread. Our baseline model predicts that areas near initial points of spread will change, and areas far from initial points of spread will not change. In this sense, the baseline model divides the study area into two strata with respect to proximity to initial points of spread. In the following subsections we describe how we build the baseline model, and how we compare it with the initial reference, final reference and simulation simultaneously. Through such simultaneous comparison of four maps, we obtain useful information for model assessment. In addition, to distinguish between our analysis and existing methods, we note highlights about the existing method of multiple resolution map comparison, which is used for analyzing allocation errors.

3.2.1 Distance analysis and proximity suitability map

The concept of our proximity suitability map is that places adjacent to initial points of spread are the most suitable candidates for future changes in the entire study area; in turn, places adjacent to those candidates are the next suitable ones; and so on. In this concept, suitability

decreases as distance from initial points of spread increases. For ease of computation, we use Manhattan distances in making the proximity suitability map. By overlaying the reference dataset of the end of validation time interval on the map of Manhattan distances from initial points of spread, we extract information about percentage of reference infestations in various classes of distance from initial points of spread. This information provides a better insight into the case study.

3.2.2 Total Operating Characteristic curves

We classify simulation and proximity suitability maps using thresholds. The classified predictions depend on their respective classification threshold. Such dependency is different for each model. We gain information about the performance of the models by analyzing their agreement and disagreement with reference data for each classification threshold. We perform and summarize this analysis using the Total Operating Characteristic (TOC) curve (Robert Gilmore Pontius Jr & Si, 2014).

The TOC curve is a tool for analyzing agreement between the ranked output of a model with a binary reference. In this sense, its overall aim is similar to that of the Relative Operating Characteristic (ROC) curve. The TOC curve is obtained by connecting points that each represent a threshold for classification of model output. Corresponding to each threshold, there is a contingency table that summarizes agreement and disagreement between model output and reference data. Misses, hits, false alarms, and correct rejections for each classification are directly identified in the TOC curve as shown in Figure 3.3. This is because of the way the plot coordinates and axes are defined. The TOC shows misses, hits, false alarms, and correct rejections as sizes; their sum equals the size of the study area. For each point of the TOC curve, the horizontal coordinate equals the sum of hits and false alarms of that point's respective contingency table, and the vertical coordinate equals hits. Similar to the ROC curve, the Area Under Curve (AUC) of the TOC curve offers a metric to summarise accuracy. In the TOC plot, a parallelogram bounds all possible TOC curves, as shown in Figure 3.3. In this plot, the AUC is the ratio of the area under TOC curve in the parallelogram to the total area of the parallelogram.




Figure 3.3. A hypothetical TOC curve. In this plot, hits, misses, false alarms, and correct rejections are measures of area. Abundance is the area of reference change

To find the threshold for each model, we identify the intersection of its TOC curve with the vertical line, $x = Abundance$, which is drawn down from the upper left corner of the parallelogram. On that vertical line, misses and false alarms are equal. Recalling that the sum of misses and hits is the quantity of reference change, and that the sum of false alarms and hits is the quantity of simulated change, it follows that at this intersection point, the quantity of change from the start to the end of the validation time interval equals the quantity of change in the reference data.

The distance threshold that we find for the proximity suitability map using the TOC curve is the key to define near and far with respect to initial points of spread. This threshold classifies the proximity suitability map into a binary map. That is, in the proximity suitability map, values less than or equal to that threshold are classified as *changed*, and other values are classified as *unchanged*. We use this classified map as the baseline model.

3.2.3 Three-map comparison at multiple resolutions

In order to assess performance of the models, existing methods compare simultaneously three maps: reference at the beginning of the validation time interval, reference at the end of the validation time interval, and simulation at the end of the validation time interval. Such three-map comparison, which is well described in literature (Robert Gilmore Pontius Jr, Boersma, et al., 2008; Robert Gilmore Pontius Jr et al., 2004), is based on the contingency table of simulated change versus reference change. In applications with change that is one-way during the

validation time interval, places that underwent change before the beginning of the study are excluded from the study area, as they are no longer candidates for change.

By classifying suitability maps, we obtain maps in which quantification error is eliminated or minimized. The error that remains in model outputs is of allocation type. This implies that the size of misses equals the size of false alarms. In this case, for each change that the model misses, the model commits a false alarm. In model assessment, we are interested in knowing the distances between such pairs of miss and false alarm errors. In comparison of two models whose quantification error is eliminated, the one whose allocation errors occur in shorter distances performs better. In order to include this consideration in model assessment, we use the method of three-map comparison at multiple resolutions (Robert Gilmore Pontius Jr, 2002; Robert Gilmore Pontius Jr, Boersma, et al., 2008; Robert Gilmore Pontius Jr et al., 2004; Robert Gilmore Pontius Jr, Peethambaram, & Castella, 2011). In this method, fine-resolution maps are coarsened to a given resolution; then, they are compared. Throughout coarsening, several fine pixels are aggregated in a coarse pixel. In the coarse pixel, information concerning quantities remain, but information concerning allocation is eliminated. This effect is shown in Figure 3.4. If two fine pixels containing a miss and a false alarm are put together in a coarse pixel, coarse resolution comparisons show that both the coarsened reference and the coarsened simulation contain one unit of change and one unit of persistence in the aggregated pixel. As such, through coarsening, for each pair of miss and false alarm errors aggregated in a larger pixel, the counts of misses and false alarms each decrease by one, and the counts of hits and correct rejections each increase by one.

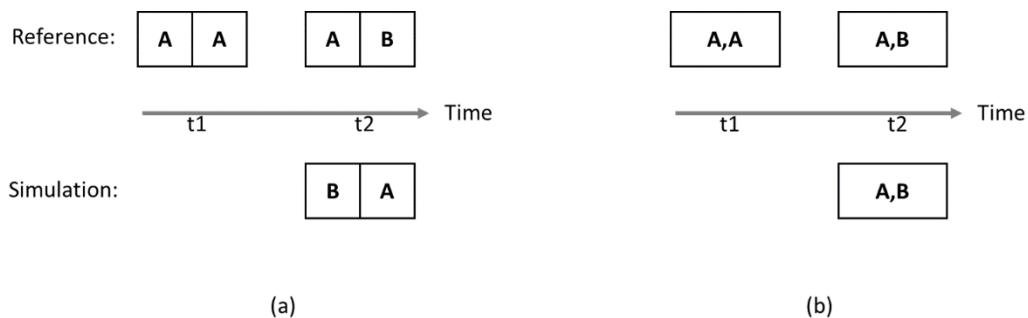


Figure 3.4. A hypothetical example showing reduction of allocation error in coarsening of resolution: (a) two fine-resolution pixels contain a miss and a false alarm; (b) when aggregated in the same coarse pixel, they indicate a hit and a correct rejection. ‘A’ and ‘B’ are states

3.2.4 Four-map comparison and analysis of components of change in partitioned study area

The baseline model divides the study area into two strata: one that is closer to initial points of spread, and another that is further from initial points of spread. For ease of reference, we call them the *near* and *far* strata, respectively. The baseline model predicts the near stratum as changed, and the far stratum as unchanged. These strata are subsets of the study area, such that their union is the entire study area and their intersection is empty. This stratification is important for us because the near stratum is more suitable than the far stratum for future change, and we expect that the simulation shows this matter. To assess the model using baseline, we simultaneously compare the four maps of initial reference, final reference, simulation, and baseline. We call this procedure the four-map comparison.

The four-map comparison is performed by first calculating the difference between the map of reference data at the beginning of the study and each of the three other maps. After this step, there will be three maps of change from beginning to end of the validation time interval. The three maps of change are: reference, simulation, and baseline. Next, we note that each of these three maps is a set of pixels, and we overlap the three sets and note their unions and intersections. Since the scope of our study is limited to models with two states, the combination of these three sets can produce up to 8 subsets. For example, the subset of pixels that indicate change in the reference set, change in the simulation set, and change in the baseline set includes Hits in proximity to the initial points of spread. Other subsets are interpreted similarly. Of the 8 subsets thus produced, 2 of them include no change in the reference set and no change in the simulation set. These two are the subsets of *Correct Rejections*, and we do not include them in the rest of our analysis, as they do not provide any information about correct or incorrect simulations involving change. The remaining 6 subsets indicate respective areas of *Misses*, *Hits*, and *False alarms* in the *near* stratum and in the *far* stratum.

Our validation method involves (1) defining a proximity baseline and (2) comparing the baseline with simulation and references. A summary of steps of these two activities is depicted in Figure 3.5. Note that for applications where scientists choose a different baseline instead of the proximity baseline, the steps for comparison of the baseline with simulation and references are the same as those shown in the figure.

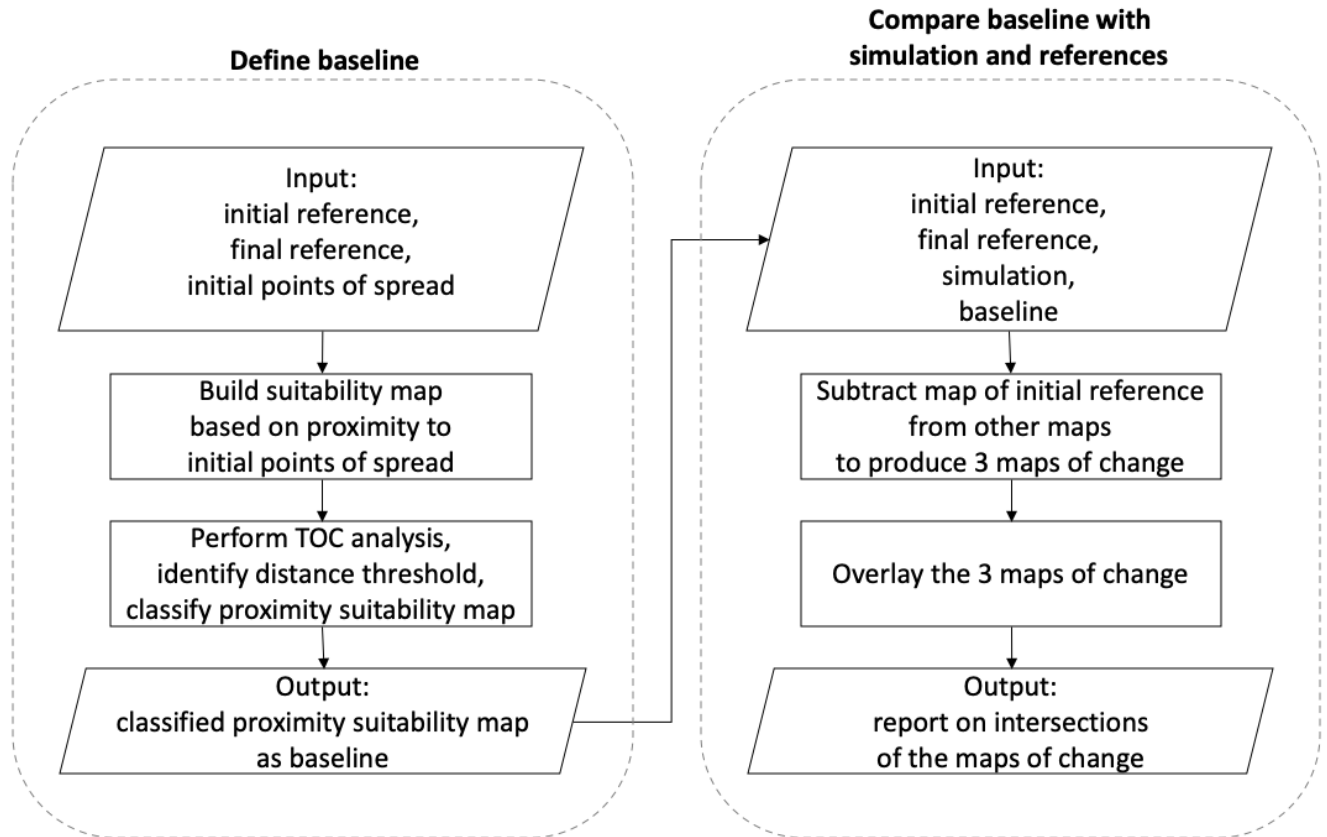


Figure 3.5. Flowchart of the proposed validation method including definition of a proximity baseline and use of the baseline in assessment of simulations.

3.2.5 Data

We assess a MPB model (Harati et al., 2020) and compare two simulations generated by different algorithms, namely, logistic regression (LR) and random forest (RF). The simulations predict the spread of MPB infestations in the forests of BC from 2008 to 2014. Each simulation is analysed in comparison with reference data of the years 2008 and 2014. In addition to the two aforementioned simulations, we define a baseline model and compare its output with reference datasets and LR and RF simulations.

All model outputs and reference datasets are rasters of the study area. Dimensions of the maps are 3516 by 4011 pixels, with pixel resolution of 400 meters. The extent of the study area is from 59°59'27" N 138°54'19" W to 48°59'53" N 114°2'37" W. The study area is the union of forest areas that were not infested at the beginning of 2008. In other words, areas that were already infested before 2008 were excluded from analysis because no change would happen in the state of those areas.

Reference datasets are binary, meaning each pixel of reference data is classified as *infested* or *not infested* at 2008 or 2014. The two MPB model outputs are suitability maps for the region that is not infested at 2008. For two pixels in the same suitability map, the one with the higher value is a more suitable candidate for infestations beyond 2008. The value of a pixel in a suitability map is comparable in terms of rank with values of other pixels in the same map. It is meaningless to compare the value of a pixel in a suitability map with the value of the corresponding pixel in the other suitability map because the suitability values are not probabilities.

The baseline model, which is also associated with a suitability map, is built as described in the Methods section. The input required for this model is the initial points of spread, places of insects in reference data of 2008, which is also a raster with the same dimensions and resolution as the other datasets used in this study.

Reference infestation data was obtained from the web portal of BC Ministry of Forests, Lands, Natural Resource Operations and Rural Development - hereinafter The Ministry (Province of British Columbia, 2015). The data was then converted to binary using a threshold. The Ministry maintains datasets of cumulative MPB attacks based on the Provincial Aerial Overview of Forest Health (Province of British Columbia, 2020). Detection of areas of MPB infestation is based on the fact that patches of trees attacked by the insect, change color in subsequent years.

3.2.6 Software tools

Our analyses were carried out using packages “raster” (Hijmans, 2019), “lulcc” (Moulds, Buytaert, & Mijic, 2015), and “TOC” (Robert G. Pontius Jr, Santacruz, Tayyebi, Parmentier, & Si, 2015) of the statistical software R (R Core Team, 2019). Maps were produced with ArcGIS for Desktop (ESRI, 2015). TOC output plot was produced using TOC Generator software package (Liu, 2020).

3.3 Results

This section presents the results of applying the methods of the paper on the example case of MPB infestations.

Table 3.2 shows the distribution of reference changes in various classes of Manhattan distance from the initial points of spread. The median Manhattan distance to nearest initial point

of spread was 3 pixels. This distance corresponds to the 6-year duration, i.e. from 2008 to 2014. This table includes two noteworthy findings: firstly, in the 6 years of validation time interval 75 % of the new infestations occurred within only 8 pixels (Manhattan distance) from initial points of spread; and secondly, the remaining 25 % of new infestations were dispersed as far as 876 pixels from initial points of spread.

Table 3.2. Percentiles of shortest Manhattan distances to initial points of spread

Manhattan distance (pixels)	Percentile (%)
1	0.0
2	25.0
3	50.0
8	75.0
876	100.0

TOC curves for two simulations as well as the baseline model are shown in Figure 3.6. It can be seen that the baseline model has a higher AUC than the other two models, and the RF model has a higher AUC than the LR model. At about 2/3 of the maximum height of the plot area, the TOC curve of the LR model includes a bend that is sharper than the RF curve. To the right of that bend in LR TOC plot, i.e., with lower suitability thresholds that classify more pixels as infested, changes in threshold cause larger changes in false alarms and smaller changes in hits compared to the other two models. Using these curves, we also selected the classification thresholds for each of the suitability maps. Each of these thresholds corresponds to the point on each curve that is under the upper-left corner of the TOC plot parallelogram. In particular, in our TOC curve analysis, the identified distance threshold for classification of the proximity suitability map was 2 pixels.




Figure 3.6. Total Operating Characteristic curves for two simulations and a proximity model. Hits and false alarms are pixel counts. Areas Under Curve are 0.79 for Logistic Regression, 0.83 for Random Forest, and 0.86 for Proximity.

Figure 3.7 shows the proximity suitability map, which was calculated based on Manhattan distances to initial points of spread. In addition, this figure shows the RF and LR suitability maps, which were the input data of our analysis. The RF and LR maps were produced in a previous study, based on infestation and geographical data (Harati et al., 2020). The maps in this figure are not classified, and each of them indicates some parts of the study area as more suitable for infestations than other parts of the study area. We used the thresholds obtained in TOC analysis to classify the maps of Figure 3.7. The results of these classifications were binary maps in which each data pixel is infested or not infested. Figure 3.8 shows the components of agreement and disagreement with reference data for the classified RF, classified LR, and baseline. The baseline map was produced by classifying the proximity suitability map using a threshold of 2 pixels. Each of the maps of Figure 3.8 is the result of comparison of maps of reference at 2008, reference at 2014, and the respective model at 2014. The figure shows that

most of the observed and simulated change happened in the northern part of the study area. The figure also shows that all models missed changes in the northeast of the study area.

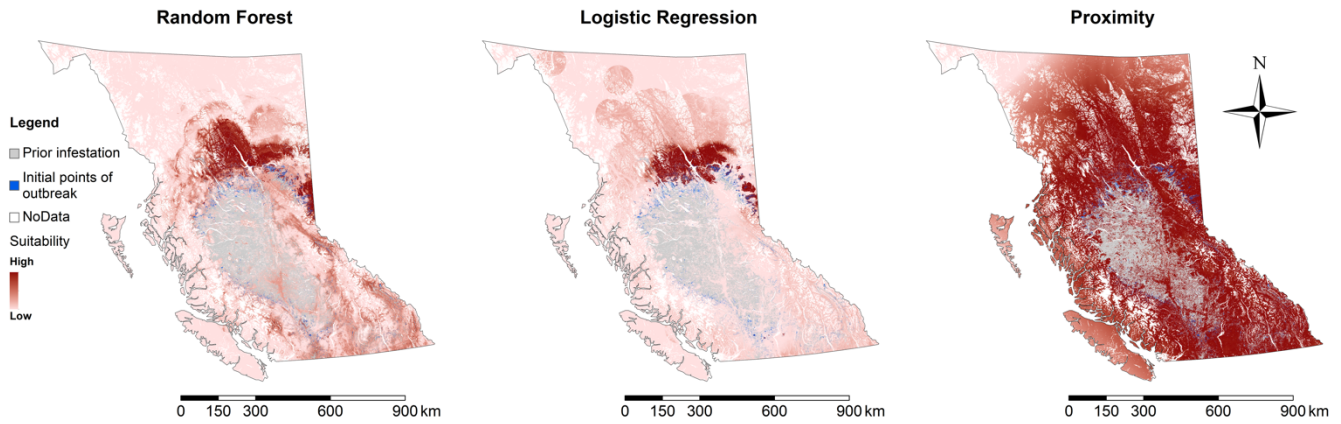


Figure 3.7. Maps of infestation suitability from 2008 to 2014 for two simulations and a proximity model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008

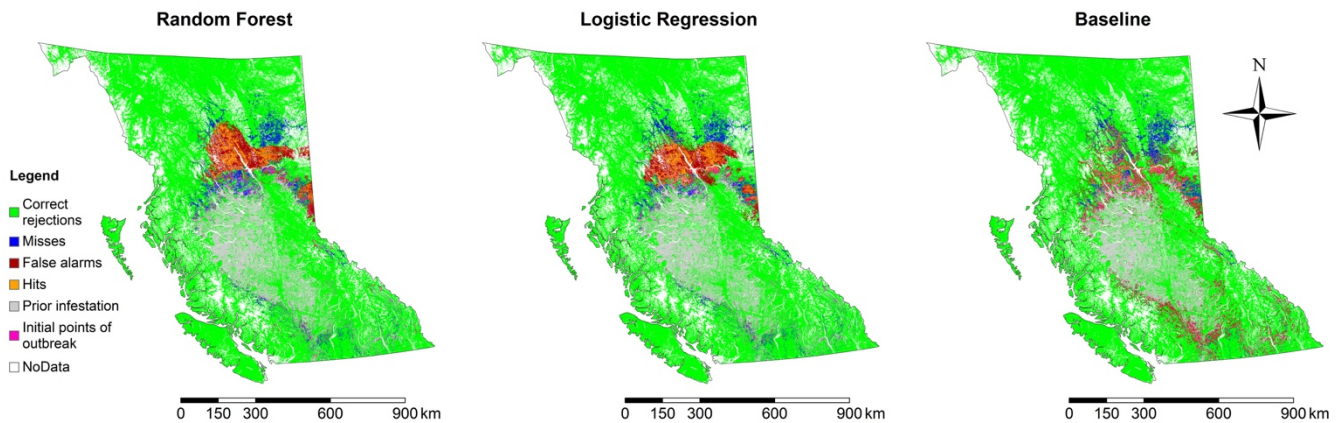


Figure 3.8. Comparison of maps of reference in 2008, reference in 2014, and prediction in 2014 for two simulations and a baseline model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008

Figure 3.9 summarizes the results of three-map comparisons at multiple resolutions for each model. This figure shows area proportions of hits, misses, false alarms, and correct rejections of each simulation at multiple resolutions. The figure shows that the reference infestation is about 6 percent of the overall study area. The figure also shows that RF and LR have zero quantity error, as both misses and false alarms are zero at the coarsest resolution.

Baseline has more false alarms than misses, which indicates that the baseline model has more simulated infestation than reference infestation.

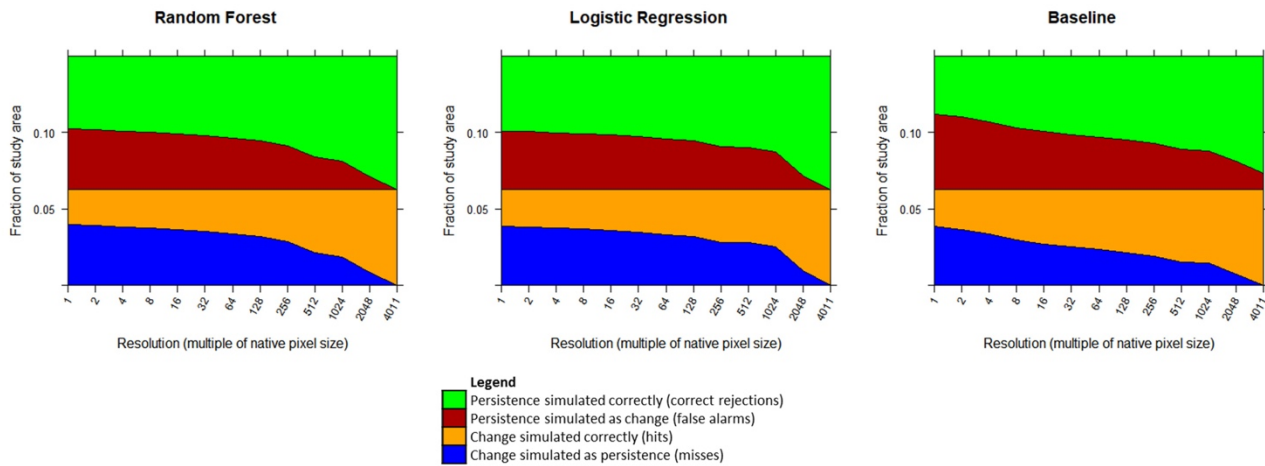


Figure 3.9. Multiple-resolution analysis of components of agreement and disagreement between maps of reference in 2008, reference in 2014, and prediction in 2014

In Figure 3.10 we demonstrate the result of partitioning of the study area using the baseline model. The near and far strata are defined by the threshold obtained in the analysis of the TOC curve of the proximity suitability map, which was 2 pixels. The baseline model predicts the near and far strata as infested and not infested, respectively.




Figure 3.10. Partitioning of the study area into *near* and *far* strata using the baseline model. Initial points of outbreak are where insects were at start (2008). Analyses exclude infestations prior to 2008

Figure 3.11 shows the components of agreement and disagreement for simulated and reference change in the two strata of the study area that are defined based on proximity to initial points of spread. These strata are the predictions of the baseline model. The misses, hits, and false alarms of each model in near stratum are areas that the baseline model predicts as infested. Conversely, misses, hits, and false alarms of each model in the far stratum are areas that the baseline model predicts as not infested. RF has more allocation error than LR. The overall sizes of misses and false alarms of RF are equal. Similarly, the overall sizes of misses and false alarms of LR are equal. In the near stratum of the study area, RF has more misses than false alarms.

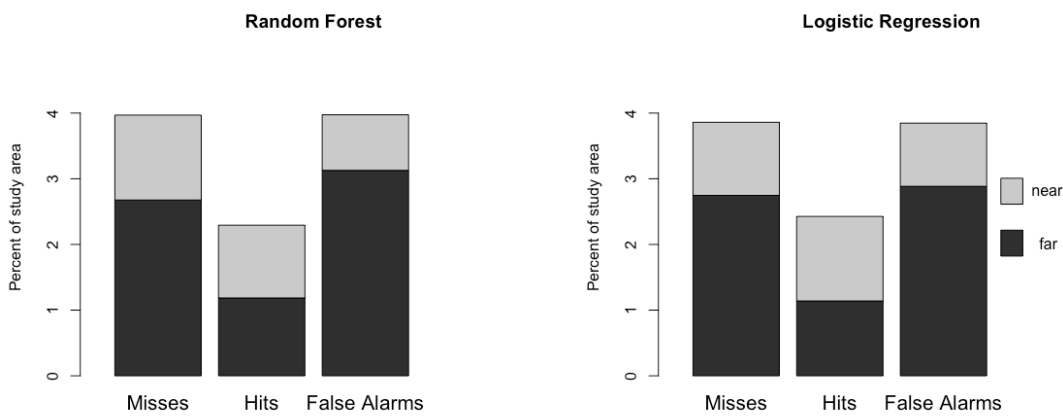


Figure 3.11. Components of change in near and far strata of the study area

3.4 Discussion

This section describes the information and insight from the results for the example case of MPB infestations. These descriptions also include results of existing methods to demonstrate how the new methods complement existing methods by providing information that would otherwise not be attainable. Afterwards, we discuss the limitations and implications of the methods developed in this study.

3.4.1 Insights about the case of study

The results of distance analysis in Table 3.2 reveal interesting information about reference data. The fact that the median distance of spread in 6 years is 3 pixels shows that many new infestations occurred near initial points of spread. This observation suggests it is reasonable to build the baseline in proximity to initial points of spread. On a different note, it is also worth mentioning that the Manhattan distance to initial points of spread ranged from 1 to 876 pixels.

The upper end of this range is much larger than the lower 75 percent of the values, which raises questions about the possible cause of such difference. Future studies can analyse this case and assess if such large distance of MPB spread has been due to, for example, weather and wind conditions, man-made interventions (e.g. MPB piggy-backing on logs transported to processing plants), or data processing errors.

In the analysis of the TOC curves, it is notable that the proximity suitability map has the highest AUC, which indicates better performance compared to the other two models. The curves also show that LR is above RF at the correct quantity, which indicates that LR is better than RF at specifying where few infestations exist. LR curve is below the others in the right side of the TOC space, which indicates that the non-LR models are better at assigning low ranking suitability values to pixels where infestation does not occur.

The distance threshold obtained in the analysis of the TOC curve of the proximity suitability map defines the near and far strata of the study area. In this case, the areas within Manhattan distance of 2 pixels from an initial point of spread are considered the near stratum. Conversely, the areas with Manhattan distance of more than 2 pixels from an initial point of spread are considered the far stratum. In this way, the baseline model is built. The baseline model predicts change in all of the near stratum, and no change in the far stratum. In other words, the baseline model predicts that the near stratum will be completely infested, and the far stratum will not be infested.

Visual inspection of Figure 3.8 shows RF and LR simulations missed many changes in the region in the northeast of the study area. This region is located on the eastern side of the Rocky Mountains. The Rocky Mountains previously served as a natural barrier blocking the spread of MPB towards the east. However, during the validation time interval, large MPB outbreaks occurred in the northeastern part of the province. We see in Figure 3.8 that the RF and LR models did not predict these outbreaks.

Figure 3.9 shows that RF and LR models have no quantity error. This is because in the analysis of their TOC curves, classification thresholds were selected to eliminate quantity error. The figure shows that for the baseline model, the elimination of errors at the coarsest resolution does not happen completely. The baseline model has some excess false alarms, which means that its quantity error was not completely removed. This is because the proximity suitability is

constructed using Manhattan distances from initial points of spread. Manhattan distances are expressed as natural numbers, and there are a large number of pixels having the same distance value. Nevertheless, the selection of the threshold has been set such that the quantification error of the classified result was minimized. That is, the quantification error for the selected threshold is smaller than the quantification error for any other threshold.

In Figure 3.9, as resolutions become coarser, misses and false alarms decrease, and hits and correct rejections increase. A closer look at the LR plot in Figure 3.9 shows that in the coarsening from cell-size 128 to cell-size 256 the curves of errors are steeper than in smaller cell-sizes. This shows that at this particular resolution, there are suddenly more pairs of miss and false alarm errors that cancel one another. In the classified LR simulation, the pairs of miss and false alarm errors that are located between 128 and 256 pixels away from one another are more noticeable than those in shorter distances from one another. This distance interval is an indicator of allocation error of the LR simulation. Following similar steps with the RF simulation, we find that its respective distance interval is between 256 and 512 pixels. This means that, compared with the LR simulation, allocation errors in the RF simulation are further away from one another. In other words, multiple-resolution analysis revealed that the LR model performs better than the RF model in terms of allocation error.

The four-map comparison provides information that the existing methods do not show. Recall that the four-map comparison includes reviewing the results of three-map comparison in the two strata defined by the baseline model (*near* and *far* strata with respect to initial points of spread, as shown in Figure 3.10). This method reveals new information about performance of the models. Figure 3.11 shows the components of change for each model in the study area, as well as their breakdown in the two strata. For each simulation, this figure summarizes useful information obtained from comparing four maps: reference 2008, reference 2014, respective simulation, and the baseline model. The figure provides information on where the misses, hits, and false alarms are in the map: how much of them is in the *near* stratum, and how much in the *far* stratum. Note that the sum of each column in Figure 3.11 corresponds with the fine-resolution information given by Figure 3.9. However, the breakdown of each column into *near* and *far* strata includes new information that cannot be found from Figure 3.9. For one thing, it shows how each of the simulations compares with the baseline model. Since the baseline model predicts no change in

the *far* stratum, it misses all reference changes in the *far* stratum. As such, the *Misses* of simulations in the *far* stratum are *Misses* of the baseline model as well. *Hits* in the *far* stratum show the strength of the LR and RF simulations over the baseline model. On the other hand, LR and RF simulations also produce false alarms in the *far* stratum, which the baseline model does not. Similar arguments can be stated about the *near* stratum. The baseline model predicts change in all of the *near* stratum, therefore it misses nothing there. *Misses* of the simulation in the *near* stratum are all *Hits* of the baseline model, and indicate relative weakness of the simulation with respect to the baseline model. On the other hand, the baseline model produces more false alarms than the simulation in the *near* stratum.

In the example case of MPB infestations, in addition to comparing the simulations with the baseline model, we can gain insight about the models by noting how they compare against one another. Because both LR and RF simulations are classified such that their quantification errors are almost eliminated, for each simulation, misses and false alarms in the entire study area are nearly equal. However, the distribution of these errors in the two strata of the study area is noteworthy. In both simulations, in the near stratum there are more misses than false alarms; and in the far stratum there are fewer misses than false alarms. This difference means that in both simulations, in the near stratum the quantity of change is predicted less than in reference data; and in the far stratum the quantity of change is predicted more than in reference data. In other words, both simulations involve errors of allocating less change than reference to the near stratum, and allocating more change than reference to the far stratum. Moreover, this error is larger in the RF simulation than in the LR simulation. The RF model, in comparison with the LR model, has the weakness of underestimating the spread of infestations near initial points of spread. This useful finding was not evident in previous analyses; rather, it is the result of the additional analysis of components of change in near and far strata.

Our analysis of allocation errors described in this section depends on the strata defined by the baseline model, which, in turn, are the result of distance analysis. As such, for a different result of distance analysis from what we calculated, our assessment of the simulations could be different from what we presented in this section.

3.4.2. Limitations and implications of the study

In this paper we developed methods to assess land change simulation with respect to the baseline that is constructed with common sense. Our motivation in this study was to give an objective answer to the question of where errors occur. A map of errors, of course, shows where errors occur. However, the interpretation of maps is subjective (van Vliet, Bregt, & Hagen-Zanker, 2011). We sought methods to extract objective information from data that was available in initial reference, final reference, simulation, and baseline maps. Our methods have two implications for assessment of land change models: (1) use of near and far strata as defined in this paper for models within the scope of this study; and (2) simultaneous comparison of the four maps of initial reference, final reference, simulation, and baseline, for models that are assessed with a baseline.

This paper demonstrated the application of our methods on an example case of forest insect infestations, but the implications of this study are not limited to this case. The methods of this study can be applied in assessment of a variety of models involving a single transition and spatiotemporal dependency. Nevertheless, the scope of this paper is one of its limits, as there are a wide range of applications that involve multiple land classes and reversible transitions, creating more complicated assessment problems. These problems deserve to be addressed in future works.

An important point to consider about the methods proposed in this paper is that even within the scope of the study, distance from initial points of spread is not the only factor that may be related to errors. In particular, in phenomena with factors that act along certain vectors, it is possible to observe patterns that defy the assumption of spread of change in proximity of previously changed places. Examples of these are the effect of wind on forest fires and forest insect infestations, and the effect of roads on expansion of cities. Fire and insect infestations can spread rapidly along wind vectors, and urban built areas can spread along roads at a faster pace than what a baseline model might suggest. Modelers should be mindful of such effects when interpreting results of model assessment methods. It is important to consider, though, that in such cases if the modelers have a reasonable idea of the other factors that influence the spread of change, they can include that idea in the construction of the baseline model, so that the baseline model agrees with common sense and is still easy to understand. Then they can use the method of four-map comparison to assess their simulation using the new baseline. In all cases, the

rationale of this paper is to use appropriate baseline models in order to gain insights concerning the performance of simulations. If the process of change is known to happen along certain vectors, then the baseline model can be constructed as predicting change along those vectors. If the process of change is known to include a combination of vector effects and proximity effects, then the baseline model can be constructed by superposition of a vector baseline and a proximity baseline.

Another matter worth mentioning is that the use of Manhattan distance involves some deviation with respect to Euclidean distance. The calculation of Manhattan distance has a computational advantage over Euclidean distance in applications where maps have a large number of rows and columns. However, with modern-day hardware and software, calculation of Euclidean distance is feasible for many applications. Modelers should be mindful of this matter in their assessments. In our case, even though we used Manhattan distance to build our proximity suitability map, the respective TOC curve had the highest AUC of all models. This curve was remarkably higher than the random line, which shows that our baseline was much more accurate than a random baseline. The curve was also higher than the RF curve at all threshold points. We built our baseline as an easy-to-understand model that makes sense, especially more relevant than a random model, because we know that infestations do not spread randomly. Use of random baselines has been criticized as irrelevant and/or misleading (Robert Gilmore Pontius Jr & Millones, 2011). The TOC curves indicate that our baseline in this application offered helpful insight. Nevertheless, the building of our baseline involved simplifications, and with simplifications come inaccuracies.

Lastly, the methods presented in this paper have implications for the improvement of methods for calculation of allocation errors. Presently, the method of three-map comparison at multiple resolutions serves this purpose. However, analysis done by the multiple resolution method is dependent on the relative coordinates of pixels with respect to a corner of the map of the study area. This means that the result of multiple resolution analysis of a map can change especially at coarser resolutions, depending on the reference point for defining coordinates of map pixels. Our proposed method in this paper does not depend on a fixed point on the map. Rather, our analysis is based on the phenomenon under study. This can be the basis for development of new methods for assessment of allocation error, while avoiding the said problem

with multiple resolution analysis. Such methods could be developed by expanding our method to partition the study area into multiple strata instead of two.

3.5 Conclusions

Existing methods for evaluation of allocation error provide information on how far allocation errors are with respect to one another, but not with respect to reference data. Our article addressed the topic of allocation of errors in simulations of a subset of land change processes with a single transition and neighborhood effects. For applications in this subset, we presented methods to obtain objective information about where errors occur, by partitioning the study area with a baseline model that is defined through distance analysis of reference data, and performing a four-map comparison including the baseline model, the simulation, the initial reference, and the final reference. These methods identify the distribution of misses, hits, and false alarms in the two strata of the study area, which are defined based on proximity to initial points of spread of the phenomenon. The methods of this paper helped us gain insight concerning the performance of two example simulation cases, and revealed information that would otherwise be unattainable. We recommend that for a one-way, geographically-spreading process, model validation should distinguish between errors that are allocated near and far from the initial points of spread.

Data availability: The datasets generated and/or analysed are available in the Open Science Framework repository, via <https://osf.io/d5em3/>.

Acknowledgements: We are thankful to the Natural Sciences and Engineering Research Council (NSERC) of Canada for partial support of this study under the Discovery Grant Program awarded to LP, and to the Université de Montréal's International Affairs Office (IAO) for their financial support through the International Partnership Development program, which allowed the collaboration between researchers from UdeM and CREAM. RMH received financial support from the NEWFOREST (PIRSES-GA-2013-612645) program of the European Union's Seventh Framework Programme. The United States National Science Foundation supported RGP through its Long Term Ecological Research Network via grant OCE-1637630 for Plum Island Ecosystems.

Competing interests: The authors declare that they have no conflict of interest.

Appendix 3.A. Post-publication remarks

This appendix has been added in the present thesis chapter after the publication of the paper. This appendix includes additional information on the contents of the chapter based on remarks and comments received between publication in journal and final submission of this thesis. Where necessary, references have been cited in this appendix and added in the chapter's bibliography.

One of the key concepts of this work is the distance analysis, which is performed using the proximity suitability map. It is important to note why the proximity suitability map is constructed. This map is used to make a baseline for comparison against simulations. The rationale behind definition of this baseline is that it should represent a naïve model with an easy-to-understand logic. In the case of the present chapter, such naïve model is defined by proximity. That is, our baseline model is built on the easy-to-understand assumption that in a one-way process of change, places adjacent to the points of spread of change are likely to change.

Once the classified baseline map is built, it can be compared with classified simulation maps. The latter, in turn, are constructed by applying a threshold on the simulation suitability maps, which are the simulation outputs. In other words, it is assumed that the models to be assessed produce matrix maps with continuous cell values. The conversion of simulation suitability maps from continuous to classified is done by applying a threshold to each map. Such thresholds are obtained in a TOC analysis.

The aim of this chapter was to introduce the methods of four-map comparison and area partition and show their application in assessment of land change models. In addition to this main theme, the established method of three-map comparison at multiple resolutions was applied on the case of study to demonstrate the information that can and cannot be obtained by this existing method. The description given here for this latter method and the discussion of its results in this sense are of secondary importance in this chapter. However, these very parts formed the beginning of the newer methods that were presented here. Indeed, this work started as my effort to validate the results of the models of Chapter 2 using multiple resolution three-map comparison. In order to understand that method and interpret its results, I built a tool that performed the same operations based on descriptions given in literature (Robert Gilmore Pontius Jr, Boersma, et al., 2008; Robert Gilmore Pontius Jr et al., 2004, 2011; Robert Gilmore Pontius

Jr, Thontteh, & Chen, 2008). As a matter of fact, the first draft of the paper that comprises this chapter was on the use and interpretation of multiple resolution three-map comparison on the MPB simulations. The text that remains of those works in the present chapter therefore summarizes my understanding of this method, and I hope that some readers find it insightful.

Finally, to help better understand the concepts of proximity suitability map, TOC threshold selection, and four-map comparison, the following schematic figures are added:




Figure S3.1. Hypothetical proximity suitability map and baseline. (a) Reference (b) Proximity suitability with Manhattan distances from initial point of spread (c) Baseline. Initial point of spread is marked in red, reference change in yellow, suitability in shades of blue, and baseline with diagonal stripes. Baseline marks the same number of cells as reference change.




Figure S3.2. Critical threshold selection with TOC curve. The vertical line drawn down from the upper left corner of the parallelogram meets the TOC curve at the critical threshold point (shown with a star)

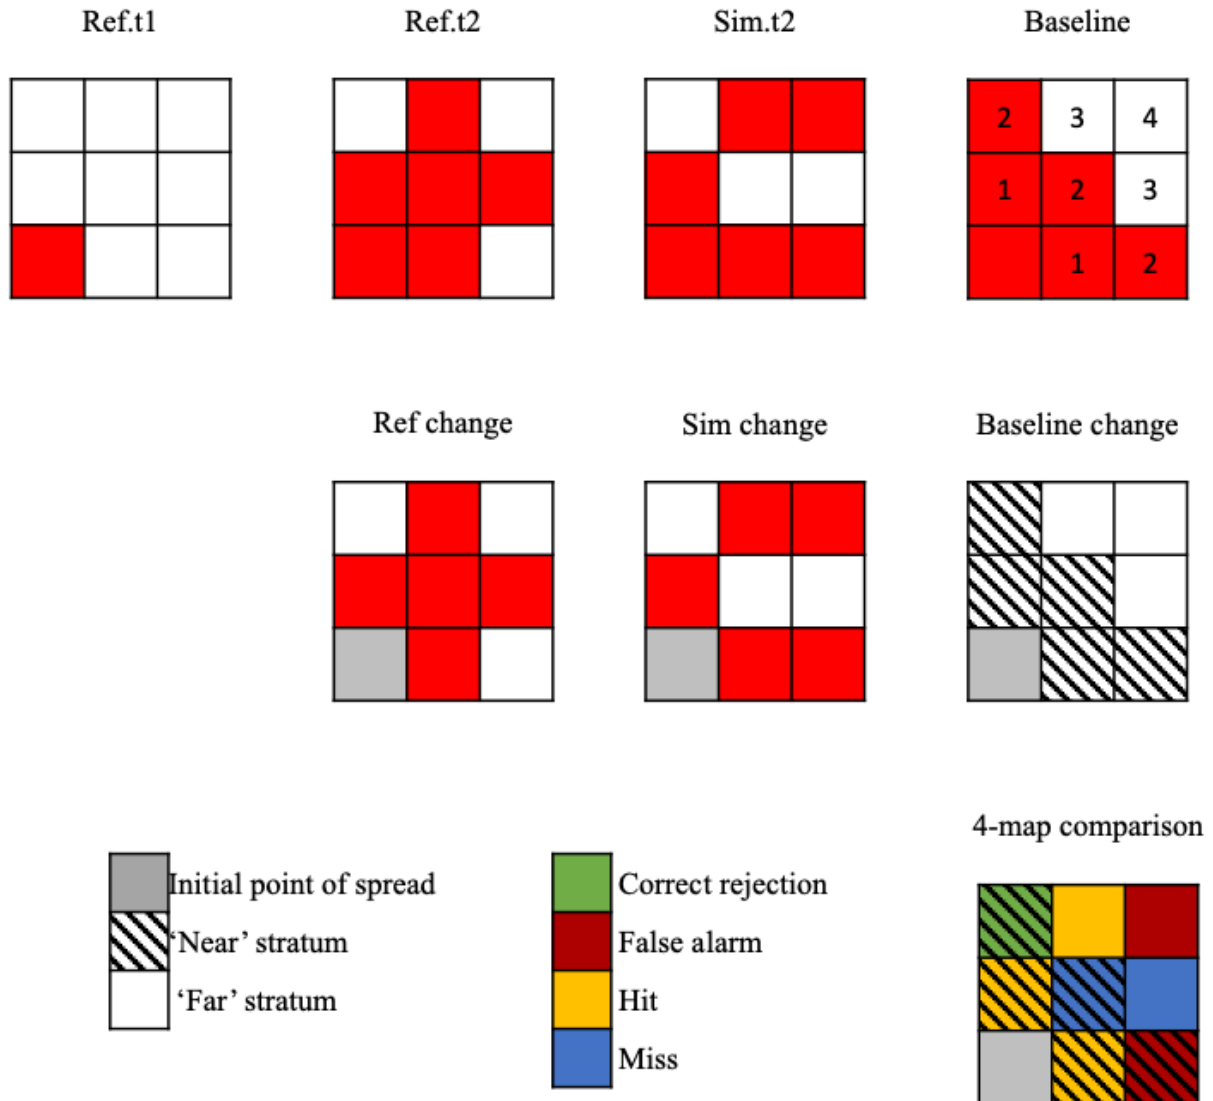


Figure S3.3. Hypothetical four-map comparison. The four maps of the first row are converted to three maps of change on the second row, then overlaid to identify components of agreement and disagreement on the third row.

References

- Batty, M., & Torrens, P. M. (2005). Modelling and prediction in a complex world. *Futures*, 37(7), 745–766. <https://doi.org/10.1016/j.futures.2004.11.003>
- Brown, D. G., Verburg, P. H., Pontius Jr, R. G., & Lange, M. D. (2013). Opportunities to improve impact, integration, and evaluation of land change models. *Current Opinion in Environmental Sustainability*, 5(5), 452–457. <https://doi.org/10.1016/j.cosust.2013.07.012>
- Brown, D. G., Walker, R., Manson, S., & Seto, K. (2012). Modeling Land Use and Land Cover

- Change. In G. Gutman, A. C. Janetos, C. O. Justice, E. F. Moran, J. F. Mustard, R. R. Rindfuss, ... M. A. Cochrane (Eds.), *Land Change Science* (pp. 395–409). Dordrecht: Springer. https://doi.org/10.1007/978-1-4020-2562-4_23
- Chen, H., & Pontius, R. G. (2010). Diagnostic tools to evaluate a spatial land change projection along a gradient of an explanatory variable. *Landscape Ecology*, 25(9), 1319–1331. <https://doi.org/10.1007/s10980-010-9519-5>
- Congalton, R. G. (2004). Putting the map back in map accuracy assessment. In R. S. Lunetta & J. G. Lyon (Eds.), *Remote sensing and GIS accuracy assessment* (pp. 1–11). CRC Press. Retrieved from https://books.google.ca/books?hl=en&lr=&id=KFUV1LL25EIC&oi=fnd&pg=PA1&dq=congalton,+putting+the+map+back&ots=MV6_VWY0mf&sig=3sp7UbLQsI7_IzIWfYjNyJnRrHU
- Cushman, S. A., Macdonald, E. A., Landguth, E. L., Malhi, Y., & Macdonald, D. W. (2017). Multiple-scale prediction of forest loss risk across Borneo. *Landscape Ecology*, 32(8), 1581–1598. <https://doi.org/10.1007/s10980-017-0520-0>
- de Sousa-Neto, E. R., Gomes, L., Nascimento, N., Pacheco, F., & Ometto, J. P. (2018). Land Use and Land Cover Transition in Brazil and Their Effects on Greenhouse Gas Emissions. In *Soil Management and Climate Change* (pp. 309–321). Academic Press. <https://doi.org/10.1016/B978-0-12-812128-3.00020-3>
- Di Gregorio, S., Serra, R., & Villani, M. (1997). A Cellular Automata Model of Soil Bioremediation. *Complex Systems*, 11(1), 31–54. Retrieved from <http://www.complex-systems.com/pdf/11-1-2.pdf>
- ESRI. (2015). ArcGIS 10.4.1 for Desktop.
- Foody, G. M. (2004). Thematic map comparison: evaluating the statistical significance of differences in classification accuracy. *Photogrammetric Engineering & Remote Sensing*, 70(5), 627–633.
- Gaudreau, J., Perez, L., & Drapeau, P. (2016). BorealFireSim: A GIS-based cellular automata model of wildfires for the boreal forest of Quebec in a climate change paradigm. *Ecological Informatics*, 32, 12–27. <https://doi.org/10.1016/j.ecoinf.2015.12.006>
- Hagen-Zanker, A. (2006). Map comparison methods that simultaneously address overlap and structure. *Journal of Geographical Systems*, 8(2), 165–185. <https://doi.org/10.1007/s10109->

006-0024-y

- Harati, S., Perez, L., & Molowny-Horas, R. (2020). Integrating neighborhood effect and supervised machine learning techniques to model and simulate forest insect outbreaks in british columbia, canada. *Forests*, *11*(11), 1–23. <https://doi.org/10.3390/f11111215>
- Hermoso, V., Morán-Ordóñez, A., & Brotons, L. (2018). Assessing the role of Natura 2000 at maintaining dynamic landscapes in Europe over the last two decades: implications for conservation. *Landscape Ecology*, *33*(8), 1447–1460. <https://doi.org/10.1007/s10980-018-0683-3>
- Hijmans, R. J. (2019). raster: Geographic Data Analysis and Modeling. R package version 2.9-5. Retrieved from <https://cran.r-project.org/package=raster>
- Lambin, E. F., Geist, H., & Rindfuss, R. R. (2006). Introduction: Local Processes with Global Impacts. In E. F. Lambin & H. Geist (Eds.), *Land-Use and Land-Cover Change* (pp. 1–8). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-32202-7_1
- Li, Z., Huffman, T., Zhang, A., Zhou, F., & McConkey, B. (2012). Spatially locating soil classes within complex soil polygons – Mapping soil capability for agriculture in Saskatchewan Canada. *Agriculture, Ecosystems & Environment*, *152*, 59–67. <https://doi.org/10.1016/j.agee.2012.02.007>
- Liu, Z. (2020). TOC Curve Generator. Retrieved from <https://lazygis.github.io/projects/TOCCurveGenerator>
- Moulds, S., Buytaert, W., & Mijic, A. (2015). An open and extensible framework for spatially explicit land use change modelling: The lulcc R package. *Geoscientific Model Development*, *8*(10), 3215–3229. <https://doi.org/10.5194/gmd-8-3215-2015>
- National Research Council. (2014). *Advancing Land Change Modeling: Opportunities and Research Requirements*. Washington, D.C.: National Academies Press. <https://doi.org/10.17226/18385>
- Natural Resources Canada. (2019). Mountain pine beetle. Retrieved June 2, 2020, from <https://www.nrcan.gc.ca/our-natural-resources/forests-forestry/wildland-fires-insects-disturban/top-forest-insects-diseases-cana/mountain-pine-beetle/13381>
- Paudel, S., & Yuan, F. (2012). Assessing landscape changes and dynamics using patch analysis and GIS modeling. *International Journal of Applied Earth Observation and Geoinformation*, *16*, 66–76. <https://doi.org/10.1016/j.jag.2011.12.003>

- Pérez, L., Dragičević, S., & White, R. (2013). Model testing and assessment: Perspectives from a swarm intelligence, agent-based model of forest insect infestations. *Computers, Environment and Urban Systems*, 39, 121–135.
<https://doi.org/10.1016/j.compenvurbsys.2012.10.004>
- Perez, L., Molowny-Horas, R., & Harati, S. (2016). Modelling Forest Insect Outbreaks: Efforts Towards an Inverse Approach to Model Calibration. In S. Sauvage, J. M. Sánchez-Pérez, & A. E. Rizzoli (Eds.), *Proceedings of the 8th International Congress on Environmental Modelling and Software (iEMSs)* (p. 688). Toulouse, France. Retrieved from
<https://scholarsarchive.byu.edu/iemssconference/2016/Stream-C/54/>
- Pijanowski, B. C., Pithadia, S., Shellito, B. A., & Alexandridis, K. (2005). Calibrating a neural network-based urban change model for two metropolitan areas of the Upper Midwest of the United States. *International Journal of Geographical Information Science*, 19(2), 197–215.
<https://doi.org/10.1080/13658810410001713416>
- Pontius Jr, Robert G., Santacruz, A., Tayyebi, A., Parmentier, B., & Si, K. (2015). TOC: Total Operating Characteristic Curve and ROC Curve. R package version 0.0-4. Retrieved from
<https://cran.r-project.org/package=TOC>
- Pontius Jr, Robert Gilmore. (2000). Quantification error versus location error in comparison of categorical maps. *Photogrammetric Engineering & Remote Sensing*, 66(8), 1011–1016.
- Pontius Jr, Robert Gilmore. (2002). Statistical methods to partition effects of quantity and location during comparison of categorical maps at multiple resolutions. *Photogrammetric Engineering and Remote Sensing*, 68(10), 1041–1050.
- Pontius Jr, Robert Gilmore. (2018). Criteria to confirm models that simulate deforestation and carbon disturbance. *Land*, 7(3), 1–14. <https://doi.org/10.3390/land7030105>
- Pontius Jr, Robert Gilmore, Boersma, W., Castella, J. C., Clarke, K., Nijs, T., Dietzel, C., ... Verburg, P. H. (2008). Comparing the input, output, and validation maps for several models of land change. *Annals of Regional Science*, 42(1), 11–37. <https://doi.org/10.1007/s00168-007-0138-2>
- Pontius Jr, Robert Gilmore, Castella, J.-C., de Nijs, T., Duan, Z., Fotsing, E., Goldstein, N., ... Veldkamp, A. T. (2018). Lessons and Challenges in Land Change Modeling Derived from Synthesis of Cross-Case Comparisons. In M. Behnisch & G. Meinel (Eds.), *Possible Urban Futures: The Impact of Planners and Developers on Urban Dynamics* (Vol. 19, pp. 143–

- 164). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-52522-8_8
- Pontius Jr, Robert Gilmore, Huffaker, D., & Denman, K. (2004). Useful techniques of validation for spatially explicit land-change models. *Ecological Modelling*, *179*(4), 445–461. <https://doi.org/10.1016/j.ecolmodel.2004.05.010>
- Pontius Jr, Robert Gilmore, & Millones, M. (2011). Death to Kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *International Journal of Remote Sensing*, *32*(15), 4407–4429. <https://doi.org/10.1080/01431161.2011.552923>
- Pontius Jr, Robert Gilmore, & Parmentier, B. (2014). Recommendations for using the relative operating characteristic (ROC). *Landscape Ecology*, *29*(3), 367–382. <https://doi.org/10.1007/s10980-013-9984-8>
- Pontius Jr, Robert Gilmore, Peethambaram, S., & Castella, J. C. (2011). Comparison of three maps at multiple resolutions: A case study of land change simulation in cho don district, Vietnam. *Annals of the Association of American Geographers*, *101*(1), 45–62. <https://doi.org/10.1080/00045608.2010.517742>
- Pontius Jr, Robert Gilmore, & Si, K. (2014). The total operating characteristic to measure diagnostic ability for multiple thresholds. *International Journal of Geographical Information Science*, *28*(3), 570–583. <https://doi.org/10.1080/13658816.2013.862623>
- Pontius Jr, Robert Gilmore, Thontteh, O., & Chen, H. (2008). Components of information for multiple resolution comparison between maps that share a real variable. *Environmental and Ecological Statistics*, *15*(2), 111–142. <https://doi.org/10.1007/s10651-007-0043-y>
- Pontius Jr, Robert Gilmore, Walker, R., Yao-kumah, R., Arima, E., Aldrich, S., Caldas, M., & Vergara, D. (2007). Accuracy assessment for a simulation model of Amazonian deforestation. *Annals of the Association of American Geographers*, *97*(4), 677–695. <https://doi.org/10.1111/j.1467-8306.2007.00577.x>
- Province of British Columbia. (2015). [dataset] BC MPB Observed Cumulative Kill - vol.12. Retrieved from <https://www.for.gov.bc.ca/ftp/hre/external/!publish/web/bcmpb/Year12/BCMPB.v12.CumKill.Observed.zip>
- Province of British Columbia. (2020). Aerial Overview Survey Methods. Retrieved June 2, 2020, from <https://www2.gov.bc.ca/gov/content/industry/forestry/managing-our-forest->

resources/forest-health/aerial-overview-surveys/methods

- R Core Team. (2019). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/>
- Rollins, M. G., Keane, R. E., & Parsons, R. A. (2004). MAPPING FUELS AND FIRE REGIMES USING REMOTE SENSING, ECOSYSTEM SIMULATION, AND GRADIENT MODELING. *Ecological Applications*, *14*(1), 75–95. <https://doi.org/10.1890/02-5145>
- Rykiel, E. J. (1996). Testing ecological models: The meaning of validation. *Ecological Modelling*, *90*(3), 229–244. [https://doi.org/10.1016/0304-3800\(95\)00152-2](https://doi.org/10.1016/0304-3800(95)00152-2)
- Tobler, W. R. (1970). A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic Geography*, *46*, 234. <https://doi.org/10.2307/143141>
- van Vliet, J., Bregt, A. K., Brown, D. G., van Delden, H., Heckbert, S., & Verburg, P. H. (2016). A review of current calibration and validation practices in land-change modeling. *Environmental Modelling & Software*, *82*, 174–182. <https://doi.org/10.1016/j.envsoft.2016.04.017>
- van Vliet, J., Bregt, A. K., & Hagen-Zanker, A. (2011). Revisiting Kappa to account for change in the accuracy assessment of land-use change models. *Ecological Modelling*, *222*(8), 1367–1375. <https://doi.org/10.1016/j.ecolmodel.2011.01.017>
- Verburg, P. H., Kok, K., Pontius Jr, R. G., & Veldkamp, A. (2006). Modeling Land-Use and Land-Cover Change. In E.F. Lambin & H. Geist (Eds.), *Land-Use and Land-Cover Change* (pp. 117–135). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-32202-7_5
- White, R. (2006). Pattern based map comparisons. *Journal of Geographical Systems*, *8*(2), 145–164. <https://doi.org/10.1007/s10109-006-0026-9>

Chapter 4

Presentation of the article

Previous chapters described a social model and an ecological model, which were independently developed. In the work described in Chapter 4 I coupled these models together and built a social-ecological model. I then used this coupled model as a virtual laboratory in which various governance scenarios were tested in order to gain insight about the dynamics of the studied social ecological system and the changes that may arise in it upon intervention. In this way, Chapter 4 brings the previous parts of this doctoral work together to address the questions of this thesis. This chapter presents an approach to the analysis of a class of governance problems in the context of sustainability, and it clarifies the said approach with an example application. This chapter demonstrates that it *is* possible to promote environmentally responsible behavior and protect a vulnerable resource against disturbance. In so doing, this chapter also showcases how we can gain and improve an insight about the world through hypothetical experiments with abstract models.

The text of this chapter has been prepared for submission to peer-reviewed journals. My coauthors in preparation of this manuscript are my supervisors, Dr. Liliana Perez and Dr. Roberto Molowny-Horas.

Emergence of environmentally responsible behavior through a recognition mechanism – insights from a conceptual social-ecological model

Saeed Harati^{1*}, Liliana Perez¹, Roberto Molowny-Horas²

¹ Laboratory of Environmental Geosimulation (LEDGE), Department of Geography, Université de Montréal, 1375 Ave. Thérèse-Lavoie-Roux, Montreal, QC H2V 0B3, Canada; saeed.harati.asl@umontreal.ca , l.perez@umontreal.ca

² Centre de Recerca Ecològica i Aplicacions Forestals (CREAF), Bellaterra, E-08193 Cerdanyola de Vallès, Catalonia, Spain; roberto@creaf.uab.es

* Corresponding author

Abstract

One of the challenges of governance for sustainable development is to engage users of an ecological resource in its protection. This task is especially difficult when the users act upon self interest, when the government cannot use force or financial incentives, and when the government has little knowledge about the ecosystem's response to intervention. These constraints raise questions about the possibility of promoting environmentally responsible behavior among users and protecting the ecological resource. Focusing on a case of forest insect infestations, we addressed these questions with a conceptual model. We coupled a land change model of forest insect infestations with a social model in which a governing agent applies a recognition mechanism to encourage several user agents to behave in an environmentally responsible way. In the recognition mechanism, the governing agent selects the criteria for acknowledging responsible user agents. We ran the coupled model using various scenarios including best-case, worst-case, and random baselines as well as a scenario with a Reinforcement Learning algorithm for the governing agent. In simulations with a synchronous start of social and ecological models, the ecological resource was deteriorated before environmentally responsible behavior emerged. However, with independent runs of the recognition mechanism in the social model prior to coupling with the ecological model, the emergence of responsible behavior was hastened in the

coupled model and the resource was saved. We noted superior performance of the Reinforcement Learning algorithm in comparison with the random baseline and within the range of results identified by the best and worst case scenarios. Thus, with a conceptual model and through scenario analysis, we gained insight about the possibility of promoting environmentally responsible behavior without using force or financial incentives.

Keywords: Social-ecological system, Agent Based Model, Reinforcement Learning, scenario, conceptual model, governance, environmentally responsible behavior, emergence, complex systems

4.1. Introduction

4.1.1. A governance problem in sustainable development

Sustainable development is defined as a development that provides the needs of the present time, without sacrificing the ability of future generations to provide their needs (Brundtland, 1987) (Brundtland, 1987). The criteria for such development are known to have social, ecological, and economic dimensions (Barbier, 1987; Brown, Hanson, Liverman, & Merideth, 1987). In the present study we are interested in a group of sustainable development problems wherein a government seeks help from users of a natural resource, which is at risk, to protect that resource. From the government's viewpoint it is ideal that the users find the motivation to cooperate with the government, but such motivation may not come out of the users' sense of altruism (Kaplan, 2000). One possible way to build that sense of cooperation may be to offer a financial incentive, but that may not be enough to create a long-lasting motivation for environmentally responsible behavior, either (de Young, 2000; Katzev & Johnson, 1987). Moreover, experience shows that using authority to enforce environmentally responsible behavior may fail (Blundell & Gullison, 2003; Feeny et al., 1990; Wagner, 2004; Wittemyer et al., 2011). Therefore, in the ideal situation in the government's view, users should voluntarily adopt a behavior that implies costs to them, without the governing entity needing to use force or financial incentives. In addition, in that ideal situation, the said behavior successfully protects the natural resource. These considerations give rise to a question: is such an ideal situation possible?

In this paper, our interest is in a particular type of the above-mentioned sustainability problems, where the ecological system is a forest resource that is attacked by a pest, and the

social system includes users of the resource and a governing entity. We assume a case where the users are logging companies and the government wants them to cut specific parts of the forest in order to control the spread of the pest. This problem is multi-disciplinary: it involves land change science (Lambin et al., 2006) in the study of changes in the ecological system; it relates to the domains of collective action (Nyborg et al., 2016; Ostrom, 2009) and social norms (Farrow et al., 2017; Nyborg et al., 2016) in the analysis of formation of a common behavior among users; it entails the field of social-ecological systems (Liu et al., 2007; Ostrom, 2009) in the endeavor to understand the dynamics that emerge through coupling the society with the forest; and, in a broader view, this problem and the question of how to address it are in the realm of complex systems (Cosens et al., 2021; Filotas et al., 2014).

4.1.2. Background from multiple disciplines

Complex systems are entities composed of elements and interactions that make the system behave as a whole, with such characteristics as self organization, non-linearity, emergence, feedback, and path-dependence (O’Sullivan, 2004). Because of these characteristics, the dynamics of complex systems involve novelty and surprise (Michael Batty & Torrens, 2005). This causes a concern in problems of sustainable development, as they typically involve intervention in or experimentation with complex systems, and particularly social-ecological systems. Due to the uncertainty and complexity that is inherent in these systems, it is not always ethically and logistically justifiable to perform trial-and-error experiments on them (Kriebel et al., 2001), as intervention in these systems may have unanticipated, irreversible and adverse effects. This concern justifies learning by modeling and simulation (M. A. Janssen & Ostrom, 2006).

Societies and ecosystems are complex systems. When a society uses a natural resource, the link between the two systems creates a larger complex system, referred to as a social-ecological system (Berkes & Folke, 2000). Social-ecological systems (SES) demonstrate complexities that cannot be understood through the lens of sociology or ecology alone (Liu et al., 2007). In a typical SES, the society receives ecosystem services (Daily, 2000; Millennium Ecosystem Assessment, 2003) and makes changes in the ecosystem. To combine social and ecological knowledge in the analysis of the complexity of SES, a framework has been developed, which accounts for governance and resource systems at larger scales, as well as users and resource units at smaller scales (Ostrom, 2009). This framework has been implicit in

sustainability studies in a variety of domains such as sustainable navigation (Parrott et al., 2011), fishery (Schlüter, Hinkel, Bots, & Arlinghaus, 2014), and forest management (Wimolsakcharoen, Dumrongrojwatthana, Le Page, Bousquet, & Trébuil, 2021).

Many studies of complex systems involve building and using models that replicate some aspects of those systems (Railsback & Grimm, 2012; Wolfram, 2002). Simulations of complex systems are often built with a bottom-up approach, using methods such as Agent-Based Models (ABM) and Cellular Automata (CA) models (Grimm et al., 2005). An ABM is made up of several computer programmed agents that interact with each other and their environment, and act upon their decision rules (Castle & Crooks, 2006). A CA model is composed of a grid of cells, where the state of each cell is defined by a rule based on the previous states of that cell and its neighbors (White & Engelen, 1993). ABMs have been used in a wide variety of complex systems studies, such as epidemiology (Perez & Dragicevic, 2009), animal movement (Bonnell et al., 2013), land development (Pooyandeh & Marceau, 2013) and forest disturbance (Katan & Perez, 2021; Perez & Dragicevic, 2010). Likewise, CA have been used in research works within fields such as land change (Lambin & Geist, 2006; National Research Council, 2014), urban growth (M. Batty, Xie, & Sun, 1999; Clarke, Hoppen, & Gaydos, 1997; de Almeida et al., 2003) and forest disturbance (Bone, Dragicevic, & Roberts, 2006; Gaudreau, Perez, & Drapeau, 2016), among others.

Emergence of behavior in a society involves the field of social norms. Several definitions of norms have been stated in social sciences literature. In one definition, norms are cultural rules that guide people in their behavior (Ross, 1973). In another definition, norms are social rules that govern the encouragement or condemnation of certain behaviors (Savarimuthu & Cranefield, 2011). Norms have also been defined in the context of institutions (Crawford & Ostrom, 1995; Ostrom, 1990) as valuations of actions regardless of the immediate consequences of the actions. Institutions can formalize norms by converting them to regulations (D. C. North, 1990). In yet another view, norms are classified as descriptive and injunctive. Descriptive norms show what others do, whereas injunctive norms show what others approve of (Cialdini, Reno, & Kallgren, 1990). A review of literature on SES governance indicates that social norms largely influence environmentally responsible behavior (Bourceret, Amblard, & Mathias, 2021). Literature also highlights that emergence of environmentally responsible behavior in a society depends on what

the individuals do and what they favor (Nyborg et al., 2016), which, by the above definitions, are the equivalents of descriptive and injunctive norms, respectively.

4.1.3. Setting, questions and objectives

In this study we are interested in a SES governance problem. We consider a setting where the government needs the participation of individual users in a management action with the aim of protecting a natural resource. In this sense, the government desires to promote environmentally responsible behavior among the individual users. To clarify the scope of the problem we state the following assumptions:

- The resource is at risk, and the state of the resource urges the government to act towards its protection.
- The expected behavior to protect the resource is costly to the individuals.
- The individuals are driven by self-interest, and not by altruism.
- The government cannot offer financial incentives for the purpose of enticing the cooperation of the individuals.
- The government cannot enforce its authority and oblige the individuals to cooperate with it.
- Although individuals have an interest in good reputation, there is no social sanction or punishment for individuals who do not demonstrate responsible behavior.
- The government's knowledge of the social system is limited. It does not know the decision criteria of the individuals.
- The government's knowledge of the ecosystem is limited. Although the government wishes to intervene in the ecosystem to protect the resource, the government is not certain about the consequences of its desired intervention.

The above assumptions represent some of the challenges of SES governance. The final statement refers to a particular complexity that is neither entirely social nor entirely ecological. Enticing the cooperation of individual users with the government in the management action is

already a social challenge. However, even if this challenge is overcome and the users fully cooperate with the government, it is not clear how the ecosystem will respond to the management action. Moreover, there are many possibilities involving partial cooperation of users with the government. These give rise to important questions about the ecosystem's response to these cases of partial action, and subsequently, about the whole SES as the government continuously tries to protect the changing ecosystem.

Considering the above-mentioned complexities and the precautionary principle (Kriebel et al., 2001) a modelling approach to these problems makes sense, because modelling allows to gain insight into a complex system without actually modifying it and risking unanticipated consequences. Indeed, SES governance literature highlights the use of conceptual models for gaining insight about complex SES problems (Marco A. Janssen & Ostrom, 2006). Therefore, in this study we model and simulate a SES. The ecological case that our model simulates is the infestation of the west Canadian forests of British Columbia (BC) by the Mountain Pine Beetle (MPB). We have chosen this setting because this study is the continuation of a series of works that we have done on MPB infestation in BC, particularly including the development and validation of a land change model that simulates the spread of MPB infestations (Harati, Perez, & Molowny-Horas, 2020; Harati, Perez, Molowny-Horas, & Pontius, 2021). We define our SES model by coupling the land change model with a conceptual social model that simulates the emergence of new behavior among a group of user agents. The social model has also been developed in a previous work (Harati, Perez, & Molowny-Horas, 2021). Our SES model is therefore comprised of a dynamic landscape model, several user agents, and a governing agent with the intention of encouraging the user agents to perform a certain behavior. In our study, the intended behavior is cutting trees within an established neighborhood around all newly observed infestations.

Based on assumptions of importance of personal motivations such as good reputation in voluntary action (Omoto & Snyder, 1995; Stern, Dietz, & Kalof, 1993) and importance of visibility of one's actions in one's behavioral choices in society (Mosler, 1993; Nyborg et al., 2016), we define a hypothetical scheme in which the government offers recognition to users who cooperate with it. This offer of recognition might motivate the users to perform the requested

action. Users assess the utility of such recognition by considering the uniqueness that it gives them, and the visibility of their recognition in their group.

Each year, the government requests the users to participate in some action. The government selects the action between two choices: (1) the costly management action that helps the government in the control of forest infestations, and (2) a no-cost action that is useless for the government. The reward for users who accept the government's request is a '*responsible user*' label, which is a recognition that may improve the social reputation of the users but has no financial or legal value otherwise. If the government requests the costly action and some users cooperate with it, some parts of the forest will be cut. This will influence the remaining healthy forest and the spread of MPB infestations in the next year.

In a previous work, we built a conceptual social model with the recognition scheme described above, without ecological complexities, and we showed that decision algorithms exist for the government to make the '*responsible user*' label valuable (Harati, Perez, & Molowny-Horas, 2021). Now, the addition of the ecological component to the model raises new questions. Specifically, is the government able to entice the cooperation of the users in the new setting? With the optimistic assumption that users always cooperate with the government, can the government's management action control the infestations and save the forest? Without the above-said optimistic assumption, can the government's management action control the infestations and save the forest? These questions form our endeavor in the present study. Consequently, our objectives in this paper are:

- To build a SES model by coupling the said social and ecological models;
- To perform hypothetical experiments by implementing management scenarios in simulations of the SES;
- To interpret the outcome of the hypothetical experiments. Subsequently, to gain insight about (1) the above-said recognition scheme and its potential for promotion of environmentally responsible behavior, and (2) the state of health of the ecological resource in response to the social dynamics that emerge from the recognition scheme.

4.2. Methods

In order to answer the questions introduced in the previous section, we take a modelling approach that couples together a conceptual social model with an ecological model, hence building a SES model. We have developed those two models independently in previous studies: the social model simulates the emergence of a new behavior in a society through a mechanism of encouragement and recognition of responsible individuals (Harati, Perez, & Molowny-Horas, 2021), and the ecological model simulates the spread of an insect infestation in a forest (Harati et al., 2020). The two models are coupled such that the decisions made in the social model influence, and are influenced by, changes in the ecological model. Then, we use this coupled model as a virtual laboratory to perform hypothetical tests to simulate different scenarios, including various baselines. These include: 3 scenarios with calculated government decisions; 3 scenarios with random government decisions; an enforced action scenario; and a business as usual scenario. Among them, the random, the enforcement, and the business as usual scenarios are used as baselines for comparison with the results of the calculated government decisions. Comparison of the scenarios involves analyzing their results in terms of cooperation of the users with the government, as well as the remaining healthy forest. With the help of baselines we learn about the range of possible outcomes of our hypothetical tests, and through comparison of baselines with management scenario simulations we gain insights about the complexities of the SES of our study.

In the description of our model we follow the ODD+D protocol, which is an extension of the ODD protocol particularly adapted for describing human decisions in ABMs (Müller et al., 2013). The ODD (Overview, Design concepts, Details) protocol is a standard for communication of information about ABMs (Grimm et al., 2006, 2010, 2020). Sections 2.1 to 2.3 below describe the model according to the ODD+D protocol. Section 2.4 describes how we use the model in the present study.

4.2.1. Overview

4.2.1.1. Purpose

The overall goal of this study is to gain insight about the possibility of emergence of environmentally responsible behavior in a SES in the absence of altruism, obligation, and financial incentives. To that end, this model has been built with the purpose of simulating a mechanism of recognition of environmentally responsible behavior in a setting of forest

disturbance. We use the model to learn about the complexities of a SES in which users of a forest resource are requested to participate in a costly action to protect the resource from a disturbance. Particularly, we intend to better understand the potential of the desire for good reputation in promotion of environmentally responsible behavior, and its implications in management and policy making. The model has been designed for scientists, policy and decision makers, and experts with an interest in developing decision support systems.

4.2.1.2. *Entities, state variables and scales*

The model consists of a social component, which is an ABM, and an ecological component, which is a spatial ecological model. The social ABM includes a governing agent, several user agents and an auxiliary agent called registrar, which has been defined for a better understanding of the model. The model includes a spatial component, which represents a forest resource attacked by insect infestations. The spatial ecological model's units are grid cells.

In each time-step the governing agent is in a state, takes an action, and receives a reward. The governing agent's state is a 2-dimensional variable which cumulatively summarizes past interventions and their results. The two components of this state variable are calculated based on actions and rewards. The governing agent's reward is the cooperation by the user agents in the management action to save the forest. The governing agent's action is to generate a binary *signal* for communication to user agents. A signal of 0 is a request for a no-cost action, and a signal of 1 is a request for a costly action. To produce *signal*, it uses a *policy*, which recommends an action for each state. To update its *policy*, it uses its memory of past states, actions and rewards. For that purpose, the governing agent has two arrays of time-discounted scores calculated for each *(state, action)* pair (Harati, Perez, & Molowny-Horas, 2021). Each user agent is allocated a forest zone where the user agent harvests. These zones are created by dividing the map of the study area into equal squares. Each square contains cells that are in the study area and cells that are not. Therefore, user zones include various numbers of cells from the study area. Moreover, each user agent is characterized by a constant decision *threshold* that it uses in a cost-benefit analysis to make a binary *decision* in response to the governing agent's *signal*. Decisions of 0 and 1 mean rejection and acceptance of the governing agent's request, respectively. The registrar is characterized by two variables, *nSum* and *nLast*, which are non-negative integers. *nSum* is the total number of times user agents decided to cooperate with the governing agent. *nLast* is the number of user agents who cooperated with the governing agent in

the last time the latter requested a costly action. In the spatial component of the model, each cell is identified by geographical data fields including coordinates, elevation, aspect, slope and ruggedness. Cells are marked with presence or absence of infestations. Cells are also marked with a mask layer that indicates presence or absence of trees. The simulation area is divided into zones and each zone is allocated to a user agent. Table 4.1 lists the model entities and their state variables.

Table 4.1. Model entities and their state variables

Entity	State variable	Description
Registrar	<i>nLast</i>	Last known number of user agents cooperating with the governing agent
	<i>nSum</i>	Cumulative number of user agents cooperating with the governing agent
Governing agent	<i>state</i>	2-dimensional summary of past interventions and results
	<i>signal</i>	Binary request (easy task or hard task)
User agent	<i>decision</i>	Binary response (refuse or accept)
Cell	<i>infestation</i>	Binary indicator of presence or absence of infestations
	<i>mask</i>	Binary indicator of presence or absence of host trees

The social ABM is influenced by an exogenous driver: spread of insect infestations in the forest. This driver is simulated in the spatial ecological model, which is coupled with the ABM. The spatial ecological model is a GIS model that simulates land change. The land change process of this study is the infestation of forests of British Columbia (BC) in western Canada by the Mountain Pine Beetle (MPB). The area where infestations are modeled is a sub-division of the Kamloops Timber Supply Area (TSA), as shown in Figure 4.1. The extents of this area are from 120°19'59"W 50°45'22"N to 119°6'0"W 51°32'40"N. In the spatial ecological model, one grid cell represents an area of 400m x 400m.

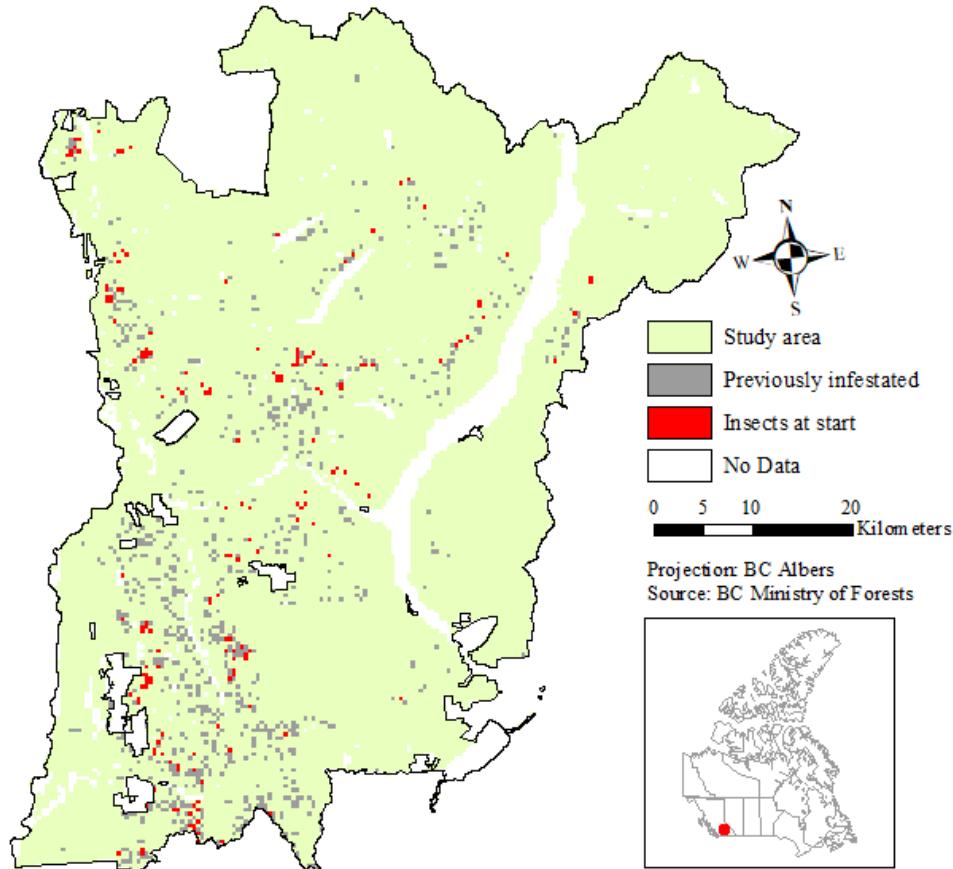


Figure 4.1. Study area of forest insect infestations simulation in British Columbia, Canada. Locations of active infestations at the beginning of the simulation period are marked as ‘Insects at start’. Locations that were infested before the start of simulations are marked as ‘Previously infested’.

One time-step in the coupled social-ecological model represents one year and the simulations ran for 10 years. In addition, separate sets of simulations were run with the social model alone, in which time-step was arbitrarily defined as one month. These simulations, which prepared the agents for later runs of the coupled social-ecological model, were performed with 10 and 20 additional steps.

4.2.1.3. Process overview and scheduling

In the coupled social-ecological model, in each time-step, infestations spread from infested cells in the previous time-step. This spread is simulated by the spatial ecological model. Newly infested grid cells remain invisible to the governing agent for one time-step after infestation and become visible in the next time-step. These grid cells act as sources of spread of infestations while they are invisible to the governing agent.

The social model's governing agent analyzes the last visible spread of infestations and calculates the cost of management action to stop the spread of infestations in each user zone. The management action is to cut a neighborhood area of the forest surrounding last visible infestations. Accordingly, the cost of action is defined proportional to the size of the said neighborhood. Then the governing agent sends a binary signal to the user agents. A signal of 1 means the governing agent requests the user agents to participate in the management action in their allocated forest zone voluntarily and at their own cost, in exchange for a '*responsible user*' label. The '*responsible user*' label only shows recognition of the user agents who cooperate with the governing agent, and has no monetary value. A signal of 0 means the governing agent requests the user agents to do an easy task with no cost for the user agents and no benefit for the governing agent, in exchange for the '*responsible user*' label. There is no difference between labels given when the governing agent's signal is 0 or 1. The governing agent uses a Reinforcement Learning (RL) algorithm in order to produce its signal, considering the past states, actions, and rewards (Harati, Perez, & Molowny-Horas, 2021). The reward for the governing agent is the cooperation by the user agents in the costly management action and thus saving the forest from the infestations.

Each user agent considers the governing agent's signal and produces a binary decision in response, which indicates whether or not the user agent accepts the governing agent's request in exchange for the '*responsible user*' label. When the governing agent's signal is 0, the requested action is of no cost and all user agents accept the governing agent's request and all user agents receive the '*responsible user*' label. When the governing agent's signal is 1, each user agent makes a decision with a cost-benefit analysis, taking into account the governing agent's calculated cost of action, history of the '*responsible user*' label, the uniqueness and visibility that they gain if they get the label, and the revenue from sales. Each user agent considers obtaining the label as an opportunity to be unique in having a recognition that some user agents do not. Such uniqueness is assessed based based on the response of user agents to the governing agent in the last known interaction. Each user agent also assumes that the acknowledgement of the label in its group depends on how much the group knows the label, which in turn depends on the cumulative number of times the label has been seen in the group. Thereupon, user agents consider a visibility score for the label.

In order to help understand interactions in the model, another agent, the registrar, is defined. The registrar observes and registers actions of the governing agent and the user agents in each time-step from the beginning of an episode of simulations. The other agents refer to the registrar in their decision making process.

Once the users make their decisions, a message is sent from the social ABM to the spatial ecological model, and modifications are correspondingly made in the forest map. These modifications include removing trees for annual harvest or management action. Specifically, each user's zone is subject to annual harvest unless that user cooperates with the governing agent when the signal is 1. In this case, that is, if the signal is 1 and the user agents cooperates with the governing agent, the neighborhood indicated for management action in the user's zone is cut. The modified landscape map is used by the spatial ecological model in the next time-step.

4.2.2. Design concepts

4.2.2.1. Theoretical and empirical background

The core idea of the social ABM is the promotion of responsible behavior using individuals' desire for respect. The theoretical basis for this idea notes that sustainability issues are problems of collective action (Ostrom, 1990), that an individual's behavior is influenced by the observation of behavior of others in the society, or descriptive norms, as stated in the theory of normative conduct (Cialdini et al., 1990), and that people care about their reputation in the society (Anderson, Hildreth, & Howland, 2015; Lazaric et al., 2020; Nolan, Schultz, Cialdini, Goldstein, & Griskevicius, 2008; Tascioglu, Eastman, & Iyer, 2017).

The two models that are coupled in this study are both taken from previous works. The social ABM has been built on the above concepts and calibrated through thousands of training iterations (Harati, Perez, & Molowny-Horas, 2021). The spatial ecological model has been developed, calibrated and tested with observed data (Harati et al., 2020). We refer the readers to these two papers for a detailed description of the models.

Complexities arise when the model's governing agent uses the '*responsible user*' label to encourage the user agents to engage in a costly action. At the beginning of the simulations, the label has not been introduced in the society of user agents and it is therefore not deemed valuable. Later on, as the label becomes more visible in society, its value increases in the calculations of the user agents. Meanwhile, an ecological disturbance causes damage to the forest

resource. The model sheds light on the complexities of the above said setting. Specifically, the model helps answer these questions about the possibility of success for the governing agent: Can the governing agent gain the cooperation of the user agents? Can they effectively control the disturbance? Can they save the resource?

The governing agent's decisions are based on bounded rationality (Simon, 1990). The governing agent does not know how the group of user agents behaves, it does not know what their decision thresholds are, and its information about the ecological system comes with a delay. The governing agent is designed in such a way that it observes the outcome of its actions, and learns to update its decision policy according to its observation. User agents make rational choices (Scott, 2000) based on information that is available to them. They do not modify their decision rule. User agents compare the utility of a suggestion with a threshold that indicates their hesitation, and make their decisions accordingly.

The social ABM uses input from the spatial ecological model. This input is the simulation of changes in a landscape, which is produced and processed through a GIS approach. This GIS approach does not take a time-series of external inputs during the simulations. However, the spatial ecological model is calibrated before the start of simulations using GIS data, which is available at grid cell level.

4.2.2.2. Individual decision-making

Subjects of decision making are the governing agent and user agents. The object of decision making of the governing agent is its binary signal, which is a variable that the governing agent communicates to user agents. The signal indicates whether the governing agent is requesting a costly action or no-cost action from the user agents. The object of decision making of each user agent is its decision, which is the user's response to the governing agent's signal request.

Since the governing agent's decisions are based on bounded rationality and it does not have perfect knowledge of the complex system that it deals with, it takes actions according to its available knowledge. Then based on the result of its action, the governing agent updates its decision policy using a RL algorithm. The governing agent's RL algorithm is a double-learning algorithm, which means it includes two arrays of scores of *(state,action)* pairs. These two arrays are updated iteratively in a convoluted manner, each based on the other.

The user agents' decisions are based on rational choice. User agents calculate the utility of cooperating with the governing agent, and compare it with an internal decision threshold. The utility that user agents calculate is a quantity between 0 and 1. If the calculated utility of a suggestion exceeds a user agent's decision threshold then the user agent accepts that suggestion. In calculation of the utility of the '*responsible user*' label, user agents take into account the *uniqueness* that they will have with the label, and the *visibility* of being associated with responsibility. They assess *uniqueness* based on the last known proportion of user agents who cooperated with the governing agent in a costly action. They assess *visibility* based on the total number of times the label has been presented in their society since the first time-step of the simulation. User agents calculate *uniqueness* and *visibility* based on the registrar's *nLast* and *nSum*.

User agents adapt to changes in their social and ecological environment. Social changes influence each user agent's perceived value of being recognized as a '*responsible user*', and ecological changes influence the size of the area where the management action is prescribed, hence influencing the cost of action required to receive the '*responsible user*' label. These variables do not change the decision rule of the user agent. The simulations shed light on the emergence of a norm of environmentally responsible behavior among user agents. On the other hand, the spread of infestations in the forest is a spatial process, which influences the governing agent's perceived state of forest health, and subsequently, cost of management action in each user zone.

All agents in the model use memory in their decisions. User agents refer to the registrar's memory. The governing agent, in addition to the memory of the registrar, uses its own built-in memory. The governing agent's RL algorithm applies a future discounting rate in the calculation of the present value of future consequences of its decisions.

The model includes some elements of uncertainty. The decision thresholds of user agents are taken from a normal distribution. The decision policy of the governing agent is defined stochastically. That is, for each (*state, action*) pair, the policy includes a number, which is used as a threshold for comparison against a random number. The decision is made according to that comparison.

4.2.2.3. *Learning*

Learning is the basis of the governing agent's RL algorithm. The RL algorithm keeps track of its states, actions, and rewards. The algorithm uses a policy to decide an action in each state. Then, based on the subsequent reward, the RL algorithm updates its policy. Through iterations, the governing agent's RL algorithm learns to adjust its policy in order to maximize its rewards. The model does not include collective learning.

4.2.2.4. *Individual sensing*

In this model, individuals are the agents in the social ABM. The model includes endogenous and exogenous variables. As for endogenous variables, user agents sense the governing agent's signal. User agents and the governing agent sense the total number of 'responsible user' labels as well as the last known number of labels given in a time-step when signal was 1. These variables are accessible to agents through the registrar. These endogenous variables are sensed without error. As for exogenous variables, the governing agent senses the changes that happen in the ecosystem. In our conceptual model, these changes are simulated by a spatial ecological model that is coupled to the social ABM. Therefore, this information is exogenous to the ABM. The sensing of environmental change is erroneous because in the definition of the model, environmental changes are not visible when they occur. The time lag between occurrence and visibility of the changes causes errors in the governing agent's sensing, thus adding to the complexity of the SES model. The governing agent and user agents sense the variables stored in the registrar, which is an auxiliary agent created for better understanding the model. The registrar, in turn, senses the governing agent's signal and each user agent's decision. These variables are sensed without error. The governing agent senses the spatial environment at global and local scales, when it calculates the overall state of health of the forest and the cost of management action in each user agent zone, respectively.

Within the social ABM, when the governing agent, users agents, or the registrar require information from another agent, they call that other agent. The agents are equipped with functions that send the requested information. Agents do not have direct access to variables of other agents. In the link between the social ABM and the spatial ecological model, each model is designed to perform some calculations, then wait for the other model to send the required information. This information is transferred by copying a file into the recipient model's inbox

directory. The model does not assume any costs associated with cognition or for gathering information.

4.2.2.5. Individual prediction

The governing agent's RL algorithm uses the data gained through experience in order to assess the values of its possible actions in the next step. The user agents consider data of the last known states in their calculations. The governing agent uses a temporal difference RL algorithm known as Double Expected SARSA (Sutton & Barto, 2018). The user agents assess the future value of obtaining the '*responsible user*' label with the assumption that the agents who previously chose a costly action in return for a label, will do so again. The predictions of the agents may be erroneous. User agents have limited ability to predict future changes in their society. Likewise, the governing agent's social prediction capability is limited. In addition, the governing agent's external input, which comes from the ecological spatial model, is designed to come with a delay.

4.2.2.6. Interaction

The model includes direct and indirect interactions among agents. Direct interactions include the communication of governing agent's signal and action cost calculations, as well as user agents' decisions. Indirect interactions occur due to user agents' desire to be better recognized than their peers, as well as through the market where all user agents sell their harvest. The governing agent's decisions and calculations depend on the history of responses from the user agents as well as the state of the ecological system. User agents' decisions depend on action costs, which are calculated through a spatial analysis. Interactions within the social ABM are communicated via the registrar. Interactions between the social ABM and the ecological spatial model are performed via file transfers, in which messages are copied into the recipient's inbox directory. The model does not include a coordination network.

4.2.2.7. Collectives

There are no collectives in this model.

4.2.2.8. Heterogeneity

User agents are heterogeneous in their decision thresholds, as well as their allocated forest zones. User agents and the governing agent are different in their decision making. The object of decision of the governing agent is the signal it sends to the user agents, and the objects of decisions of user agents are their responses to the governing agent. The governing agent uses a

RL algorithm in its decision, whereas user agents compare the utility of a suggestion with a threshold.

4.2.2.9. Stochasticity

The decision thresholds of the users are drawn from a normal distribution. The decision policy of the governing agent is stochastic.

4.2.2.10. Observation

In each time-step, the governing agent's signal, the proportion of user agents who cooperated with the governing agent, and remaining proportions of infested, non-infested, and harvested forest land are collected for analysis. In addition, for testing and verification of the model, all communications between the social ABM and the ecological spatial model are saved. Among the user agents, cooperation with the governing agent despite its cost is a behavior that emerges through simulations. In addition, saving forest areas from infestations is an emergent effect in the simulations.

4.2.3. Details

4.2.3.1. Implementation details

The social ABM was developed in Java, using features of REPAST (M. J. North et al., 2013). The spatial ecological model was developed in R (R Core Team, 2019). Please see the 'Data and code availability' section for links to model code and results.

4.2.3.2. Initialization

The social ABM consists of a governing agent, nine user agents, and a registrar agent. The governing agent's policy is defined by the results of a previous study (Harati, Perez, & Molowny-Horas, 2021), wherein the governing agent's RL algorithm was trained through interaction with the same number of user agents. There is no history of decisions of user agents, therefore the last known number of user agents cooperating with the governing agent is zero. In the previous study that defined the social ABM (Harati, Perez, & Molowny-Horas, 2021), three sets of simulations were run with mean user agent decision thresholds of 0.3 (low), 0.5 (medium) and 0.7 (high). In the present study, assuming that the governing agent does not have any information about the decision thresholds of user agents, we initialized the governing agent with the policy obtained from training with medium level user agent decision thresholds. The said training was the output of the previous study (Harati, Perez, & Molowny-Horas, 2021). In the

spatial ecological model, locations of insects at the start of simulations are extracted through GIS analysis of infestation data (BC Ministry of Forests, 2015).

There are some differences between various runs of the same series of simulations. The decision thresholds of user agents in the social ABM are drawn from a truncated normal distribution with pre-set mean and standard deviation. In new runs, new thresholds are drawn from the same normal distribution. Therefore, user agents change in new runs. For the governing agent, within the same set of simulations, the decision policy is updated based on rewards earned in the previous episode of runs.

4.2.3.3. Input data

In each time-step, the social ABM uses input from the ecological spatial model. This input is based on simulations of spread of infestations in the forest. As the simulated infestations spread further in the forest, the data transmitted to the social ABM changes over time.

4.2.3.4. Submodels

The social ABM includes a RL algorithm for the governing agent and a simple threshold decision-making algorithm for user agents. These are explained in detail in a previous work (Harati, Perez, & Molowny-Horas, 2021). The ecological spatial model is based on a logistic regression algorithm, which is explained in detail in another previous work (Harati et al., 2020). The social and ecological models are coupled through a mechanism that we call flip-flop. The social model requires inputs about the state of the forest and newly spread infestations, which is calculated in the ecological model. Conversely, the ecological model requires inputs on actions of user agents, which change land cover. In the flip-flop mechanism, each of the models runs its algorithm up to the moment it requires input from another model. Then it enters a loop in which it waits and observes an inbox directory that is allocated to that model in the computer's hard disk. In the meantime, the other model continues its calculations and eventually produces an output message file and sends it to the above-mentioned inbox directory. As soon as the message file is copied into the inbox directory, the first model notices the change in the contents of its inbox, exits the waiting loop, reads the file and resumes computing. In this way, models take alternative turns of running and pausing, hence the name 'flip-flop'. This strategy has enabled us to facilitate the exchange of information between two different algorithms (i.e. the social and the ecological models) that have been implemented in two different computer languages (Java and

R, respectively). Figure 4.2 depicts the concept of this coupling mechanism, and Figure 4.3 shows a more detailed view of the coupled model.




Figure 4.2. The flip-flop mechanism. Each of the two models requires input from the other model. The models communicate with each other via inbox directories. Arrows show direction of data transfer. Numbers beside arrows show the order of operations.

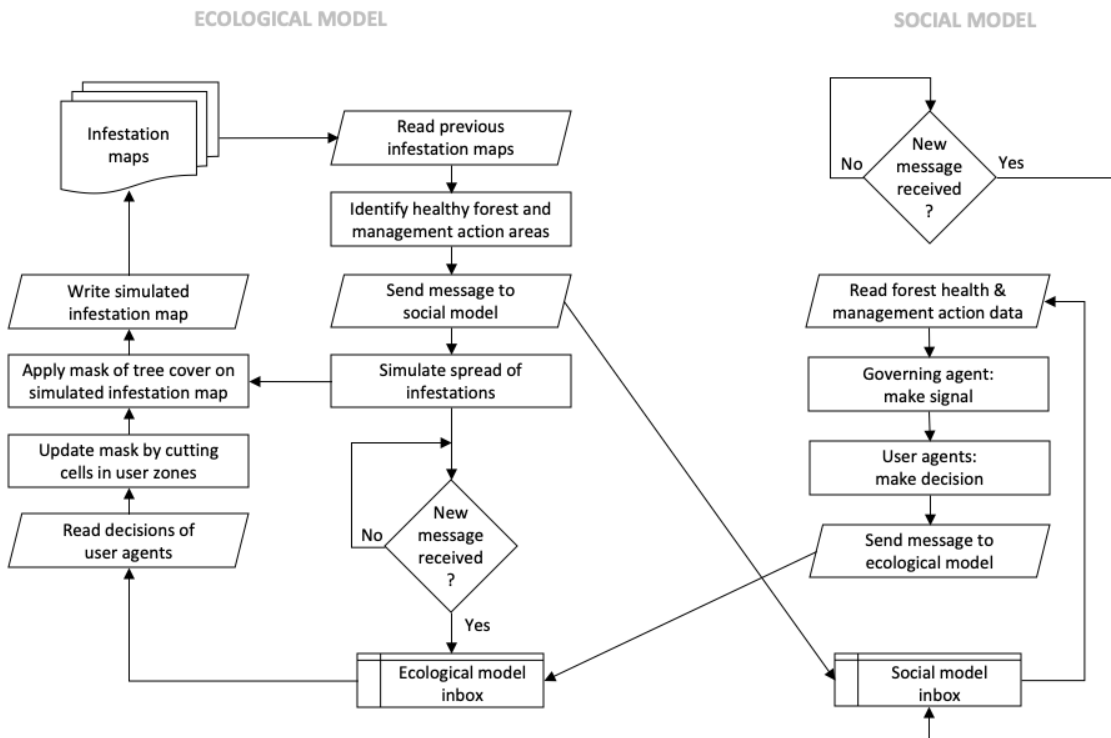


Figure 4.3. The coupled model. In each time-step, the ecological model begins with reading previous infestation maps and ends with writing a newly simulated infestation map; the social model begins with waiting for a message from the ecological model and ends with sending and message to the ecological model.

The management action that we consider in this study is cutting cells in a neighborhood of newly observed infestations. The size of this neighborhood is a parameter that needs to be defined. Based on insight obtained about spatial spread of MPB infestations in a previous study (Harati, Perez, Molowny-Horas, et al., 2021), in the present study we used Moore neighborhoods of size 4 to simulate the above-said management action. Considering that the cell-size in the model is 400 meters, the said neighborhood will be a square with side length of 3.6 kilometers. The rationale for this hypothetical neighborhood is that newly infested cells are not immediately detected. By the time infested cells change color and become observable, the infestation spreads further in the area.

The subject of calibration of the social ABM is the decision policy of its RL algorithm. This policy was learned previously (Harati, Perez, & Molowny-Horas, 2021) through 50 sets of 4000 training episodes each in a configuration with medium-level decision thresholds for user agents. Each set of 4000 episodes resulted in one (1) learned policy, thus, there were a total of fifty (50) learned policies. The mean of those 50 policies was used as the starting policy in the simulations of the present study. The spatial ecological model was calibrated using observed infestation data of years 2002-2004 for BC. Details of the model and its calibration are described in the corresponding previous work (Harati et al., 2020).

In addition to the social and ecological components, we initially included an economic submodel in the study. This submodel simulates a virtual market where user agents sell their harvested wood. In each time-step, the price of the market is adjusted according to the changes in wood supply. After several test runs of the coupled model we noted that the economic submodel largely influenced the results. Specifically, from the first time-step onward, the earning from additional harvest would become a strong incentive for user agents to accept the governing agent's requests of additional harvesting. This economic incentive would counter the user agents' perceived cost of action, and overshadow the social dynamics of emergence of environmentally responsible behavior. However, since the aim of the present study is to gain insight about the role of reputation in society as a driver for emergence of environmentally responsible behavior, we kept the economic submodel in the simulations but decided to limit the magnitude of the economic incentive with a cap.

As both the social ABM and the spatial ecological model are taken from previous works, we only made modifications to code and parameters for the coupling of the two models and the runs of this study. The social model's parameters include the number of user agents, mean and standard deviation of decision thresholds of user agents, future discounting rate, number of time-steps in one episode, and number of episodes. Each episode of simulations was run with a new set of user agents. In addition, in the present study we added a new parameter for the number of preparation steps before the social model is coupled with the spatial ecological model. The parameters of the coupling of the two models are the business-as-usual harvest ratio, which is the ratio of the study area that the user agents would harvest regardless of disturbance management, and the size of the neighbourhood of newly visible infestations, in which the management action of cutting cells is defined. Finally, the parameter applied to the economic submodel is the cap that limits its incentive effect. This means that the incentive associated with wood sales revenue is limited by a cap in the simulations, before it is considered in the cost-benefit decision making of the user agents. Table 4.2 shows the model parameters and their values. Note that in this table all values are dimensionless except for the management action neighborhood size, which is in grid cells.

Table 4.2. Model parameters

Parameter	Value(s)
Number of user agents	9
Mean decision threshold of user agents	0.7
Standard deviation of decision thresholds of user agents	0.08
Future discounting rate	0.1
Number of time-steps in one episode	10
Number of preparation time-steps	0, 10, 20
Number of episodes	50
Business-as-usual harvest ratio	0.01
Management action neighborhood	Moore, size 4
Economic incentive cap	0.1

4.2.4. Simulation scenarios

In order to gain insight about emergence of environmentally responsible behavior in the setting described above, we ran several rounds of simulations with different scenarios. We defined scenarios so that the comparison of their results provides useful information about the subject of study. Below are descriptions of these scenarios:

- The simplest scenario is Business As Usual (BAU), in which there is no intervention from the governing agent to control the disturbance. In each time-step, user agents harvest a proportion of their allocated zones. That proportion is the business-as-usual harvest ratio, which is a model parameter. The spatial ecological model iteratively simulates the spread of infestations, noting that harvested grid cells cannot become infested anymore. In this scenario, the governing agent's RL algorithm is not engaged. This scenario indicates a case where there is no government intervention to control the insect disturbance, or a case where user agents never cooperate with the governing agent. Hence, this scenario serves as a baseline for comparison with main simulations.
- Another baseline in our study is a scenario in which all user agents always cooperate fully with the governing agent. This scenario, which we named Enforce, provides a best case for the social component of our study. The Enforce scenario shows the effectiveness of the management plan in the control of the disturbance. In this scenario the governing agent's RL algorithm is not engaged.
- Our main scenarios, which we named Suggest, are those in which the governing agent is active and uses its RL algorithm. In Suggest scenarios, the governing agent suggests that if user agents cooperate with it in the management action then they will receive '*responsible user*' labels. User agents then analyze the governing agent's suggestion and make their decisions. In terms of cooperation of user agents with the governing agent, the Suggest scenarios are between BAU and Enforce. The neighborhood of management action is a Moore neighborhood of the newly visible infestations. The size of this neighborhood is a model parameter. We ran simulations with a neighborhood of size 4 grid cells. We also defined preparation runs, in which the ecological model is not engaged. Instead, agents in the social model interact with each other, which results in increased visibility and value of the '*responsible user*' label. Thereupon, the following three scenarios were defined:
 - Suggest scenario with 0 preparation time-steps
 - Suggest scenario with 10 preparation time-steps
 - Suggest scenario with 20 preparation time-steps

- Corresponding to each Suggest scenario, we defined another baseline, in which the governing agent behaves randomly instead of using its RL algorithm. In these scenarios, which we named Random, user agents analyze and respond to the governing agent's signals, as in the Suggest scenarios. The calculation of the state of health of the resource and the costs of management action in user agent zones are performed similar to the Suggest scenarios. The only difference between Suggest and Random scenarios is in the decision making mechanism of the governing agent. In this sense, by showing what could be achieved with a naïve model, Random scenarios serve as a baseline to indicate the power of the sophisticated RL algorithm of the governing agent. Thereupon, the following three scenarios were defined:
 - Random scenario with 0 preparation time-steps
 - Random scenario with 10 preparation time-steps
 - Random scenario with 20 preparation time-steps

4.3. Results

Figure 4.4 shows the mean maps of remaining infestations in the simulated scenarios at the final time-step. It can be seen that, without preparation, the Suggest and Random scenarios are similar to BAU. On the other hand, with addition of preparation steps, less infestation remains in the study area. The figure also shows that in Suggest scenarios less infestation remains in comparison with Random scenarios.

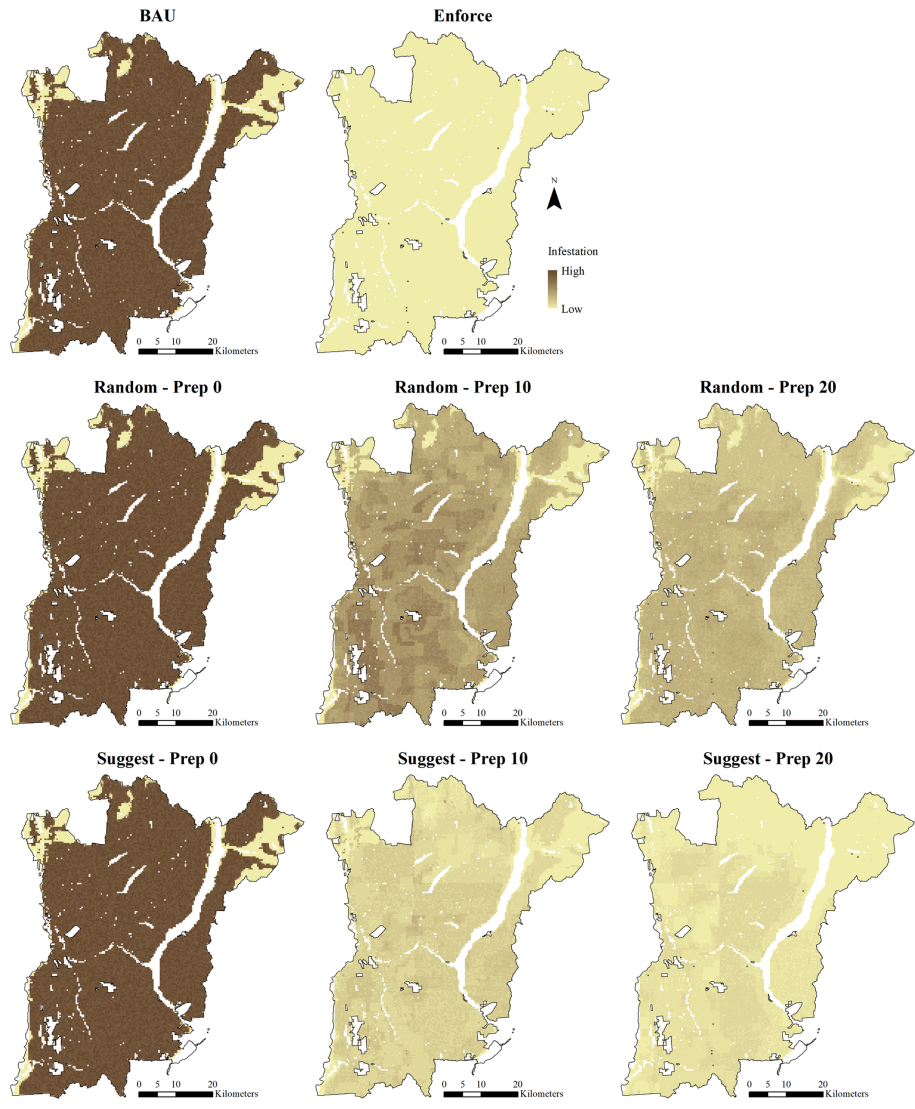


Figure 4.4. Maps of mean remaining infestation after the final time-step of simulations. The Enforce scenario was run once, and all other scenarios were run 50 times. For each scenario, 'High' infestation in a cell means the presence of infestation in the cell in all runs of that scenario; 'Low' infestation means the absence of infestation in the cell in all runs of that scenario.

Figure 4.5 shows the mean ratio of cooperation of user agents with the governing agent over time-steps of Suggest and Random scenarios. Without preparation of the user agents, both Suggest and Random scenarios end with nearly no cooperation at all. Therefore, in these cases no management action is done to control the infestations, which explains why the maps of no-preparation scenarios are similar to the map of BAU. With preparation, cooperation ratio

increases in both Suggest and Random scenarios, with Suggest scenarios showing higher cooperation than Random. Nonlinear behavior is observed in the curves of Suggest and Random scenarios with 10 steps of preparation, which shows sudden emergence of cooperation with the governing agent.

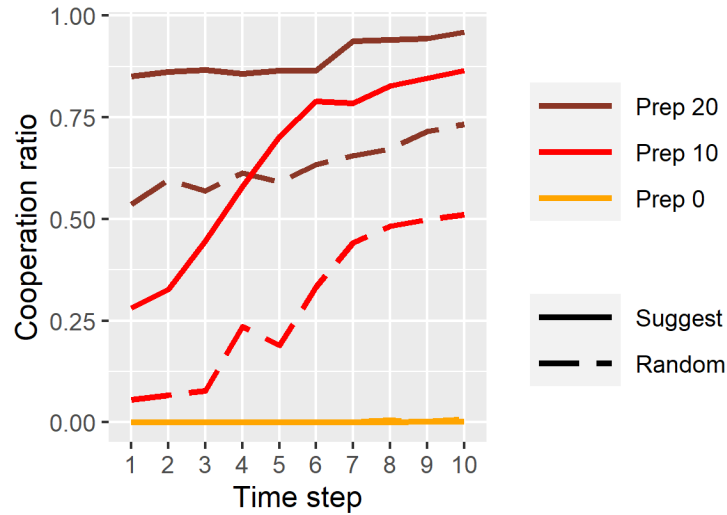


Figure 4.5. Mean cooperation ratio over time for Suggest and Random scenarios with 0, 10 and 20 preparation steps. Each scenario was run 50 times.

Figures 4.6 and 4.7 show the mean proportions of study area that are covered by healthy and infested forest, as well as the area that is cut, in each time-step. The baselines scenarios BAU and Enforce, which are shown in Figure 4.6, indicate the maximum amount of forest that can be saved from infestation if the management action is successfully implemented. The Random and Suggest scenarios, shown in Figure 4.7, demonstrate interim situations where the management action is partly implemented. The plots of Random and Suggest scenarios also show the results of adding preparation steps in the simulations.




Figure 4.6. Proportions of healthy, infested and cut areas in BAU and Enforce baseline scenarios.




Figure 4.7. Proportions of healthy, infested and cut areas in Random baseline and Suggest scenarios. Each plot represents 50 runs.

Table 4.3 gives a quantitative summary of the proportions of healthy, infested and cut area at the end of the simulations. Note that the Enforce scenario was run only once, because it involves no stochasticity in decisions of agents. This is why there is no deviation in the results of this scenario. The mean values in this table correspond to the final time-step in the plots of Figures 4.6 and 4.7. The standard deviations reveal several things about variations of the results. Firstly, variations are minimal in the BAU, as well as in Random and Suggest scenarios with no preparation. These are the scenarios in which user agents rarely cooperate with the governing agent, and therefore, the management action is not implemented. Variations in results increase substantially when preparation steps are included in simulations. Secondly, variations in the proportion of healthy area are smaller than variations in proportions of infested and cut areas. Thirdly, variations in Suggest scenarios are smaller than variations in Random scenarios. This is particularly evident in scenarios with 20 preparation steps. This table shows that, in comparison with the Random baseline scenarios, the Suggest scenarios result in a higher proportion of healthy forest at the end of the simulations, with smaller variations in results.

Table 4.3. Mean and standard deviation of the final proportions of healthy, infested and cut areas in simulations. Each scenario was run 50 times, except for Enforce, which run once.

Scenario	Healthy		Infested		Cut	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
BAU	0.076	0.001	0.827	0.001	0.097	0
Enforce	0.295	0	0.001	0	0.704	0
Random-Prep 0	0.075	0.003	0.825	0.008	0.100	0.010
Random-Prep 10	0.057	0.037	0.395	0.292	0.547	0.297
Random-Prep 20	0.187	0.110	0.245	0.321	0.568	0.243
Suggest-Prep 0	0.076	0.002	0.826	0.002	0.098	0.004
Suggest-Prep 10	0.135	0.084	0.126	0.249	0.739	0.239
Suggest-Prep 20	0.257	0.089	0.059	0.166	0.683	0.110

It can be seen in Figures 6 and 7 as well as Table 4.3 that results of the no-preparation scenarios are similar to BAU. On the other hand, as shown in Figure 4.5, in runs with more preparation steps, the percentage of cooperation of user agents with the governing agent increases. Such increase is larger when the governing agent's decisions are made by the RL algorithm, i.e., in Suggest scenarios. Comparing the zero-preparation and 20 step-preparation scenarios in Figures 4.6 and 4.7, it is evident that the 20 time-steps of preparation lead to an increase in the remaining healthy forest, even when the governing agent's decisions are random. Considering all cases with preparation of user agents, it can be observed that more healthy cells are saved in Suggest scenarios, i.e. with RL decision making algorithm for the governing agent than in Random scenarios, i.e. with random decisions. The case of 10-step preparation with random decisions of the governing agent is particularly noteworthy. In this scenario, first the infestation spreads to large areas, and then the management action begins, which involves cutting large areas around the newly observed infestations. Consequently, the remaining healthy forest in this scenario is even slightly smaller than BAU.

Tables 4.4 and 4.5 summarize non-parametric Wilcoxon signed rank tests that were performed to statistically analyze simulation results. Scenarios were compared in terms of proportions of remaining healthy forest after the final time-step of their respective simulations. As can be seen in Table 4.4, the proportions of remaining healthy forest in scenarios Suggest-Prep10, Suggest-Prep20 and Random-Prep20 are greater than BAU. Likewise, in comparison with Enforce, all scenarios lead to significantly less proportions of remaining healthy forest, except for Suggest-Prep20. In other words, the result of Suggest-Prep20 is very similar to our best cooperation case baseline. Table 4.5 shows the comparison of Random and Suggest scenarios. It is seen that without preparation steps, the results of Random and Suggest scenarios

are not significantly different. On the other hand, in scenarios with preparation steps, the proportion of remaining healthy forest is greater in Suggest scenarios than in Random scenarios.

Table 4.4. Wilcoxon signed rank test statistic and p-value for comparison of remaining healthy forest proportions of scenarios with BAU and with Enforce (n=50)

	Scenario	Null hypothesis	Alternative hypothesis	Statistic	p-value
Comparison with BAU	Random-Prep 0	Scenario = BAU	Scenario > BAU	638.0	0.50
	Random-Prep 10			295.0	0.99
	Random-Prep 20			1133.0	8.8×10^{-07}
	Suggest-Prep 0			669.5	0.38
	Suggest-Prep 10			1064.0	1.9×10^{-05}
	Suggest-Prep 20			1244.0	2.4×10^{-09}
Comparison with Enforce	Random-Prep 0	Scenario = Enforce	Scenario < Enforce	0	3.8×10^{-10}
	Random-Prep 10			0	3.8×10^{-10}
	Random-Prep 20			300	4.8×10^{-04}
	Suggest-Prep 0			0	3.8×10^{-10}
	Suggest-Prep 10			15	9.5×10^{-10}
	Suggest-Prep 20			903	0.99

Table 4.5. Wilcoxon signed rank test statistic and p-value for comparison of remaining healthy forest proportions between Random and Suggest scenarios (n=50)

Null hypothesis	Alternative hypothesis	Statistic	p-value
Random-Prep0 = Suggest-Prep0	Random-Prep0 < Suggest-Prep0	555.5	0.29
Random-Prep10 = Suggest-Prep10	Random-Prep10 < Suggest-Prep10	140.0	8.0×10^{-07}
Random-Prep20 = Suggest-Prep20	Random-Prep20 < Suggest-Prep20	109.0	5.7×10^{-03}

4.4. Discussion

4.4.1. Insights about the case of study

Our simulation results reveal several points that deserve further attention and discussion. One such point is the importance of the first time-step in the result of the simulations. In the comparison of the scenarios in Figure 4.7, those which have a smaller proportion of infested cells at the end of the first time-step have a larger proportion of healthy forest remaining at the end of the final time-step. Future infestations as well as the size of the areas to cut thus depend on infestations in the first time-step. Therefore, the more newly infested trees are cut in the first time-step, the more healthy forest remains in future time-steps. As an implication of the importance of success at the first step, and expanding this discussion beyond the scope of the present study, if the government has access to limited financial resources to provide incentives for the users, then our insight suggests that focusing the incentives at the beginning may contribute substantially to success in the end. The study of such cases and building hybrid

models for them, where the government uses both incentives and recognition, can be areas for future research.

Our results in Figure 4.7 show that forest cover is eliminated not only by infestations, but also by the considered management action, which is to cut trees in order to stop the progress of infestations. Depending on the scenario, the proportion of the forest that is cut may even be larger than the proportion that is infested. Particularly, when infestations spread in a large area and user agents decide to cooperate with the governing agent, user agents will cut a large zone in the forest such that a small proportion of healthy forest will remain. Likewise, in the Enforce scenario, control of the epidemic is also achieved at the cost of cutting a large area of the forest, but in this case the proportion of remaining healthy forest is larger, and the participation of user agents is enforced, not voluntary.

An assumption in the simulations of this study is that for the user agents the perceived cost of cooperation with the governing agent is high. This cost appears as a threshold in the decision process of user agents, and it is compared against the calculated utility of cooperation with the governing agent. In the study that presented and described our social model (Harati, Perez, & Molowny-Horas, 2021), user agents' perceived cost of cooperation with the governing agent is defined in three levels of low, medium and high, where high cost of an action means low motivation to perform that action. Therefore, the present study involves a challenge for the social model, because our user agents are defined to be hesitant towards cooperation with the governing agent. In such challenging problems with high cost of the desired behavior, the governing agent may succeed in promoting the desired behavior with appropriate decision making algorithms but not with random decisions (Harati, Perez, & Molowny-Horas, 2021). The technique that we added in the present study was the inclusion of additional preparatory steps. We found that preparatory steps lead to the emergence of the desired behavior among user agents, even if the governing agent's decisions are made randomly. This shows that recognizing responsible users and introducing them to the society is a powerful mechanism with the potential to create new behavior norms.

The governing agent's RL algorithm uses a policy for making decisions, and updates that policy based on the rewards it receives as a result of its decisions. The governing agent is not aware of the decision thresholds of user agents. In the study that defined the social model

(Harati, Perez, & Molowny-Horas, 2021), in each of the cases with low, medium and high user agent thresholds, the governing agent's RL algorithm was trained and calibrated iteratively. Therefore, corresponding to each case a policy was calculated. In our present study, still assuming that the governing agent does not know if the decision thresholds of user agents are low, medium or high, we have parameterized the governing agent with the policy calculated previously in training the RL algorithm with medium-threshold user agents (Harati, Perez, & Molowny-Horas, 2021). Note that this initial parameterization does not match our choice of high decision thresholds for user agents in the present study. That is a price that has to be paid for the reasonable assumption that the governing agent is not aware of decision thresholds of the user agents.

Another interesting matter about the simulations is that they show that, in all scenarios, infestations spread rapidly at first and slow down later, such that by the final time-steps little or no change is noticeable in the proportion of the study area that is infested. As the spread of infestations stops, there will be no areas to cut around observed infestations, and the composition of the study area does not change anymore. It is noteworthy that in BAU the largest spread of infestations happens in the first three time-steps. Therefore, in the Suggest scenarios, if no management action is taken in these initial time-steps, then a large part of the forest is destroyed by infestations.

A feature of our simulations was the definition of the management action. Cutting neighborhoods of infestations is only one of the possible actions to control infestations (Maclauchlan & Brooks, 1994). Our inspiration for choosing this action came from a previous work, in which we performed spatial analyses of spread of MPB infestations and validation tests on our land change model (Harati, Perez, Molowny-Horas, et al., 2021). In that work, we noted that most of the new MPB infestations occur in the vicinity of previous locations of attacked trees. Therefore, in the present study we defined neighborhoods of management action around newly observed infestations. The action was to cut the cells in these neighborhoods. Our choice of neighborhood size was also inspired by distance analysis in that previous work (Harati, Perez, Molowny-Horas, et al., 2021) as well as the consideration that the governing agent cannot detect new infestations in the first time-step after their occurrence. This means that infestations spread further in a larger neighborhood before the governing agent realizes their previous locations.

Through the course of MPB infestations, BC eventually adopted the policy of increasing the allowable annual cut, first in order to suppress the infestations, and later to facilitate salvage harvest in infested areas (Forest Practices Board, 2007, 2009). Such increase was smaller than the simulations of the present study. In severely infested areas of the province, the allowable annual cut was increased by 80 percent of the pre-infestation levels (*idem*), which was less than one third of a percent of the forest area (BC Ministry of Forests, 2003). Therefore, at the largest increase of the allowable annual cut, still less than one percent of the forest has been harvested per year.

Cutting trees in large scale can be a practical challenge. According to an analysis of data of year 2008, with that year's rate of harvesting, it would take 22 years to cut the pine trees that were killed by infestations up to 2008 (Forest Practices Board, 2009). Moreover, although the government wanted the added harvesting to be concentrated in severely infested areas to control the pest, the forest industry preferred to harvest from other locations and especially from forest stands with mixed species (Forest Practices Board, 2007). Furthermore, the decision to increase the allowable annual cut raised concerns about possible ecological impacts (Forest Practices Board, 2007, 2009), which is beyond the scope of the present work and can be studied in future research.

4.4.2. Insights about governance of SES

Our most important finding in this study is that it is possible to create a strong motivation for effective action towards protection of natural resources in a SES by encouraging users – without financial incentives, enforcement or punishment. This point is of particular importance because the role of punishment as a basis for the formation of norms of environmentally responsible behavior has been emphasized in the literature of SES (Farrow, Grolleau, & Ibanez, 2017) and social sciences (Axelrod, 1986). Our results show that even without punishment, recognition of responsible behavior through the mechanism of our model can create a force towards emergence of a norm of responsible behavior. Our approach, which is only one of the possible approaches to the problem of collective action, is based on the theory of normative conduct (Cialdini et al., 1990). Moreover, from model results presented in Figure 4.7 we gain the insight that even with existence of the potential for action towards protection of a natural resource, uncalculated governance decisions about using that potential may cause adverse effects on the resource. The simulations show an example of this type. The difference between model

results with and without an intelligent algorithm highlights the importance of well-thought-through decision making in governance of SES.

Another general insight that we gained from the simulations pertains to the temporal differences between social and ecological processes, which add to the complexity of a SES. For example, from our simulations we noted that formation of an environmentally responsible behavior norm takes time. We also noted, throughout the simulations produced, that the largest damage made by the ecological disturbance occurred in the beginning time-steps of the study. Therefore, if the efforts to promote a new norm of responsible behavior begin at the onset of the ecological disturbance, then by the time the behavior emerges in the society and actions to protect the resource start, much of the resource is already lost. This means that it is important to prepare the society and promote environmentally responsible behavior before there is any need for action to protect the natural resource.

4.4.3. Limitations of the study and prospects for future works

Our model is defined in a specific scope. This scope might as well be considered a limitation for the model. That is, for example, our model does not account for the effect of other social processes than what we included in it. Future works may use other social models instead of ours and insert them in our SES model. Potential research efforts may as well couple other ecological models to the social model. In these cases, the coupling mechanism of our model may be modified according to the needs of other applications and assumptions. In our study we noted that the economic model, especially in the first time-step, may produce a high financial incentive for the user agents to cooperate with the governing agent. However, our aim was to study the power and effect of the social mechanism for creation of costly behaviors. Therefore, we applied a cap on the output of the economic model, so that the task of promoting cooperation remained challenging for the social model. Future works may, depending on their objectives, run the economic model without any limits, and simulate other complexities. Another matter that may be considered in future works is the interactions among user agents. In the present study our goal was to gain useful insight for managerial and government decision making. Therefore, our RL decision making algorithm was placed in the governing agent. Future works can equip user agents with more intelligent algorithms as well, and expand their analyses to include interactions between user agents and collaboration among them.

4.5. Conclusion

In this study we connected two social and ecological models, which were developed independently, through a coupling mechanism to build a conceptual social-ecological model. By using our model we carried out tests that allowed us to perform ‘what-if’ analyses with several scenarios of SES management. Our simulations showed that in a society where individuals or companies (i.e. “agents”) care about their reputation, it is possible to promote environmentally responsible behavior through an encouragement mechanism, without use of force, without use of financial incentives, and only by recognition of responsible individuals in the society. In the management of a SES under disturbance, it is important to note that during the time before the emergence of environmentally responsible behavior, the ecological system may be damaged by the disturbance, as demonstrated in our zero-preparation scenario simulations. It is therefore important to prepare the society in advance for engagement in environmental protection and ecological conservation action. However, allocating time for preparation of the society for environmental action does not guarantee success in management of disturbances in a SES. This was demonstrated in our simulations of Suggest and Random scenarios, which showed the difference between the outcomes of two sequences of governance decisions with and without careful consideration of users’ response and the state of the ecosystem. By defining the BAU and Enforce scenarios to represent zero and full cooperation of users, we built reference cases to demonstrate the range of possible outcomes of the model. Then, by comparing these reference cases with our simulations of Suggest and Random scenarios, we were able to identify where our simulation results stand in the space of possible outcomes. Through these hypothetical experiments, we were able to gain insights about a complex system using a conceptual model.

Funding: This research was partially funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada through the Discovery Grant number RGPIN/05396-2016 awarded to L.P., R.M.-H. received financial support from the European Union’s Seventh Framework Programme through NEWFOREST programme number PIRSES-GA-2013-612645.

Data and code availability: Model code files are available at <https://github.com/s-harati/model-FlipFlop1-ProMEERB> (accessed December 16, 2021). Datasets of model input, calibration, and results are available at <https://osf.io/x8huq/> (accessed December 16, 2021).

Conflict of interest: The authors declare no conflict of interest.

Acknowledgements: We are thankful to the Université de Montréal's International Affairs Office (IAO) for their financial support through the International Partnership Development program, which allowed the collaboration between researchers from UdeM and CREA.F.

References

- Anderson, C., Hildreth, J. A. D., & Howland, L. (2015). Is the desire for status a fundamental human motive? A review of the empirical literature. *Psychological Bulletin*, *141*(3), 574–601. <https://doi.org/10.1037/a0038781>
- Axelrod, R. (1986). An evolutionary approach to norms. *The American Political Science Review*, *80*(4), 1095–1111.
- Barbier, E. B. (1987). The Concept of Sustainable Economic Development. *Environmental Conservation*, *14*(2), 101–110. <https://doi.org/10.1017/S0376892900011449>
- Batty, M., Xie, Y., & Sun, Z. (1999). Modeling urban dynamics through GIS-based cellular automata. *Computers, Environment and Urban Systems*, *23*(3), 205–233. [https://doi.org/10.1016/S0198-9715\(99\)00015-0](https://doi.org/10.1016/S0198-9715(99)00015-0)
- Batty, Michael, & Torrens, P. M. (2005). Modelling and prediction in a complex world. *Futures*, *37*(7), 745–766. <https://doi.org/10.1016/j.futures.2004.11.003>
- BC Ministry of Forests. (2003). *British Columbia's Forests And Their Management*.
- BC Ministry of Forests. (2015). [dataset] BC MPB Observed Cumulative Kill - vol.12.
- Berkes, F., & Folke, C. (2000). Linking social and ecological systems for resilience and sustainability. In F. Berkes, C. Folke, & J. Colding (Eds.), *Linking Social and Ecological Systems: Management Practices and Social Mechanisms for Building Resilience* (pp. 1–26). Cambridge University Press.
- Bone, C., Dragicevic, S., & Roberts, A. (2006). A fuzzy-constrained cellular automata model of forest insect infestations. *Ecological Modelling*, *192*(1–2), 107–125. <https://doi.org/10.1016/j.ecolmodel.2005.09.013>
- Bonnell, T. R., Campenni, M., Chapman, C. A., Gogarten, J. F., Reyna-Hurtado, R. A., Teichroeb, J. A., ... Sengupta, R. (2013). Emergent Group Level Navigation: An Agent-Based Evaluation of Movement Patterns in a Folivorous Primate. *PLoS ONE*, *8*(10). <https://doi.org/10.1371/journal.pone.0078264>
- Bourceret, A., Amblard, L., & Mathias, J. D. (2021). Governance in social-ecological agent-based models: A review. *Ecology and Society*, *26*(2). <https://doi.org/10.5751/ES-12440-260238>
- Brown, B. J., Hanson, M. E., Liverman, D. M., & Merideth, R. W. (1987). Global sustainability: Toward definition. *Environmental Management*, *11*(6), 713–719. <https://doi.org/10.1007/BF01867238>
- Brundtland, G. H. (1987). *Report of the World Commission on environment and development: "our common future"*.
- Castle, C., & Crooks, A. (2006). *Principles and concepts of agent-based modelling for developing geospatial simulations*. London, UK: Centre for Advanced Spatial Analysis (UCL).
- Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology*, *58*(6), 1015–1026. <https://doi.org/10.1037/0022-3514.58.6.1015>
- Clarke, K. C., Hoppen, S., & Gaydos, L. (1997). A self-modifying cellular automaton model of

- historical urbanization in the San Francisco Bay area. *Environment and Planning B: Planning and Design*, 24(2), 247–261. <https://doi.org/10.1068/b240247>
- Cosens, B., Ruhl, J. B., Soininen, N., Gunderson, L., Belinskij, A., Blenckner, T., ... Similä, J. (2021). Governing complexity: Integrating science, governance, and law to manage accelerating change in the globalized commons. *Proceedings of the National Academy of Sciences*, 118(36), e2102798118. <https://doi.org/10.1073/pnas.2102798118>
- Crawford, S. E. ., & Ostrom, E. (1995). A grammar of institutions. *American Political Science Review*, 89(3), 582–600.
- Daily, G. C. (2000). Management objectives for the protection of ecosystem services. *Environmental Science & Policy*, 3(2000), 333–339. [https://doi.org/10.1016/S1462-9011\(00\)00102-7](https://doi.org/10.1016/S1462-9011(00)00102-7)
- de Almeida, C. M., Batty, M., Monteiro, A. M. V., Câmara, G., Soares-Filho, B. S., Cerqueira, G. C., & Pennachin, C. L. (2003). Stochastic cellular automata modeling of urban land use dynamics: Empirical development and estimation. *Computers, Environment and Urban Systems*, 27(5), 481–509. [https://doi.org/10.1016/S0198-9715\(02\)00042-X](https://doi.org/10.1016/S0198-9715(02)00042-X)
- Farrow, K., Grolleau, G., & Ibanez, L. (2017, October 1). Social Norms and Pro-environmental Behavior: A Review of the Evidence. *Ecological Economics*. Elsevier B.V. <https://doi.org/10.1016/j.ecolecon.2017.04.017>
- Filotas, E., Parrott, L., Burton, P. J., Chazdon, R. L., Coates, K. D., Coll, L., ... Messier, C. (2014). Viewing forests through the lens of complex systems science. *Ecosphere*, 5(January), 1–23. <https://doi.org/10.1890/ES13-00182.1>
- Forest Practices Board. (2007). *Tree Species Harvested In Areas Affected By Mountain Pine Beetles Special Report Table of Contents*.
- Forest Practices Board. (2009). *Biodiversity conservation during salvage logging in the Central Interior of BC*. Retrieved from <https://www.bcfpb.ca/wp-content/uploads/2016/04/SR35-Salvage-Logging.pdf>
- Gaudreau, J., Perez, L., & Drapeau, P. (2016). BorealFireSim: A GIS-based cellular automata model of wildfires for the boreal forest of Quebec in a climate change paradigm. *Ecological Informatics*, 32, 12–27. <https://doi.org/10.1016/j.ecoinf.2015.12.006>
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., ... DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126. <https://doi.org/10.1016/j.ecolmodel.2006.04.023>
- Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, 221(23), 2760–2768. <https://doi.org/10.1016/j.ecolmodel.2010.08.019>
- Grimm, V., Railsback, S. F., Vincenot, C. E., Berger, U., Gallagher, C., DeAngelis, D. L., ... Ayllón, D. (2020). The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism. *Journal of Artificial Societies and Social Simulation*, 23(2), 1–7. <https://doi.org/10.18564/jasss.4259>
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., ... DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750), 987–991. <https://doi.org/10.1126/science.1116681>
- Harati, S., Perez, L., & Molowny-Horas, R. (2020). Integrating neighborhood effect and supervised machine learning techniques to model and simulate forest insect outbreaks in british columbia, canada. *Forests*, 11(11), 1–23. <https://doi.org/10.3390/f11111215>
- Harati, S., Perez, L., & Molowny-Horas, R. (2021). Promoting the emergence of behavior norms

- in a principal–agent problem—an agent-based modeling approach using reinforcement learning. *Applied Sciences (Switzerland)*, *11*(18). <https://doi.org/10.3390/app11188368>
- Harati, S., Perez, L., Molowny-Horas, R., & Pontius, R. G. (2021). Validating models of one-way land change: an example case of forest insect disturbance. *Landscape Ecology*, *36*(10), 2919–2935. <https://doi.org/10.1007/s10980-021-01272-0>
- Janssen, M. A., & Ostrom, E. (2006). Empirically Based , Agent-based models. *Ecology And Society*, *11*(2). Retrieved from <http://www.ecologyandsociety.org/vol11/iss2/art37/>
- Janssen, Marco A., & Ostrom, E. (2006). Governing Social-Ecological Systems. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of Computational Economics (vol. 2)* (pp. 1465–1509). Elsevier B.V. [https://doi.org/10.1016/S1574-0021\(05\)02030-7](https://doi.org/10.1016/S1574-0021(05)02030-7)
- Katan, J., & Perez, L. (2021). ABWiSE v1.0: Toward and Agent-Based Approach to Simulating Wildfire Spread. *Natural Hazards and Earth System Sciences*, 1–27. <https://doi.org/10.5194/nhess-2021-179>
- Kriebel, D., Tickner, J., Epstein, P., Lemons, J., Levins, R., Loechler, E. L., ... Stoto, M. (2001). The precautionary principle in environmental science. *Environmental Health Perspectives*, *109*(9), 871–876. <https://doi.org/10.1289/ehp.01109871>
- Lambin, E. F., & Geist, H. (Eds.). (2006). *Land-Use and Land-Cover Change*. Berlin, Heidelberg: Springer Berlin Heidelberg. <https://doi.org/10.1007/3-540-32202-7>
- Lazaric, N., Le Guel, F., Belin, J., Oltra, V., Lavaud, S., & Douai, A. (2020). Determinants of sustainable consumption in France: the importance of social influence and environmental values. *Journal of Evolutionary Economics*, *30*(5), 1337–1366. <https://doi.org/10.1007/s00191-019-00654-7>
- Liu, J., Dietz, T., Carpenter, S. R., Alberti, M., Folke, C., Moran, E., ... Taylor, W. W. (2007). Complexity of Coupled Human and Natural Systems. *Science*, *317*(5844), 1513–1516. <https://doi.org/10.1126/science.1144004>
- Maclauchlan, L. E., & Brooks, J. E. (1994). *Strategies and tactics for managing the mountain pine beetle Dendroctonus ponderosae*.
- Millennium Ecosystem Assessment. (2003). Ecosystems and their services. In *Ecosystems and Human Well-being: A Framework for Assessment* (pp. 49–70). Washington, DC: Island Press. <https://doi.org/10.1007/s13398-014-0173-7.2>
- Mosler, H.-J. (1993). Self-dissemination of environmentally-responsible behavior: The influence of trust in a commons dilemma game. *Journal of Environmental Psychology*, *13*(2), 111–123. [https://doi.org/10.1016/S0272-4944\(05\)80144-6](https://doi.org/10.1016/S0272-4944(05)80144-6)
- Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., ... Schwarz, N. (2013). Describing human decisions in agent-based models – ODD + D, an extension of the ODD protocol. *Environmental Modelling & Software*, *48*, 37–48. <https://doi.org/10.1016/j.envsoft.2013.06.003>
- National Research Council. (2014). *Advancing Land Change Modeling: Opportunities and Research Requirements*. Washington, D.C.: National Academies Press. <https://doi.org/10.17226/18385>
- Nolan, J. M., Schultz, P. W., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2008). Normative Social Influence is Underdetected. *Personality and Social Psychology Bulletin*, *34*(7), 913–923. <https://doi.org/10.1177/0146167208316691>
- North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CBO9780511808678>
- North, M. J., Collier, N. T., Ozik, J., Tataru, E. R., Macal, C. M., Bragen, M., & Sydelko, P.

- (2013). Complex adaptive systems modeling with Repast Symphony. *Complex Adaptive Systems Modeling*, 1, 3. <https://doi.org/10.1186/2194-3206-1-3>
- Nyborg, K., Anderies, J. M., Dannenberg, A., Lindahl, T., Schill, C., Schlüter, M., ... de Zeeuw, A. (2016). Social norms as solutions. *Science*, 354(6308), 42–43. <https://doi.org/10.1126/science.aaf8317>
- O’Sullivan, D. (2004). Complexity Science and Human Geography. *Transactions of the Institute of British Geographers*, 29(3), 282–295. <https://doi.org/10.1111/j.0020-2754.2004.00321.x>
- Omoto, A. M., & Snyder, M. (1995). Sustained helping without obligation: Motivation, longevity of service, and perceived attitude change among AIDS volunteers. *Journal of Personality and Social Psychology*, 68(4), 671–686. <https://doi.org/10.1037/0022-3514.68.4.671>
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge, UK: Cambridge University Press.
- Ostrom, E. (2009). A General Framework for Analyzing Sustainability of Social-Ecological Systems. *Science*, 325(July), 419–422. <https://doi.org/10.1126/science.1172133>
- Parrott, L., Chion, C., Martins, C. C. A., Lamontagne, P., Turgeon, S., Landry, J. A., ... Dionne, S. (2011). A decision support system to assist the sustainable management of navigation activities in the St. Lawrence River Estuary, Canada. *Environmental Modelling and Software*, 26(12), 1403–1418. <https://doi.org/10.1016/j.envsoft.2011.08.009>
- Perez, L., & Dragicevic, S. (2009). An agent-based approach for modeling dynamics of contagious disease spread. *International Journal of Health Geographics*, 8(1), 50. <https://doi.org/10.1186/1476-072X-8-50>
- Perez, L., & Dragicevic, S. (2010). Modeling mountain pine beetle infestation with an agent-based approach at two spatial scales. *Environmental Modelling & Software*, 25(2), 223–236. <https://doi.org/10.1016/j.envsoft.2009.08.004>
- Pooyandeh, M., & Marceau, D. J. (2013). A spatial web/agent-based model to support stakeholders’ negotiation regarding land development. *Journal of Environmental Management*, 129, 309–323. <https://doi.org/10.1016/j.jenvman.2013.07.028>
- R Core Team. (2019). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/>
- Ross, H. L. (1973). *Perspectives on social order*. New York: McGraw-Hill.
- Savarimuthu, B. T. R., & Cranefield, S. (2011). Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, 7(1), 21–54. <https://doi.org/10.3233/MGS-2011-0167>
- Schlüter, M., Hinkel, J., Bots, P. W. G., & Arlinghaus, R. (2014). Application of the SES framework for model-based analysis of the dynamics of social-ecological systems. *Ecology and Society*, 19(1). <https://doi.org/10.5751/ES-05782-190136>
- Scott, J. (2000). Rational Choice Theory. In G. Browning, A. Halcli, & F. Webster (Eds.), *Understanding Contemporary Society: Theories of the Present* (pp. 126–138). SAGE Publications Ltd. <https://doi.org/10.4135/9781446218310.n9>
- Simon, H. A. (1990). Bounded Rationality. In J. Eatwell, M. Milgate, & P. Newman (Eds.), *Utility and Probability* (pp. 15–18). London: Palgrave Macmillan UK. https://doi.org/10.1007/978-1-349-20568-4_5
- Stern, P. C., Dietz, T., & Kalof, L. (1993). Value Orientations, Gender, and Environmental Concern. *Environment and Behavior*, 25(5), 322–348. <https://doi.org/10.1177/0013916593255002>

- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). Cambridge, Massachusetts: MIT Press.
- Tascioglu, M., Eastman, J. K., & Iyer, R. (2017). The impact of the motivation for status on consumers' perceptions of retailer sustainability: the moderating impact of collectivism and materialism. *Journal of Consumer Marketing*, 34(4), 292–305. <https://doi.org/10.1108/JCM-03-2015-1351>
- White, R., & Engelen, G. (1993). Cellular automata and fractal urban form: a cellular modelling approach to the evolution of urban land-use patterns. *Environment and Planning A*, 25(8), 1175–1199. <https://doi.org/10.1068/a251175>
- Wimolsakcharoen, W., Dumrongrojwatthana, P., Le Page, C., Bousquet, F., & Trébuil, G. (2021). An agent-based model to support community forest management and non-timber forest product harvesting in northern Thailand. *Socio-Environmental Systems Modelling*, 3, 17894. <https://doi.org/10.18174/sesmo.2021a17894>

Conclusions

In my doctoral work I addressed a sustainability governance issue in the context of an ecological disturbance. This disturbance was the spread of insect infestations in a forest, and the governance issue was the engagement of a group of individuals in protecting their common resource. The definition of the problem in this doctoral work was one of its noteworthy features. I proposed a particular mechanism of interactions between government and resource users, and asked if it is possible to use that mechanism in order to elicit the participation of the users in the government's intended task. The mechanism involves the government offering a recognition label to users in exchange for an action that is decided by the government in each time step. Although comparable interaction structures with different objectives were found in the literature (Bone and Dragičević, 2010; Ostrom, 2009; Yang and Wu, 2009), the use of this particular recognition mechanism for encouraging user participation in resource protection was a novelty in the present work.

My overall approach in this study was to break the problem into components, to simplify and analyze each component, and to combine them again to gain insight about the original problem. Throughout the works, I assumed the viewpoint of a government that desires to protect and conserve an ecological resource with the help of users of the resource. In Chapter 1 I showed with an abstract model that the recognition mechanism may lead to emergence of a new behavior that the government desires. In Chapters 2 and 3 I described the construction and validation of a land change model, which simulates the ecological phenomenon of this study. Then, in Chapter 4 I coupled the two models to build a virtual laboratory, in which I performed several virtual experiments. In this sense, Chapter 4 demonstrates the final step of bringing everything together and answering the research problem, whereas the works reported in Chapters 1-3 set the ground for the analysis described in Chapter 4. Nevertheless, all parts of the work involved innovative features, led to notable findings, and opened up new questions for further research, as described below.

Novelty and findings

Chapter 1

Chapter 1 marked that elicitation of cooperation from users is possible, even in costly tasks. The surprising aspect of this finding was that the model that generated these results assumed users to act upon self-interest and not based on altruism. A remarkable feature of this

model was that it did not provide the governing agent with decision making heuristics. Rather, the governing agent learned by itself to first promote the *'responsible user'* label by giving free labels to all user agents, and then to ask user agents to perform a costly action in exchange for the label. To do so, the governing agent had to guess the decision criteria of user agents, of which it did not have any prior information. Interestingly, the governing agent's RL algorithm managed to solve both problems, that is, to develop a nearly two-stage policy, and to find when to switch from one stage to another. These achievements owe in part to such features of the algorithm as a 2-dimensional representation of abstract states and a stochastic policy.

In addition to the surprising conclusion that indicates the possibility of promoting a costly behavior without force or financial incentives, the conceptual model of Chapter 1 provided an important output. The model was coded to save its trained policy. That policy was later served as a starting point for runs of the coupled social-ecological model in Chapter 4. This means that the coupled model of Chapter 4 would not need to go through thousands of episodes of training. It is important to note that the assumption of a governing agent's lack of knowledge about decision criteria of user agents was maintained, since what was transferred to the model of Chapter 4 was the mean policy of Chapter 1 model's runs with medium-level user agents.

Interestingly, it was noted that similar research questions appear in literature of other domains, which may seem irrelevant to sustainability. These include marketing (Delre et al., 2007; Rand and Rust, 2011) and innovation diffusion (Delre et al., 2010; Kiesling et al., 2012). This shows that the model of Chapter 1 can have diverse applications. It also shows that this study in particular and sustainability studies in general can benefit from the interaction with a diverse variety of other fields such as the above mentioned.

Chapter 2

One of the noteworthy features of the spatial model of Chapter 2 was that it integrated several distance functions – namely, uniform, linear, inverse, squared inverse – as explanatory variables to reflect the effect of distance from previous infestations on the probability of infestation in future. This novel design compensated for the insufficient spatial information about patterns of spread of the studied insect. The model trained itself using calibration data and weighted the distance functions in such a way that their combination produced spatial patterns that best matched observation data.

Chapter 2 also presented an example case for comparison of two modelling techniques – Random Forest (RF) and Logistic Regression (LR). In the model of Chapter 2, these techniques produce regression maps. That is, the value that they generate for each cell is regressed over explanatory variables. Therefore, the generated cell values are initially continuous. These values are then converted to binary through comparison with a final cut-off threshold. Errors of the simulations were identified through comparison of such cut-off images with reference data. One of the questions that arose then was which modelling technique had a better performance.

Various metrics were used in order to compare the RF and LR models. The overall accuracy of all models was high, with RF being the highest. The larger part of the landscape were areas that were far from infestations and remained unchanged in all simulations as well as in reference data. In other words, the overall accuracy awarded the models not for the changes that they predicted, but for their correct rejections. Because prevalence of change in the study area was small, it was expectable that the simulations produce a lot of correct rejections regardless of the simulation algorithm, and thus high overall accuracy values. On the other hand, values of the figure of merit were substantially lower than overall accuracy because the figure of merit does not take into account the correct rejections. However, the values of the figure of merit raised more questions. It was particularly curious that RF simulations, which had the highest overall accuracy, showed the lowest figure of merit. It was only through consideration of the components of figure of merit – that is, hits, misses and false alarms – that the comparison of RF and LR simulations became clear. RF made fewer hits, more misses and fewer false alarms than LR. In fact, RF predicted less change overall than LR. Conversely, LR predicted more change than RF, and some of those excess simulated changes were actually correct, contributing to the higher hits of LR. In a similar fashion, some of the excess changes simulated by LR were actually incorrect, contributing to the higher false alarms of LR. In terms of total errors, RF performed better than LR, which is why RF simulations had higher overall accuracy. Finally, the error in simulation of quantity of change was smaller in RF than in LR. This means that the difference between misses and false alarms was smaller for RF. Overall, this example showed that single integrated metrics are not good indicators for assessment of models, because they combine the effects of several informative factors in one value, which could be interpreted in multiple ways. For example, a high overall accuracy could indicate high hits and low error, or it could simply indicate that the prevalence of change was small compared to the size of the study

area. Instead of such combined metrics, it is more useful for model assessment to use multiple simple metrics.

Another noteworthy point about Chapter 2 was that the similarity projections of the future spread of infestations with those made by the BC Ministry of Forests. In both cases, it was predicted that the provincial infestations will subside. On first thoughts, one might expect a decline in the population of the insect due to limited stock of host trees and food. It is true that resource limit imposes a limit to population and can eventually annule the rate of increase of a population, creating an S-curve. However, in projections of both Chapter 2 and the Ministry, spread of infestations nearly stops while there is still substantial amount of hosts available. In other words, not all available food is predicted to be consumed by the beetle population. Indeed in the simulations the infestations become less and less intense before they reach areas that are far from initial points of spread.

Chapter 3

The novelty of the work reported in Chapter 3 lies in the argument that in assessment of land change models, the places of occurrence of errors are important and should be taken into consideration. This was an addition to previously existing methods, which were based on counts of various types of errors. For the particular case of my study, which involved the spread of a phenomenon in space, I divided the study area into two partitions based on distance to locations of insects at the start of the study. Then, I showed that by looking at the distribution of errors between those two partitions of the study area it is possible to gain information that was otherwise hidden from the analyses of existing methods. In this way, I performed an error analysis taking into account where – that is, in what partition of the study area – errors occurred.

Another innovative feature of Chapter 3 was that it introduced a method for simultaneous comparison of four maps. In existing methods, assessment of land change models through comparison of three maps – reference at beginning and final times and simulation of the final time of the study interval – was a well-known validation methodology in land change science (National Research Council, 2014; Pontius Jr et al., 2011, 2004; Pontius Jr and Millones, 2011; Varga et al., 2019). Land change literature also mentions the use of baselines for comparison against simulations (van Vliet et al., 2016). In this sense, the baseline map is a fourth map to be used in model assessment. However, the way existing methods take baselines into account is that

they perform a three-map comparison with the baseline and references, then they perform another three-map comparison with the simulation and references, and finally they compare the results of the above two steps. In other words, the assessment in existing methods does not include all four maps in the same analysis. As a result, some information that is available in the maps is lost in aggregation. For example, although the existing methods report the count of correct predictions of change in the model and a similar count for the baseline, they do not indicate to what extent the above two counts overlap. The originality of Chapter 3 in this regard was that it offered a validation method of four-map comparison to answer questions such as the above. As shown with an example case in Chapter 3, these new methods help improve model assessment by providing useful information that was otherwise unattainable with previous methods.

Chapter 4

Chapter 4 integrated the works of previous chapters to build a coupled social-ecological model and answer the questions raised in the introduction of this thesis. In other words, the previous chapters served to make the analysis done in Chapter 4 possible. In this sense, findings and conclusions of Chapter 4 are the main parts of the conclusions of this thesis. The most important feature of this chapter was that it demonstrated with an example that is possible to create the motivation to save a vulnerable resource by encouraging self-interested users to perform costly actions towards resource protection, without applying force and without using financial incentives. Moreover, this finding was in spite of insufficient knowledge about the behavior of the ecological system in response to the intended intervention. In addition to answering questions about the particular case of MPB infestations, this study presented an overall approach to problems of its kind, that is, problems that involve ecological complexity and uncertainty on the one hand, and require social participation in costly actions on the other.

It is worth emphasizing that aside from the social complexities that affect a SES, the uncertainty about the ecological response to interventions is an important government challenge that Chapter 4 addressed with its “*Enforce scenario*”. This scenario assumed full cooperation of users with the government, and simulated complete implementation of the government’s intended intervention in the ecosystem. In these simulations, the interventions modified the landscape in each time-step, and those modifications affected the subsequent ecological changes

in the system. This is a tool that provides valuable insight to policy and decision makers about the potential consequences of their actions.

In Chapter 4, by building a model, an environment for virtual experiments was created. This is one of the functions of conceptual models, which allow to perform ‘what-if’ analyses and are thus employed in SES studies (Janssen and Ostrom, 2006). In this study a conceptual model was used for learning about the complexities of a SES and interventions in it. However, in this effort, it is important to keep in mind that a conceptual model does not make predictions about the real world. It cannot do so, because it runs in a hypothetical setting. Basically, it is not the objective of a conceptual model to produce precise predictions; rather, these models are built and used to help us better understand phenomena. In this sense, rather than predicting how a system will change, a conceptual model makes it possible to say that with the given assumptions, the simulated consequences are reasonable. This is how a conceptual model leads to insights about systems and phenomena.

In addition to building the coupled model, one of the significant features of Chapter 4 was the way it used the model to develop information that helps to answer questions of the study. This involved running the model multiple times under different conditions that set baselines for comparison against the main simulations. Defining easy-to-understand baselines made it possible to learn about the range of possible outcomes of the system, and to judge where the simulations stand in the space of possibilities. In addition, comparing the results of simulations with and without the RL algorithm resulted in an insight about the power and limits of the simulation model.

The coupled SES model of this study is a hybrid model composed of spatial and agent-based components. While ABM are well suited for the simulation of SESs, literature notes that one of the issues with these models in the study of governance of SESs is to represent governance dynamics in a way that is both adequate and simplified (Bourceret et al., 2021). Addressing this challenge is the modellers’ job. Based on model objectives, modellers choose a level of detail in which the model can simulate the envisioned complexity of the system, while being as simple as possible. This is the principle of parsimony in modelling (Batty and Torrens, 2005), which was central to the design of models of this study.

Remarks

The hypothetical setting described in Chapter 1 raises a question about the interactions among agents and promotion of the *responsible user* label. It might seem from the definition of the decision process of user agents that the value of the label increases automatically through iterations, because the more times the label is awarded, the more valueable it is deemed among user agents. Based on this reasoning, the emergence of the expected norm of responsible behavior might seem a trivial consequence of the way the setting has been defined. Indeed in Chapter 4 it was shown that even with random actions of the governing agent, still a considerable proportion of user agents may choose the label, given sufficient iterations. However, this doesn't mean that the promotion of the label is guaranteed by the definition of the model. To shed light on this matter, the process of decision making of the user agent should be viewed in detail.

As mentioned above, one of the factors by which user agents evaluate the label is the label's visibility, or the total number of times the label has been awarded. There is another factor that the user agents consider as well; that is the uniqueness that they will gain if they opt to have the label. Calculation of the value of the label for each agent includes the product of visibility and uniqueness scores. The answer to the point raised in this paragraph lies in the uniqueness score. For each user agent, this score is an inverse function of the perceived number of *other* user agents who will opt to have the label in the next time-step. In other words, it is highest if the user agent assesses that it will be the only user agent with the label in the next time-step. Conversely, the uniqueness score is lowest if the user agent assesses that all other user agents will have the label in the next time-step. In order to keep the model simple, the user agent assumes that of the number of user agents who opt for the label in the next time-step will be equal to the last known number of user agents who did so. As described in Chapters 1 and 4, this quantity is one of the model's variables, which is called *nLast*. In order to have the last known number of *other* user agents who received the label in an iteration, each user agent subtracts its own previous decision from *nLast*. Therefore, the formula for calculation of the uniqueness score is:

$$uniqueness = \frac{1}{1 + nLast - Decision}$$

where the 1 in the denominator is added to avoid division by zero.

If $nLast$ is 0, then each of the user agents will assess that receiving the label will make it highly unique. In this case, each user will calculate a uniqueness score of 1 for the label. Now, suppose exactly one of the user agents decides to have the label (in the model's terms, its decision is 1). In the next time-step, $nLast$ will be 1. The user agent who received the label previously, will still calculate uniqueness to be 1. However, for all other user agents (whose decisions were 0), the calculated score of uniqueness will be 0.5. This example shows how a change in two consecutive time-steps can affect many user agents in such a way that their assessment of the value of the label is substantially reduced.

Another point that the discussions of Chapters 1 and 4 highlighted was that even though the responsible behavior is adopted by some of the user agents in the end, a proper algorithm for the governing agent may lead to much earlier emergence of the responsible behavior among user agents. This was particularly shown via comparison of results of Reinforcement Learning models against random baselines. This matter is independent of the subject discussed in previous paragraphs of this section, because said the difference between simulations is due to the decision algorithms of the governing agent, whereas the described details about the value of the label are related to the decision function of the user agents.

The land change model of Chapter 2 is was trained with data of the whole province of British Columbia (BC). The reason for this decision was that during the study period, it was reasonable to assume that the infestations in the province were endogenous. This means that the province as a system evolved as a result of its own dynamics. In other words, to model the dynamics of the system at the province level, necessary information was available. In contrast, focusing on a smaller area would likely involve unknown incoming flux of insects from outside, which would pose a serious challenge to model calibration.

The model of Chapter 2 – which was used later on in Chapter 4 – was trained with data from the province of BC. In this sense, the province was a case study for the simulation methods of Chapter 2. Likewise, BC data was used for demonstrating the validation methods presented in Chapter 3. Hence BC was a case study for this chapter as well. BC data was also used in Chapter 4, albeit in a somewhat different sense. While Chapters 2 and 3 aimed at precise simulation and model assessment, the goal of Chapter 4 was to demonstrate a process of gaining insight about complexities of a system from a set of conceptual experiments. Unlike Chapter 2, which makes

solid claims about where infestations may go, Chapter 4 does not make any claim regarding precise prediction of where successive infestations spread nor does Chapter 4 prescribe a particular management technique for the control of infestations.

Notwithstanding the above, the hypothetical simulations of Chapter 4 are related to on the ground reality of BC. The province indeed tried to suppress the infestations by increasing its allowable annual cut – in other words, calling for increased harvesting (Forest Practices Board, 2009, 2007). In a few years, BC realized that the adopted suppression policy is ineffective, as infestations spread in large areas and killed massive amounts of pine trees. The province then continued its increased harvesting policy, though this time with a different objective: to sell the dead wood before it loses value (*idem*). This second objective, as well, happened to be hard to achieve. The volume of infested wood was too much to harvest in a short time. In fact, it was estimated that only salvage harvesting of infested areas by 2008 would take more than two decades, which is longer than the time dead trees keep their market value (Forest Practices Board, 2009). Another challenge for the local government was that logging companies were not always interested in harvesting in the locations that the government's infestation management policy indicated (Forest Practices Board, 2007). In this sense, it can be seen that the government faced two challenges of quantity and location of harvest. These two challenges are related to the management action introduced in the hypothetical simulations of Chapter 4: increasing cuts in areas identified by the governing agent. In reality, both the intensity and duration of the increased harvesting policy were smaller than the simulations of Chapter 4. In other words, the simulations of Chapter 4 shed light on what possibilities could arise if the government defined that policy with a larger allowable annual cut, and continued with that policy. As marked above, those simulations do not serve as actual predictions, but rather as sources of insight. In the example of Chapter 4, the simulations show that if the government enforced a policy of cutting neighborhoods of observed infestations, then more than half of the study area would have to be harvested, and less than half of the area's forests would remain healthy and intact. In other scenarios where the above said policy is not enforced, even smaller areas of healthy forest would remain. Decision makers can then use this information as insight and decide whether or not it is acceptable to lose more than half of the forest in pursuing this particular management action.

Challenges and opportunities for future research

Computational challenges

In the works of Chapter 4, I confronted challenges and limits that can be considered in future works. Most importantly, constraints on computational power restricted the size of the study area as well as the number of iterations for each run. In addition, tests, parameterization and verification of the model were time consuming and challenging. Another problem was that the communication between social and ecological models – which involved transfer of files between directories – was sometimes interrupted by Operating System (OS) processes. As a result of such interruptions, some messages between the two models would be read incompletely, which would change the calculations of all subsequent time-steps. In order to avoid these problems, I placed checks and controls in various parts of the code, allocated additional time to file transfer and all operations where the OS could cause a lag, and read through all outputs and logs of transferred messages between models after each simulation, making sure that there were no misreadings.

Other challenges and opportunities

One of the simplifications of the abstract model in Chapter 1 was the assumption of binary signals of the governing agent. This means that the governing agent's tool for learning about its social environment and guiding the user agents to the desired behavior was limited to only two actions. This simplification substantially facilitated the model development and analysis. However, it is as well a limit of the model. In practice, a government may find a variety of tools available, each with a different level of utility for the users. The social model of Chapter 1 can be improved to include government signals with more than 2 levels. That is, instead of only offering two options of 'costly' (difficult) and 'no cost' (easy) tasks, the improved governing agent may offer various other options with levels of cost (or difficulty) that are between the above two. In the ultimate case where there are infinitely many levels, the signal actually becomes a continuous numeric variable, which allows the model to integrate actions with any economic weight. This will ideally help the governing agent identify preferences of the user agents more quickly than the case of Chapter 1.

Another point to note about Chapter 1 is that, despite the assumption of no financial incentives, the model presented in this chapter has the potential to be used in problems where governments should agree on the amount and timing of incentives. This is an avenue for future

research work. User agents in the model calculate the utility of cooperating with the governing agent based on a number of factors. These factors might as well include financial incentives.

In Chapter 2 it was desirable to work at the scale of the whole province. However, such a large scale was accompanied by an added cost: it made the spatial model very computationally intensive, especially with its Random Forest algorithm, which was a novelty in the simulation of the studied insect infestations. This was one of the limits of the section of study that is reported in Chapter 2. An area for future works on this model can be to reduce its computational load. More experiments can be done with the model if higher computational power is available. Furthermore, future works can study the effect of new variables such as wind, precipitation and other climatic and topographic variables on the model's predictive power.

Another area for future work on the spatial model of Chapter 2 is to account for more complicated infestation states. Presently, the model only considers binary states. The available reference data, on the other hand, gives proportions of infested areas of grid cells. Therefore, the current version of the model loses some information when reading its input as binary. Future efforts can rebuild the model to work with multi-level or with continuous infestation states. This should allow richer analysis to be performed using the same input data.

The analyses of Chapter 3 are limited to land change processes with a single transition between two land classes. Therefore, the methods of Chapter 3 still need to be improved in order to be applicable to all land change applications. This can be one of the areas of future work on the methods of this chapter. Specifically, solutions should be found for problems with multiple land classes, as well as problems with bi-directional land transitions. Another aspect where the methods of Chapter 3 can be improved is the partitioning of the area. So far, only two strata are defined and analyzed in the method. By having a larger number of strata, more information may be extracted in the comparison of maps. Once this problem is solved, the new method will also be an improvement to the existing methods of multiple resolution map comparison.

In defining the scope of this study, I assumed that the government cannot successfully impose new regulations if the society is not prepared to cooperate with the government. As shown in Chapter 4, with an appropriate sequence of decisions and given sufficient time for preparing the society, it is possible to guide the society to a position where it cooperates with the government towards protection of a natural resource. After this stage, that is, beyond the scope of

the present study, other pathways may be taken, each of which deserves to be studied further. Specifically, it is possible that the emergent behavior becomes a norm that involves social sanction for non-compliance (Axelrod, 1986). Alternatively, the government may introduce regulations that enforce the newly emerged behavior (North, 1990; Savarimuthu and Cranefield, 2011). The present study helps better understand these pathways by shedding light on their preliminary stage. Future works can build on this understanding and answer questions about what happens after the emergence of the intended behavior, and how systems evolve.

Particularly, a subject of future works can be the simulation of a similar SES with the additional feature that norms of behavior are strengthened by social sanction – that is, once a norm is established, members of the society punish those who do not comply with the norm. As another particular example, a subject of future works can be the a SES simulation wherein an environmentally responsible behavior emerges and then the government introduces laws to impose that behavior. These two examples are similar in that they involve punishment of non-compliance with the emergent norm of environmentally responsible behavior. The difference between these two examples is that in one of them punishment is done by society members, and in the other by the government.

Final words

This doctoral work was an endeavor to understand the complexities of a SES and an aspect of governance for sustainable development. It involved me in multiple disciplines and introduced me to various challenges. Throughout this multipartite project I exercised the approach of breaking a problem into its components, simplifying and analyzing them, and combining the analyses under the umbrella of the main problem again. I also exercised the approach of modelling in the study of complex systems. I dealt with some of the challenges of conceptual models as well as empirically grounded models. Through this experience, I gained an insight about the power and problems of these approaches in the study of complex systems. In particular, I noted some improvements that are to be made in my approach and methods in future works. In the light of this understanding, I intend to continue working on the problems of sustainable development in a complex systems framework, building and testing models of complex systems, and analyzing data obtained by observation and simulation. I intend to work with improved methods to approach more complex problems, specifically, problems involving multiple resources with a network of interdependencies between the resources as well as a

variety of stakeholders. In future works, I will apply the knowledge and experience that I gained in my doctoral studies to answer questions about sustainability of such complex systems.

References

- Axelrod, R., 1986. An evolutionary approach to norms. *Am. Polit. Sci. Rev.* 80, 1095–1111.
- Batty, M., Torrens, P.M., 2005. Modelling and prediction in a complex world. *Futures* 37, 745–766. <https://doi.org/10.1016/j.futures.2004.11.003>
- Bone, C., Dragičević, S., 2010. Simulation and validation of a reinforcement learning agent-based model for multi-stakeholder forest management. *Comput. Environ. Urban Syst.* 34, 162–174. <https://doi.org/10.1016/j.compenvurbsys.2009.10.001>
- Bourceret, A., Amblard, L., Mathias, J.D., 2021. Governance in social-ecological agent-based models: A review. *Ecol. Soc.* 26. <https://doi.org/10.5751/ES-12440-260238>
- Delre, S.A., Jager, W., Bijmolt, T.H.A., Janssen, M.A., 2010. Will it spread or not? the effects of social influences and network topology on innovation diffusion. *J. Prod. Innov. Manag.* 27, 267–282. <https://doi.org/10.1111/j.1540-5885.2010.00714.x>
- Delre, S.A., Jager, W., Bijmolt, T.H.A., Janssen, M.A., 2007. Targeting and timing promotional activities: An agent-based model for the takeoff of new products. *J. Bus. Res.* 60, 826–835. <https://doi.org/10.1016/j.jbusres.2007.02.002>
- Janssen, M.A., Ostrom, E., 2006. Governing Social-Ecological Systems, in: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics* (Vol. 2). Elsevier B.V., pp. 1465–1509. [https://doi.org/10.1016/S1574-0021\(05\)02030-7](https://doi.org/10.1016/S1574-0021(05)02030-7)
- Kiesling, E., Günther, M., Stummer, C., Wakolbinger, L.M., 2012. Agent-based simulation of innovation diffusion: A review. *Cent. Eur. J. Oper. Res.* 20, 183–230. <https://doi.org/10.1007/s10100-011-0210-y>
- National Research Council, 2014. *Advancing Land Change Modeling: Opportunities and Research Requirements*. National Academies Press, Washington, D.C. <https://doi.org/10.17226/18385>

- North, D.C., 1990. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/CBO9780511808678>
- Ostrom, E., 2009. A General Framework for Analyzing Sustainability of Social-Ecological Systems. *Science* (80-.). 325, 419–422. <https://doi.org/10.1126/science.1172133>
- Pontius Jr, R.G., Huffaker, D., Denman, K., 2004. Useful techniques of validation for spatially explicit land-change models. *Ecol. Modell.* 179, 445–461. <https://doi.org/10.1016/j.ecolmodel.2004.05.010>
- Pontius Jr, R.G., Millones, M., 2011. Death to Kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *Int. J. Remote Sens.* 32, 4407–4429. <https://doi.org/10.1080/01431161.2011.552923>
- Pontius Jr, R.G., Peethambaram, S., Castella, J.C., 2011. Comparison of three maps at multiple resolutions: A case study of land change simulation in cho don district, Vietnam. *Ann. Assoc. Am. Geogr.* 101, 45–62. <https://doi.org/10.1080/00045608.2010.517742>
- Rand, W., Rust, R.T., 2011. Agent-based modeling in marketing: Guidelines for rigor. *Int. J. Res. Mark.* 28, 181–193. <https://doi.org/10.1016/j.ijresmar.2011.04.002>
- Savarimuthu, B.T.R., Cranefield, S., 2011. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent Grid Syst.* 7, 21–54. <https://doi.org/10.3233/MGS-2011-0167>
- van Vliet, J., Bregt, A.K., Brown, D.G., van Delden, H., Heckbert, S., Verburg, P.H., 2016. A review of current calibration and validation practices in land-change modeling. *Environ. Model. Softw.* 82, 174–182. <https://doi.org/10.1016/j.envsoft.2016.04.017>
- Varga, O.G., Pontius Jr, R.G., Singh, S.K., Szabó, S., 2019. Intensity Analysis and the Figure of Merit's components for assessment of a Cellular Automata – Markov simulation model. *Ecol. Indic.* 101, 933–942. <https://doi.org/10.1016/j.ecolind.2019.01.057>
- Yang, L., Wu, J., 2009. Scholar-participated governance as an alternative solution to the problem of collective action in social-ecological systems. *Ecol. Econ.* 68, 2412–2425. <https://doi.org/10.1016/j.ecolecon.2009.03.021>