Université de Montréal

Twitter and Social Bots:
An Analysis of the 2021 Canadian Election

Par
Gabrielle Desrosiers-Brisebois

Département de science politique
Faculté des arts et des sciences

Mémoire présenté en vue de l'obtention
du grade de maîtrise (M. Sc.) en science politique

Décembre 2022

Université de Montréal

Département de science politique, Faculté des arts et des sciences

Ce mémoire intitulé

Twitter and Social Bots:
An Analysis of the 2021 Canadian Election

Présenté par
Gabrielle Desrosiers-Brisebois

A été évalué par un jury composé des personnes suivantes

Patrick Fournier
Président-rapporteur

Jean-François Godbout
Directeur de recherche

Reihaneh Rabbany
Co-directrice de recherche

Ruth Dassonneville
Membre du jury

Résumé

Les médias sociaux sont désormais des outils de communication incontournables, notamment lors de campagnes électorales. La prévalence de l'utilisation de plateformes de communication en ligne suscite néanmoins des inquiétudes au sein des démocraties occidentales quant aux risques de manipulation des électeurs, notamment par le biais de robots sociaux. Les robots sociaux sont des comptes automatisés qui peuvent être utilisés pour produire ou amplifier le contenu en ligne tout en se faisant passer pour de réels utilisateurs. Certaines études, principalement axées sur le cas des États-Unis, ont analysé la propagation de contenus de désinformation par les robots sociaux en période électorale, alors que d'autres ont également examiné le rôle de l'affiliation partisane sur les comportements et les tactiques favorisées par les robots sociaux. Toutefois, la question à savoir si l'orientation partisane des robots sociaux a un impact sur la quantité de désinformation politique qu'ils propagent demeure sans réponse. Par conséquent, l'objectif principal de ce travail de recherche est de déterminer si des différences partisanes peuvent être observées dans (i) le nombre de robots sociaux actifs pendant la campagne électorale canadienne de 2021, (ii) leurs interactions avec les comptes réels, et (iii) la quantité de contenu de désinformation qu'ils ont propagé. Afin d'atteindre cet objectif de recherche, ce mémoire de maîtrise s'appuie sur un ensemble de données Twitter de plus de 11,3 millions de tweets en anglais provenant d'environ 1,1 million d'utilisateurs distincts, ainsi que sur divers modèles pour distinguer les comptes de robots sociaux des comptes humains, déterminer l'orientation partisane des utilisateurs et détecter le contenu de désinformation politique véhiculé. Les résultats de ces méthodes distinctes indiquent des différences limitées dans le comportement des robots sociaux lors des dernières élections fédérales. Il a tout de même été possible d'observer que les robots sociaux de tendance conservatrice étaient plus nombreux que leurs homologues de tendance libérale, mais que les robots sociaux d'orientation libérale étaient ceux qui ont interagi le plus avec les comptes authentiques par le biais de retweets et de réponses directes, et qui ont propagé le plus de contenu de désinformation.

Mots-clés : médias sociaux, comptes automatisés, désinformation, élections canadiennes de 2021, Twitter

Abstract

Social media have now become essential communication tools, including within the context of electoral campaigns. However, the prevalence of online communication platforms has raised concerns in Western democracies about the risks of voter manipulation, particularly through social bot accounts. Social bots are automated computer algorithms which can be used to produce or amplify online content while posing as authentic users. Some studies, mostly focused on the case of the United States, analyzed the propagation of disinformation content by social bots during electoral periods, while others have also examined the role of partisanship on social bots' behaviors and activities. However, the question of whether social bots' partisan-leaning impacts the amount of political disinformation content they generate online remains unanswered. Therefore, the main goal of this study is to determine whether partisan differences could be observed in (i) the number of active social bots during the 2021 Canadian election campaign, (ii) their interactions with humans, and (iii) the amount of disinformation content they propagated. In order to reach this research objective, this master's thesis relies on an original Twitter dataset of more than 11.3 million English tweets from roughly 1.1 million distinct users, as well as diverse models to distinguish between social bot and human accounts, determine the partisan-leaning of users, and detect political disinformation content. Based on these distinct methods, the results indicate limited differences in the behavior of social bots in the 2021 federal election. It was however possible to observe that conservative-leaning social bots were more numerous than their liberal-leaning counterparts, but liberal-leaning accounts were those who interacted more with authentic accounts through retweets and replies and shared the most disinformation content.

Keywords: Social media, Social bots, Disinformation, 2021 Canadian Election, Twitter

Table of Contents

Résumé i
Abstract ii
List of Figures v
List of Tables v
Remerciments vii
Introduction 1
Chapter 1. Background and Related Works 6
1.1. Defining Social Bots 6
1.1.1. What Are Social Bots? 7
1.1.2. Social Bots' Techniques 10
1.2. Social Bots and Electoral Campaigns 14
1.2.1 Involvement of Social Bots in Electoral Campaigns 15
1.2.2. Why Should We Be Concerned About Social Bots? 20
1.3. Social Bots' Partisan Behaviors 22
1.3.1. Social Bots' Presence 22
1.3.2. Social Bots' Activities 27
1.3.3. Social Bots' Disinformation 30
1.4. The Case of the 2021 Federal Canadian Election 34
1.5. Conclusion 39
Chapter 2. Methodology 40
2.1. Data 40
2.2. Bot Detection 41
2.3. Party Leaning Classification 48
2.3.1 Measuring Social Bots' Strategies and Interactions 56
2.4. Disinformation Detection 59
2.5. Conclusion 64
Chapter 3. Results 65
3.1. Botometer Results 65
3.2. Party Leaning Classification 68
3.3. Social Bots' Strategies During The 2021 Campaign 72
3.4. Interactions With Human Accounts 79

3.5. Disinformation Results 82
3.6. Conclusion 91
Chapter 4. Discussion and Conclusion 92
4.1. General Discussion 93
4.1.1. Limits and Future Works 94
References 99
Appendix 119

List of Figures

Figure 2.1. Bot Score Distribution 47
Figure 3.1. Example of Unrelated Tweet Found in the Dataset 77
Figure 3.2. Examples of Canadian Election Disinformation Tweets 86
Figure 3.3. Examples of COVID-19 Disinformation Tweets 87
Figure 3.4. Chrystia Freeland's Manipulated Tweet...90

List of Tables

Table 2.1. Summary of the 2021 Canadian Election English Dataset 40
Table 2.2. Summary of Profile-Based Classification 52
Table 2.3. Results From the Profile Classifier 53
Table 2.4. Results From the Activity Classifier 55
Table 3.1. Summary of Bot Scores' Availability 67
Table 3.2. Summary of Users and Tweets' Distribution 67
Table 3.3. Summary of Users' Distribution Per Partisan Group 69
Table 3.4. Summary of Tweets' Distribution Per Partisan Group 70
Table 3.5. Astroturfing Ratios for Humans 73
Table 3.6. Astroturfing Ratios for Social Bots 74
Table 3.7. Top 20 Hashtags Generated by Human and Social Bot Accounts 76
Table 3.8. Top 20 Mentions Generated by Human and Social Bot Accounts 78
Table 3.9. Social Bot Effectiveness in Human Engagement..80
Table 3.10. Summary of Tweets Per Topic 82
Table 3.11. Summary of Disinformation Tweets Per Topic 84
Table 3.12. Summary of Disinformation Ratios 84

Table A1. List of Keywords and Hashtags Used to Collect the Dataset.....................................119
Table A2. List of Keywords Used to Infer Partisan Leaning..119
Table A3. List of Keywords and Hashtags Used to Identify Disinformation 120
Table A4. List of Added Keywords and Hashtags to Identify COVID-19 Tweets......................121
Table B1. Total Number of Tweets Per Group 122
Table C1. Distribution of the Most Popular Keywords in the 2021 Canadian Election Disinformation Tweets by Humans...123
Table C2. Distribution of the Most Popular Keywords in the 2021 Canadian Election Disinformation Tweets by Social Bots...123
Table C3. Distribution of the Most Popular Keywords in the COVID-19 Disinformation Tweets by Humans..124
Table C4. Distribution of the Most Popular Keywords in the COVID-19 Disinformation Tweets by Social Bots..124

Table D1. Partisan Divisions with a Threshold of 0.3 125
Table D2. Partisan Divisions with a Threshold of 0.5 125

Remerciments

Je tiens tout d'abord à remercier mes directeurs de recherche, Professeur Godbout et Professeure Rabbany, sans qui la réalisation de ce mémoire n'aurait pas été possible. Je suis extrêmement reconnaissante pour le temps que vous m'avez accordé, ainsi que pour vos commentaires et suggestions qui m'ont permis de structurer mes idées et pousser mes réflexions. Merci également à MM. Jean-François Godbout et Martin Papillon de m'avoir offert l'opportunité d'être auxiliaire d'enseignement pour leurs cours au baccalauréat. Ceci a été une expérience enrichissante dans mon parcours.

Je me dois également de remercier les membres du *Complex Data Lab*, qui m'ont été d'une grande aide dans la construction de la méthodologie et qui ont toujours été disponibles et volontaires pour répondre à mes questions. Merci aussi à Aengus Bridgman qui a généreusement accepté de partager avec moi sa banque de données liée à l'élection fédérale canadienne de 2021, sur laquelle reposent les analyses effectuées dans ce mémoire.

Je souhaite souligner l'apport financier du Conseil de recherches en sciences humaines (CRSH) et de l'Institut québécois d'intelligence artificielle (MILA) qui m'ont permis de pleinement concentrer mes énergies sur mes études, ainsi que sur la rédaction de ce mémoire.

Enfin, un merci particulier à ma famille et mes amis qui ont été présents à chacune des étapes de mon parcours pour me soutenir, m'encourager et m'écouter parler de politique pendant de (trop) longues heures. Je suis très chanceuse de pouvoir compter sur des personnes aussi aimantes et dévouées dans ma vie, sans votre soutien la réalisation de ce mémoire n'aurait pas été possible!

Introduction

Social media have become essential tools in modern political communication, especially during electoral periods. Over time, these platforms transformed the way people communicate, acquire information, and run electoral campaigns (Dimitrova and Matthes 2018; Gorodnichenko, Pham and Talavera 2021). The unique structure of digital platforms allows political actors to promote their agenda without the intermediary of journalists and traditional news organizations, in addition to targeting specific segments of the electorate (Bossetta 2018). The 2016 U.S. presidential election, as well as the United Kingdom European Union membership referendum – also known as the Brexit referendum – that took place the same year, highlighted the importance of social media platforms in the electoral tactics deployed by political and social actors (Enli 2017; Hall, Tinati and Jennings 2018). As a result, the use of social media as a campaign tool raised concerns among journalists (e.g., Gebelhoff 2019; Bensinger 2020), academics (e.g., Kim et al. 2018; Garrett 2019) and politicians (e.g., Canada, House of Commons 2018; European Parliament 2021) alike. Their concerns are mainly related to the automation of political communication on social media and the potential impact of content conveyed by social bots, or automated software robots, on the quality and integrity of democratic discourse.

Technological advances in recent years have enabled the automation of political communication on social media. While beneficial in some respects, the automation of communication has nevertheless given way to the intrusion of malicious automated accounts into online discussions. Their involvement in online conversations about social issues has been documented repeatedly, notably in discussions around vaccines (Broniatowski et al. 2018; Yuan, Schuchard and Crooks 2019; Zhang et al. 2022), the COVID-19 pandemic (Ferrara 2020a; Uyheng and Carley 2020; Yang, Torres-Lugo and Menczer 2020), gun control (Ozer, Yildirim and Davulcu

2019; Schuchard et al. 2019), climate change (Al-Rawi, Kane and Bizimana 2021; Chen et al. 2021a; Marlow, Miller and Roberts 2021), and social movements like Black Lives Matter (des Mesnards et al. 2022; Jones, Nurse and Li 2022). The use of automated accounts has also been reported in a plethora of electoral campaigns around the world, such as in the U.S. (Bessi and Ferrara 2016; Deb et al. 2019; Luceri et al. 2019), the U.K. (Howard and Kollanyi 2016; Bastos and Mercea 2019), France (Ferrara 2017; Abdine et al. 2022), Germany (Brachten et al. 2017; Keller and Klinger 2019), Spain (Pastor-Galindo et al. 2020), and Canada (Rheault and Musulan 2021).

Extensive literature on social media now focuses on social bots. Bots can be briefly defined as automated computer algorithms used to produce or amplify online content while posing as real users. Although automated accounts can be found on all online platforms, Twitter represents a particularly fertile ground for them (de Lima Salge and Berente 2017). Despite Twitter's efforts to remove such accounts from its platform, their presence is estimated to be between 9 to 15% of all active English-speaking users (Varol et al. 2017). Nonetheless, such a large volume of automated accounts represents an important problem since researchers have identified them as key vectors of disinformation to manipulate public opinion, most notably in U.S. electoral campaigns (Bessi and Ferrara 2016; Kollanyi, Howard and Woolley 2016; Shao et al. 2018). A prime example of such coordinated mass manipulation of political information is the case of Russian interference during the 2016 U.S. presidential election opposing Donald Trump and Hillary Clinton. Researchers have studied this foreign interference (e.g., Badawy, Ferarra and Lerman 2018; Linvill et al. 2019), but it was also investigated in the *Report On The Investigation Into Russian Interference In The 2016 Presidential Election* (2019) conducted by special counsel Robert Mueller. This investigation has demonstrated, among other things, that the St. Petersburg-based Internet Research Agency (IRA) conducted a vast social media operation “designed to provoke and amplify political and social discord in the United States” by operating a “network of automated Twitter accounts […] that enabled the IRA to amplify existing content on Twitter” (2019, 4; 26).

Even if studies have shown the prevalence of automated accounts on social media platforms, their impact on electoral processes remains disputed among scholars. On the one hand, some researchers (e.g., Kushin and Yamamoto 2011; Murthy et al. 2016) argue that the content propagated by bots has a minimal persuasive effect, while on the other hand, some researchers (e.g., Ferrara 2020a; Pescetelli, Barkoczi and Cebrian 2022) argue that the unique structure of social networks and the sophistication of automated accounts could exert a significant influence on public opinion. A few scholars even suggested that automated accounts can influence the outcome of electoral campaigns. For example, Jamieson (2018) and Gunther, Beck, and Nisbet (2019) argued that fake Twitter accounts and disinformation propagated by them played a role in Donald Trump’s victory in 2016. Such claims have enormous implications for the integrity and transparency of the democratic process, but they are not entirely unrealistic. Studies have shown the influence of social media on political behaviors and attitudes, such as voter turnout (Bond et al. 2012; Jones et al. 2017), civic engagement (Boulianne 2015; Larson et al. 2019), and opinion formation (Messing and Westwood 2012; Barnidge, Gil de Zúñiga and Diehl 2017). Therefore, assuming that bots may influence users’ voting preferences is not incongruous.

Furthermore, U.S.-based studies (e.g., Deb et al. 2019; Chang et al. 2021; Chang and Ferrara 2022) identified differences in bots’ behaviors, attitudes, and interactions with authentic accounts according to their partisan leaning. However, whether bots influence election results remains challenging to answer, as we still do not know whether the partisan leaning of these automated accounts plays a role in the amount of disinformation content they propagate online. This is particularly true for Canada, where few studies have focused on bots' behaviors during election periods (e.g., Beskow and Carley 2020; Bellutta, King and Carley 2021; Rheault and Musulan 2021). In this context, examining automated accounts' behaviors is relevant through the following research questions:

i. *Did social bot activity during the 2021 Canadian federal election differ based on the partisan leaning of automated accounts?*
ii. *Can we find different strategies among partisan social bot groups during the 2021 electoral campaign?*
iii. *Did one partisan group spread more disinformation content than another?*

In an effort to answer these questions, I use Twitter data related to the 2021 Canadian federal election in order to (1) identify active automated accounts, (2) classify these accounts based on their partisan leaning, (3) analyze their behaviors and strategies by looking at the content of their tweets and retweets, and (4) detect disinformation content in their posts.

In other words, the objective of this study is to determine if, in the Canadian political context, the partisan leanings of social bots have the potential to influence their interactions with humans and the amount of disinformation they propagate. Previous studies have quantified the number of social bots active during election periods (Bessi and Ferrara 2016; Varol et al. 2017), analyzed the strategies favored by these accounts (Luceri et al. 2019; Ferrara 2020b) and even looked at their role in the propagation of disinformation content (Shao et al. 2018). Although not exclusive to the U.S., the literature on social bots is primarily focused on this country, mainly because of its global power and influence, the high frequency of elections, and the access to large databases in the English language. The conclusions derived from the American case can certainly guide theoretical positioning. Still, as shown by Brachten et al. (2017) in their study of the 2017 German Bundestag elections, the results obtained in other countries may be entirely contrary to what is found in the U.S. It is therefore relevant from a practical and theoretical point of view to explore the behaviors and strategies deployed by bots in the Canadian context, which is different from the U.S. in many respects. In knowing the tactics favored by social bots in the 2021 federal election, the Canadian government could, on the one hand, better understand the scope of the problem and, on the other hand, implement better strategies to address the practices of automated accounts.

This thesis will have the following structure. In the first chapter, I define the concept of social bots. I also present an overview of the various techniques employed by these accounts and their involvement in electoral campaigns. The second chapter introduces the methodology used to identify automated accounts, users' partisan leaning, and disinformation content. In the next chapter, I present and analyze the 2021 Canadian federal election Twitter dataset results. The final chapter addresses the implications of such findings for the Canadian context and reflects on the hypotheses developed in this thesis.

Chapter 1. Background and Related Works

As highlighted in the introduction, the growing presence of automated accounts and their engagement in electoral campaigns raise concerns about their impact on the quality of the democratic process. However, these types of digital accounts are diverse and multifaceted. This chapter aims to clarify what social bots are and understand how they have been used in different elections by considering the relevant studies on the subject. Therefore, this chapter is divided into four sections. The first section of this literature review focuses on defining the concept of social bots, their characteristics, and the techniques they employ during political events. The second section is dedicated to the literature on the impacts of social bots during electoral campaigns. The third part focuses on the partisan activities of social bots in online political discussions. Finally, the fourth section aims to provide context for the Canadian case and the 2021 election, which is the focus of this research.

1.1. Defining Social Bots

What are social bots? Although simple, this question has yet to reach a consensus among scholars, as research on this topic is relatively recent (Gorwa and Guilbeault 2020; Martini et al. 2021). Social bots' terminology is still not well-defined and remains somewhat ambiguous. This ambiguity mainly stems from several different terms referring to the same concept (Stieglitz et al. 2017). Nevertheless, some researchers have proposed definitions to understand what does and does not constitute a social bot. Therefore, the following section aims to present the different definitions and categorizations of social bots that scholars have put forward over the past years.

1.1.1. What Are Social Bots?

Bots, a shortened term that comes from “software robots”, are increasingly prevalent online. Their presence is so widespread that it is estimated that over 40% of the total volume of web activities originate from bots, meaning that a large portion of Internet traffic is not generated by humans (Imperva 2022, 7). The primary ability of bots is to automate tasks over the Internet. In return, this automation enables them to perform repetitive actions much faster and in greater quantities than humans can achieve. The extent of this phenomenon is mainly explained by the fact that bots are versatile, relatively easy and cheap to produce (Howard, Woolley and Calo 2018; Deb et al. 2019). For example, an increasing amount of open-source code is available on the Internet, especially on GitHub, the largest online code repository, for setting up and deploying bots (Kollanyi 2016).

Bots have been broadly defined and understood as computer-generated software programs which seek to perform automated tasks online (Hagen et al. 2022). However, automated software programs come in various forms and are used for different purposes. On the Internet, “good” and “bad” bots coexist, and their designation depends on several factors, in particular, the way in which they are used (Dunham and Melnick 2009). Among the different categories of bots, “web bots” and “chatbots” are considered benign. These bots have been present on Web platforms since the 1990s to automate various tasks but have different attributes. More precisely, web bots are mainly used to crawl, index and scrape web pages. With the rapid increase of web pages, this category of bots has become an essential component of search engines (Gorwa and Guilbeault 2020). These scripts are also used to automate information such as news or weather reports and commercial advertising (Stieglitz et al. 2017; Hagen et al. 2022). However, one of the main differences between web bots and other bots is that they are not intended to interact with users. In contrast, chatbots are automated programs designed to support a human-like dialogue through an online interface (Deryugina 2010). These automated dialogue programs are now commonly used in messaging applications like Facebook Messenger, WeChat, and Slack, as well as on companies' websites, to facilitate customer service (Xu et al. 2017; Gorwa and Guilbeault 2020). Therefore, the two categories of bots described above defy the common perception that all bots are designed to perform harmful actions. Indeed, web bots and chatbots are not considered harmful since they execute legitimate actions without the intention of deceiving the masses (Ratkiewicz et al. 2011; Orabi et al. 2020).

However, the rapid growth of the Web has also led to the appearance of malicious bots whose objective is to disrupt the stability of the Internet and fraud users. Spambots, whose sole purpose is to coordinate spam attacks, were one of the most common forms of malicious bots of the Web 1.0 (Gorwa and Guilbeault 2020). Spambots can be used to post online messages or to spread advertisements and malware in large volumes. For example, the attacks carried out mainly consisted of email spam campaigns, Web link farms, fake reviews, and Distributed Denial-of-Service (DDoS) attacks to overwhelm specific internet servers. However, coordinating these spam operations required human intervention (Geer 2005; Ferrara 2019). In other words, before the rise of the Web 2.0, hackers had to buy, host, and promote their Internet domains, which required human operators and resulted in substantial costs (Hayati et al. 2009; Ferrara 2019).

The emergence of social media platforms such as Facebook and Twitter in the mid-2000s provided malicious bots with new grounds to exploit. As a result, new forms of online automation have emerged in the Web 2.0. To describe these patterns of automation specific to social media platforms, the term "social bot" was adopted by many researchers (e.g., Wagner et al. 2012; Bessi and Ferrara 2016; Shao et al. 2018). As opposed to the previous categories of bots, social bots are automated social actors that attempt to mimic human behavior (Abokhodair, Yoo and McDonald 2015). Other researchers consider this definition of social bots too broad as it does not allow for adequate distinction between different types of automated accounts. For instance, according to Abokhodair, Yoo, and McDonald's (2015) definition, chatbots, which also aim to mimic the online responses of humans, would then be considered social bots. Therefore, Ferrara and his colleagues (2016) have argued that social bots are a distinct form of automated accounts, because they are "computer algorithm[s] that automatically produce content and interact with humans on social media, trying to emulate and possibly alter their behavior" (96). Based on this definition, social bots are characterized, on the one hand, by their imitation of human behavior, and, on the other hand, by a desire to influence social media users. Woolley (2016) also used the term social bot in reference to "software programs designed to mimic human social media users on platforms", but he particularly insists that they are deployed "to manipulate public opinion and disrupt organizational communication" (1). Similarly, in their review of the literature published in 2017, Stieglitz et al. determined that social bots are characterized by their high level of malicious intent and imitation of human behavior. In this sense, the objective of imitating human behavior is to go undetected in order to interact with real users and eventually influence them on a range of issues. This implies that social bots can automatically post, repost, like or reply to messages of their own accord on different social media platforms (Orabi et al. 2020).

The term *Sybil* has also been identified in the literature (e.g., Paradise, Puzis and Shabtai 2014; Goga, Venkatadri and Gummadi 2015; Davis et al. 2016) to refer specifically to automated accounts that operate under fake identities. Stieglitz et al. (2017) noted that the terms social bot and Sybil are not distinguishable and are often used interchangeably. Some researchers also use terms other than social bot in reference to essentially the same phenomenon. For example, in their study of the role that automated Twitter accounts played in Venezuelan political conversations, Forelle et al. (2015) use the term political bots to refer to a subset of social bots designed to accomplish political tasks that vary across different regime types. Woolley (2016) and Woolley and Howard (2016) also mention that political bots are by-products of social bots solely designed to interfere in political contexts, such as elections, crises, and conflicts. Political bots can be used by government entities, political campaigns, or civic activists (Woolley and Howard 2016, 4885). The term Twitter bot has also been used alternatively in analyses to refer to political bots active on Twitter (Howard and Kollanyi 2016; Bastos and Mercea 2019). However, once again, the terminology surrounding these different terms could be clearer and more precise since Twitter bot has also been used to refer to automated accounts involved in vaccine debates in the U.S. (Broniatowski et al. 2018).

The purpose of this section was to briefly outline the evolution of the different types of automated accounts, in addition to defining the unique characteristics of social bots. As pointed out, social bots' terminology is still not well defined. Thus, different terms are used in reference to the same concept. To clarify this terminological ambiguity, this thesis will employ the term social bot, defined by Ferrara et al. (2016, 96) as a *computer algorithm that automatically produces content and interacts with humans on social media, trying to emulate and possibly alter their behavior*.

1.1.2. Social Bots' Techniques

As indicated earlier, social bot accounts are now increasingly used in the context of online political activities. The reasons behind using social bots are numerous, and the objectives they seek to pursue are just as varied. For instance, studies showed that social bots have been used to amplify the visibility and popularity of political figures (Murthy et al. 2016), manipulate public opinion (Gorodnichenko, Pham and Talavera 2021), and even disrupt political discourse on social media platforms (Woolley 2016). These objectives are achieved through various large-scale techniques. As such, three main approaches used by social bots to interfere in online political discussions have been identified in the literature.

The first one is called *astroturfing*. Astroturfing is a centrally organized top-down strategy aiming to deceive Internet users by emulating bottom-up activities (Kovic et al. 2018; Dubois et McKelvey 2019; Keller et al. 2020). In other words, social bots who employ an astroturfing approach seek to simulate artificial grassroots movement by amplifying messages on social media platforms. By pretending to be authentic users, social bots involved in astroturfing campaigns can create a false impression of consensus around a political position or a specific candidate. Signs of astroturfing have been detected in several studies. In their study of the run-up to the 2010 U.S. midterm elections, Ratkiewicz and his colleagues (2011) were among the first researchers to use a machine learning approach based on users' network relations (i.e., retweets, replies, and mentions) to observe astroturfing patterns in political discussions. Their study found that a network of automated accounts was actively evolved in astroturfing practices, which were most notably deployed against the Democratic candidate for U.S. Senate for Delaware, Chris Coons. Similarly, Bessi and Ferrara (2016) also found signs of astroturfing during the 2016 U.S. presidential election. By performing sentiment analysis on tweets containing hashtags related to Donald Trump or Hillary Clinton, the authors demonstrated that tweets produced by social bots supporting Donald Trump had almost no negative sentiment toward the Republican candidate. Therefore, social bots posted significantly more positive tweets about Donald Trump during the campaign, thereby fomenting an artificial impression of overwhelming support for this candidate (Bessi and Ferrara 2016, 8). Astroturfing has also been detected in election campaigns outside the U.S. For example, Keller et al. (2020) confirmed, via network activities analysis, that the South Korean National

Intelligence Service (NIS) resorted to astroturfing to boost support for the conservative presidential candidate Park Geun-hye. Social bots were also involved in astroturfing practices in recent European electoral campaigns. For instance, social bots were particularly active in the days before the Brexit vote took place on June 23rd, 2016, and mostly amplified messages supporting the Vote Leave campaign (Bastos and Mercea 2019). In his analysis of the 2017 French presidential election, Ferrara (2017) demonstrated that social bots actively participated in the amplification of the *MacronLeaks* disinformation campaign.[1] It is essential to mention that social bots are not only used during election campaigns; political actors can also use them to lead discussions about the government. One of the most striking examples is the case of China, where the governmental authorities practice a permanent astroturfing strategy in order to control public opinion without resorting to censorship as much as they might otherwise (King, Pan and Roberts 2017). It was documented that astroturfing is mainly used in China to dilute negative comments regarding the regime (Miller 2016). By using a dataset of more than 43,000 posts from known astroturfer accounts from the Jiangxi province, King, Pan and Roberts (2017) found that inauthentic accounts were used to propagate pro-regime and anti-Western sentiment online. Finally, it is important to remember that astroturfing strategies have been used in many electoral campaigns worldwide over the last few years (Brachten et al. 2017; Stieglitz et al. 2017). By producing large amounts of content with hashtags and keywords related to a particular candidate or political issue, social bots attempt to build and simulate a form of online consensus to sway public opinion.

The second approach is called *smoke screening*. This same approach is sometimes referred to as Twitter bombing when talking about activities specifically occurring on the Twitter platform.

[1] The *MacronLeaks* refer to the forgery of documents and the hack of more than 20,000 emails from Emmanuel Macron's campaign, which were leaked on social media platforms on May 5th, 2017 – only two days before the second round of the presidential election – in order to undermine the candidacy of the En Marche! candidate (Jeangène Vilmer 2019, 3).

This strategy aims to disrupt online debates by disseminating irrelevant messages or hashtags (Hasert and Hooffacker 2019). Hence, smoke screening consists of covering up tweets related to a particular topic by flooding hashtags with unrelated but similar content (e.g., by using the hashtag #elxn44, which is related to the election campaign, in order to overwhelm the platform with a narrative in favor of one of the party leaders). In other words, this strategy utilizes popular hashtags to steer users' attention away from certain topics (Brachten et al. 2017; Stieglitz and Brachten 2018; Marcellino et al. 2020). The most important example of smoke screening used in a political context occurred during the civil conflicts related to the Arab Spring. During the Arab Spring movements, governments such as Morocco, Iran, Bahrain, and Syria "hijacked" trending hashtags to counter pro-revolution narratives (York 2011). For example, in their study of a group of 130 active automated accounts related to the Syrian Civil War, from April to December 2012, Abokhodair, Yoo and McDonald (2015) discovered that the #Syria hashtag was taken by pro-regime forces who overwhelmed the pro-revolution narrative conveyed through this hashtag. Evidence of smoke screening was also found during the 2010 Massachusetts Special Election called to fill the seat vacated by the death of Senator Ted Kennedy. Automated accounts were created to respond to tweets containing general information about the Senate election with spam attacks against Democrat nominee Martha Coakley (Metaxas and Mustafaraj 2010; Metaxas and Mustafaraj 2012). In this case, smoke screening was used to overload generic election-related keywords with harmful content about the Democratic Party candidate.

The third approach social bots use to interfere in online political discourses is *misdirecting*. Misdirecting and smoke screening are closely related since both methods seek to influence public opinion by manipulating their perception of political issues. These two approaches work in the same way, with the only difference being that smoke screening uses similar content to a hashtag posted, whereas misdirecting does not (Schindler, Opuszko and Stöbesand 2021). Therefore, misdirecting refers to using context-related hashtags to spread messages unrelated to the hashtag used (e.g., by using the hashtag #elxn44, but talking about something completely unrelated to the 2021 Canadian federal election). Simply put, misdirecting is employed when social bots use contextual hashtags, but a completely different topic is being reported (Stieglitz and Brachten 2018; Schindler, Opuszko and Stöbesand 2021). With this strategy, social bots massively post messages unrelated to the hashtags to which they refer to guide the public's attention away from a topic and therefore make other issues the subject of discussion (Brachten et al. 2017). This strategy was also discovered in the same group of social bots tweeting about the Syrian Civil War. For instance, social bots used the hashtag #Syria to post messages about cinema and movies to hide content related to the Syrian regime's war activities (Abokhodair, Yoo and McDonald 2015). Lastly, it has been shown that social bots can adopt various strategies to interfere in online political discussions; however, the tactics they employ depend on the objectives they pursue. As such, astroturfing is privileged by malicious actors who try to increase the visibility of a candidate, or create an impression of consensus around an opinion, whereas smoke screening and misdirecting are favored strategies to deflect users' attention toward other issues. While the implications of these approaches are different, their potential influence raises concerns, as will be discussed further in the following section.

1.2. Social Bots and Electoral Campaigns

The rise of social bots as a communication tool is now a global phenomenon, which has been detected in several democratic elections (Woolley 2016). The next sections review the involvement of social bots in various electoral campaigns and their potential destabilizing effect on democracies.

1.2.1 Involvement of Social Bots in Electoral Campaigns

Although social bots became a more common term to the public amid the 2016 presidential U.S. election, their usage in political contexts can be traced back several years. Indeed, one of the first papers to focus on the activities of social bots during electoral campaigns is Ratkiewicz et al. (2011), who studied the 2010 U.S. midterm elections. With a corpus of approximately 600,000 tweets containing political keywords, the authors performed analyses based on tweets' content and users' network connections via machine learning algorithms to examine the dissemination of political information on online social networks. They identified a network of social bots heavily active in discussions surrounding the 2010 midterm elections. Despite a relatively small sample, two trends nonetheless emerged: (1) social bots showed signs of hyperpartisan behaviors (i.e., supporting candidates while attacking others), and (2) social bots were involved in sharing thousands of URLs leading to disinformation websites.

A 2018 study by Shao and his colleagues compiled a large dataset of more than 14 million tweets from mid-May 2016 to the end of March 2017, which covered the 2016 U.S. presidential election period. In that study, Shao et al. (2018) showed that social bots played a central role in trying to manipulate online public opinion. By looking at all the articles in their corpus and cross-checking their validity with reputable third-party news and fact-checking organizations, the authors found that social bots' tweets were overrepresented in propagating unreliable information. In other words, despite the fact that social bots represented about 6% of the accounts in the sample, they were responsible for spreading 31% of all tweets and 34% of all the articles linked to low-credibility sources (Shao et al. 2018, 3). This illustrates that a small number of accounts were responsible for amplifying disinformation narratives. Equally important, the retweet network of the 227,363 users who retweeted messages with links to low-credibility articles demonstrated that humans are particularly vulnerable to manipulation narratives which contain disinformation as they retweeted social bots who post low-credibility content almost as much as they retweeted other authentic accounts.

In another example related to the 2016 U.S. presidential election, Bessi and Ferrara (2016) conducted a study on the behavior of social bots based on 20 million tweets from approximately 2.8 million unique users. This large dataset was collected from a list of pertinent keywords and hashtags associated with the election and the presidential candidates. They estimated that roughly 400,000 accounts were social bots and that these accounts generated close to one-fifth of the entire conversation. As highlighted previously, the authors showed that social bots participated in astroturfing tactics in order to amplify positive sentiment surrounding Donald Trump. In addition, sentiment analysis performed via the *SentiStrength* algorithm showed that most tweets produced by automated accounts and directed at Hillary Clinton were negative. The analysis of Bessi and Ferrara (2016) also showed that to promote a candidate, social bots can simultaneously amplify positive messages about their candidate while promoting negative messages about their opponent. A study by Stella, Ferrara and De Domenico (2018) reported that social bots increased users' exposure to negative and inflammatory narratives. In this study, Stella, Ferrara and De Domenico (2018) collected a dataset of almost 4 million tweets from around 1 million distinct users who posted messages related to the 2017 Catalan independence referendum. Based on this data, they conducted a sentiment and a semantic network analysis, which revealed that social bots targeted the accounts of influential independence supporters with negative content. Therefore, social bots can be utilized as communication tools to exacerbate tensions between political groups and promote negative sentiment aimed at opponents.

Several social bot studies focused on the 2016 Brexit campaign to determine the impact of automated accounts on this referendum. Using an original dataset of 313,832 unique users and more than 1.5 million tweets selected on pro-Leave, pro-Remain, and neutral hashtags, Howard and Kollanyi (2016) analyzed social bots' activities in online Brexit discussions. Using frequency and clustering techniques based on hashtag use, the authors found that pro-Leave social bots were much more present than pro-Remain ones and generated significantly more tweets. Moreover, the referendum campaign was subject to heavy automation as seven of the top ten accounts that tweeted the most about Brexit were identified as social bots. Such a high level of automation around a polarizing political issue like a referendum might emphasize the visibility of certain narratives and therefore affect the distribution of content production across users' networks (Howard and Kollanyi 2016). Following the publication of Howard and Kollanyi's article, Bastos and Mercea (2019) decided to examine the retweet cascades of social bots during the U.K. referendum. The authors also relied on tweets collected from a list of pertinent hashtags associated with the campaign. However, their dataset is larger than Howard and Kollanyi's (2016), with approximately 10 million tweets and more than 800,000 distinct users. From this data, they could map the structure of users' networks and examine their activities. On the one hand, their results are consistent with those of Howard and Kollanyi (2016), as they also found that social bots mainly spread messages favoring the pro-Leave campaign. On the other hand, they also showed that automated accounts served as false amplifiers "by aggregating and retweeting content tweeted by seed users" (Bastos and Mercea 2019, 51). Unlike the case of the 2016 U.S. election, the authors did not find evidence that social bots were involved in a widespread disinformation campaign regarding Brexit. Furthermore, when looking at the diffusion of information on Twitter during both the Brexit referendum and the 2016 U.S. presidential election, Gorodnichenko, Phamb and Talavera (2021) demonstrated that linked-minded users (i.e., pro-Leave, pro-Remain, liberal or conservative) were more likely to interact with each other. As a result, social bots' ability to interact with real users on Twitter depends on whether social bots' information is consistent with humans' partisan or ideological preferences. The authors, therefore, mention that social bots have the potential to exacerbate online polarization by facilitating "echo chambers" and reinforcing humans' preexisting beliefs about political candidates or issues.

Numerous papers proved that politicians used social bots to boost their online popularity artificially. In his 2016 article, Woolley identified prime examples of social bots being used to bolster candidates' popularity and relevance. Woolley (2016) performed a qualitative content analysis on a corpus of new articles focusing on social bot usage in more than a dozen countries. He found that boosting politicians' follower numbers is mainly common in Western states. In 2012, for instance, Lee Jasper, a U.K. candidate for the Respect Party, resorted to the use of automated accounts in order to give an "impression of the popularity of his campaign" (Downes 2012). The same tactics were also found in the campaign of the 2012 U.S. presidential candidate, Mitt Romney. Indeed, over a period of only 24 hours, the Republican nominee gained roughly 117,000 new Twitter followers, representing an increase of approximately 17% (Coldewey 2012). After analyzing these newly acquired followers, researchers concluded that the vast majority were fake and did not come from an organic growth of followers (Coldewey 2012). During the 2013 federal election in Australia, two independent social media analyses revealed that roughly 40% of the 50,000 most recent followers of the four most popular Australian politicians on Twitter (Julia Gillard, Tony Abbott, Kevin Rudd, and Malcolm Turnbull) were not authentic accounts (Butt and Hounslow 2013). A more recent example of such practices was also found during the 2016 U.S. presidential election when automated accounts were created to impersonate Latino voters who followed and tweeted in support of Donald Trump (Howard, Woolley and Calo 2018).[2] One hypothesis is that these accounts were put together to give a sense that the Republican presidential candidate had good support among minority communities and understood them (Howard, Woolley and Calo 2018). Some might argue that artificially increasing the online popularity of political actors is a benign practice that only enhances candidates' visibility. However, adding tens of thousands of new followers can make an account more trustworthy and influential among genuine users and help attract a new crowd of real users (Cresci et al. 2015). Therefore, this usage of social bots can mislead users and create a false impression of popularity towards a candidate or a campaign.

In summary, social bots tend to be deployed during sensitive political moments such as elections or referendums (Howard and Kollanyi 2016). However, their increasing presence on social media platforms is problematic, especially since social bots' strategies aim to manipulate public perceptions and interfere in political discussions. Their impact on political campaigns is varied and can take many forms. As pointed out previously, social bots can try to influence discourses around a campaign or a candidate by posting negative messages (Ratkiewicz et al. 2011; Bessi and Ferrara 2016; Stella et al. 2018), propagating disinformation stories (Ratkiewicz et al. 2011; Shao et al. 2018), fostering divisive narratives to disrupt political communication (Gorodnichenko, Phamb and Talavera 2021), amplifying the visibility of a campaign (Howard and Kollanyi 2016; Bastos and Mercea 2019), or artificially boosting a candidate's popularity (Coldewey 2012; Downes 2012; Howard, Woolley and Calo 2018). These findings lead us to further explore the potential consequences of social bots on democracies.

[2] For example, social bot accounts with names such as Pepe Luis Lopez, Francisco Palma, and Alberto Contreras were designed to impersonate Latino supporters of Donald Trump (Howard, Woolley and Calo 2018, 81).

1.2.2. Why Should We Be Concerned About Social Bots?

Even though it is not yet possible to establish a link between the manipulative strategies of social bots and their influence on social media users, the fact remains that automated accounts and their growing sophistication represent a risk to the democratic process, both in theory and in practice. Several researchers voiced their concerns about the political use of social bots on social media platforms (e.g., Bessi and Ferrara 2016; Ross et al. 2019; Stella et al. 2018; Cantini et al. 2022). As pointed out by Ross et al. (2019), there is no limit to the number of social bots that can be deployed to flood social media platforms with directed messages to reinforce the perception of a widely held opinion. Therefore, this potential for large-scale manipulation of public opinion through the automation of political communication is worrying. Indeed, distorting the reality of social media users by giving a false impression of consensus around a candidate or a political issue could ultimately truncate voters' decision-making ability (Bessi and Ferrara 2016; Yang et al. 2019a). This is of particular concern, given the increasing sophistication of social bots, and their improved ability to mimic human behavior (Pozzana and Ferrara 2020). For instance, a survey conducted by the *Pew Research Center* (n = 4,581) revealed that more than half of the respondents who were aware of the presence of social bots online stated they did not feel confident they could differentiate between real and inauthentic accounts (Stocking and Sumida 2018). Guilbeault and Woolley (2016) and Ross and his colleagues (2019) used the concept of the "spiral of silence" to inform against the negative impacts of social bots on democracy. By adopting astroturfing strategies, which aim to amplify the visibility of certain content to the detriment of others, social bots stifle certain discourses, thus preventing people or groups from expressing their opinions around social or political issues. This spiral of silence can lead to less discussion and diversity in politics, which is harmful to liberal democracies based on pluralism (Plattner 2010).

The implication of social bots in sharing disinformation stories also poses a severe risk to the quality of political debates. Mass manipulation of public opinion through disinformation narratives is of particular concern as false information is spread "farther, faster, deeper, and more broadly" than authentic information on social media (Vosoughi, Roy and Aral 2018, 1150). Moreover, by increasing the exposure of social media users to negative political content and disinformation narratives, social bots could therefore play a role in accelerating polarization levels in democracies (Stella et al. 2018; Cantini et al. 2022). Using a synthetic experiment calibrated to Twitter to simulate information exchanges inside social media networks, Azzimonti and Fernandes (2022) discovered that significant disinformation and polarization arose among networks in which only a minority of 15% of accounts included in the experiment believed false political news. Disinformation propagated by automated accounts can prevent the aggregation of accurate information and consensus among different groups in the population, which can result in increasing levels of polarization in democratic societies. The specific involvement of social bots in the propagation of disinformation stories in the context of election campaigns represents a serious risk to the quality of political discourse. It has the potential to damage social cohesion by amplifying polarization levels. It is also important to specify that the consequences of the diffusion of erroneous content by social bots exceed the scope of politics. For example, many scholars (e.g., Allem and Ferrara 2018; Himelein-Wachowiak et al. 2021) have raised concerns about the impact of social bots on public health safety, particularly regarding vaccines. Allem and Ferrara (2018) highlighted the potential of social bots to "drown out medically sound social media messages from medical experts or health campaigns" (1006). This is of particular concern during a pandemic, where vaccine hesitancy has been documented, for example, in the United States (Fridman, Gershon and Gneezy 2021) and Canada (Lavoie et al. 2021). It is, therefore, essential to deepen our understanding of the behaviors of social bots, especially concerning the dissemination of inaccurate content, since such content can undermine a political community's "capacity to engage in communication characterized by the use of facts and logic, moral respect, and democratic inclusion" (McKay and Tenove 2021, 703).

1.3. Social Bots' Partisan Behaviors

The following section presents relevant studies examining partisan differences regarding social bots. These partisan distinctions address their respective quantity, their interactions with humans, and the propagation of disinformation content on social media. Based on the results presented in these studies, three research hypotheses are put forward to guide the rest of this research.

1.3.1. Social Bots' Presence

The proliferation of social bots on social media platforms during political events has led researchers to examine their partisan affiliation. While it is essential to recognize that social media users are not representative of the general population (Mellon and Prosser 2017), it is nevertheless relevant to understand how social bots of different partisan leanings behave on these platforms and how they interact with authentic accounts. Before continuing, we must distinguish between the concepts of partisanship and ideology, which are often amalgamated. On the one hand, partisanship refers to siding with a political party. For example, in the United States, individuals can identify as Republicans or Democrats. It is also possible for some individuals to not side with a particular party and be Independent (Petrocik 2009). On the other hand, ideology refers to a set of preferences that guide the positioning of individuals on a range of issues (Marietta 2012). In most Western states, where the main political cleavage is between the economic left and right, individuals position themselves on an ideological spectrum ranging from liberal to conservative (Farneti 2012). However, heightened levels of elite polarization in the United States now mean that elected Republicans and Democrats are more ideologically divergent from each other than in previous decades (Hare and Poole 2014). The increased ideological cohesion of American political parties resulted in partisan sorting, which strengthened the relationship between citizens' ideological and partisan identifications (Abramowitz 2010; Lupton, Smallpage and Enders 2020). However, this alignment between partisanship and ideology in the U.S. two-party system cannot be generalized to all Western states, most of which operate in multi-partisan systems. To avoid terminological confusion, the terms "conservative-leaning" and "liberal-leaning" will be used in this research. These two terms are more flexible and serve as shortcuts to infer the ideology of individuals without assuming their partisan identity (van Ditmars 2022).

In recent years, studies have focused on the impact of the partisan orientation of social bots on their online behavior and activities. Social bots of different partisan leanings were found to be active in election campaigns in various Western democracies (see Deb et al. 2019; Pastor-Galindo et al. 2020). However, studies focusing on identifying social bots and assigning partisan affiliation to Twitter accounts demonstrated that conservative-leaning social bots were more numerous than their liberal counterparts. For instance, Bessi and Ferrara's (2016) study discussed earlier found that social bots were more present among Republican supporters than Democrat supporters during the 2016 U.S. presidential election. In this study, the authors were able to attribute a partisan affiliation to a sample of more than 24,000 Twitter accounts by looking at the most frequently used hashtags in each user's tweets.[3] With this technique, Bessi and Ferrara (2016) showed that social bots made up a little over 12% of the sample of Republican supporters, while they made up roughly 9% of the sample of Democrat supporters. In their study of the 2020 U.S. presidential election, Chang et al. (2021) adopted a similar strategy for classifying users according to their partisan

[3] To identify Republican supporters, the hashtags used were: #donaldtrump, #trump2016, #neverhillary, #trumppence16, and #trump. To identify Democrat supporters, the hashtags used were: #hillaryclinton, #imwithher, #nevertrump, and #hillary (Bessi and Ferrara 2016, 7).

leaning. As such, they used campaign-related hashtags to discriminate between users who favored left or right-leaning political discourse. In this study, 85,000 social bots were found in the sample of over 8.5 million right-leaning users, whereas 18,000 social bots were present in the sample of more than 2.5 million left-leaning users. Once again, these results indicate that conservative-leaning social bots represented a larger proportion of automated accounts involved in recent U.S. elections.

This situation has also been observed in the Canadian context. For example, Rheault and Musulan (2021) leveraged a dataset of approximately 1.7 million users who tweeted about the 2019 Canadian federal election. Instead of focusing on tweets' content to infer users' party affiliation, Rheault and Musulan (2021) relied on unsupervised learning methods. In other words, the authors used the UMAP clustering technique to reveal the different partisan subgroups present in their dataset. This approach was based on the premise that users who are ideologically close to each other behave more homogeneously (Barberá 2015). Rheault and Musulan (2021) showed that the bot density percentage for the Conservative cluster was the highest. Furthermore, they demonstrated that the People's Party of Canada (PPC) – a minor libertarian political party – experienced a disproportionate level of social bot density. It is, however, pertinent to mention that the authors only used a set of 505 candidates with known party affiliations to validate the accuracy of their predictive model. Political candidates have been proven easier to classify (Conover et al. 2011). Thus, exclusively relying on this type of data as a validation measure could produce different results when used on data from the general public.

Identifying who is behind coordinated social bots' campaigns is generally impossible to determine. Nevertheless, some researchers have proposed hypotheses to explain why right-leaning social bots seem to be more prevalent during electoral campaigns. Frost (2020) suggested that the activism of automated accounts on social media platforms is most pronounced among conservative-leaning groups, from which populist politics emerge non-exclusively, but for the most part. Populism, which is rooted in a distrust of political elites and expressed as the true voice of the people and social bot activism have complementary aims (Müller 2016). Indeed, Frost (2020) emphasized that "populism is easily set off by bot-based strategies that ventriloquize 'true voice', and bots, in turn, rely on heightened content for their circulation, so a natural synergy emerges" (9). In other words, since populist parties claim to represent the "silent majority" as the basis of their legitimacy and main appeal, they have a considerable incentive to automate their communication to give an impression of a popular grassroots movement (Mudde 2019; Silva and Proksch 2021).

In line with this argument, conservative-leaning social bots' enablers should be more numerous on social media platforms because they represent a communication tool better suited to them to exploit and amplify the "true voice of the people". Hence, right-leaning social bots' enablers should have higher incentives than their left-leaning counterparts when it comes to using social bots in political campaigns. Various studies conducted in France (Ferrara 2017), the United States (Onuchowska, Berndt and Samtani 2019), as well as at the European Union level (Silva and Proksch 2021) showed that conservative-leaning parties benefited more from the interventions of social bots than other party families. For example, it was estimated that EU MPs from radical right-wing parties had inflated their followers by up to 5%, particularly among the most popular and louder anti-EU politicians (Silva and Proksch 2021, 321).

Moreover, studies conducted by Wojcik and Hughes (2019) and Freelon (2019) showed a partisan imbalance among U.S. Twitter users. On the one hand, by surveying a representative sample of 2,791 individuals who shared their Twitter handles, Wojcik and Hughes (2019) found that U.S. Twitter users are more likely to identify as Democrats, be more educated, younger, and have higher incomes than the overall U.S. population. On the other hand, Freelon (2019) established that individuals on the right side of the ideological spectrum in the U.S. represented only one-third of the platform's users. In addition to containing a larger share of Democrats than Republicans, an analysis led by the *Pew Research Center* (Shah et al. 2020) also showed that Democrats make up the majority of active tweeters on the platform and that the most active Democrats tweet more often than their most active Republican counterparts. For instance, the median Democrat user in the top 10% of tweeting activity generated more than 1,600 tweets during the 10-month study period, while the median Republican in the top 10% of tweeting activity produced approximately 800 tweets (Shah et al. 2020). Therefore, these studies indicate that left-leaning users are more prevalent on Twitter and tweet more messages than conservative-leaning users. In sum, one potential reason for the greater presence of conservative-leaning social bots is that populism, increasingly present among right-wing parties, requires the impression of massive popular support. In this sense, conservative-leaning social bots' developers might recognize the need to bolster their support on online platforms in order to be consistent with their rhetoric of embodying the opinions and concerns of ordinary citizens, not elites, as well as to counterbalance the partisan disparities in the prevalence of users from the left and the right.

Moreover, Yan et al. (2021) suggested that right-leaning social bots might be more present on social media platforms because individuals who identify as conservatives are more often deceived by social bot accounts. In fact, Yan and colleagues (2021) conducted an online experiment on 656 participants to assess the effects of partisan identification on bot detection accuracy. Interestingly, they found that Republicans were more likely to be fooled by conservative social bots, notably because of a stronger sense of in-group favoritism. As a result, social bots' enablers could see conservative networks as easier to manipulate, leading them to be more present within these networks. Hence, this could be reflected in the fact that their usefulness is greater within conservative online communities than within liberal ones. This ability of automated accounts to respond to concrete needs for increased visibility and popularity, which are mainly needed on the side of right-wing parties, as well as the fact that conservative users are more likely to be deceived by social bot accounts, could therefore explain the greater presence of conservative-leaning social bots during election periods. With these considerations in mind, I derive my first hypothesis:

> **Hypothesis 1:** *Social bots who lean on the conservative side should be more numerous than those leaning on the liberal side.*

1.3.2. Social Bots' Activities

The scientific literature on social bots has also explored the role of partisan affiliation in their interactions with authentic accounts. Many large-scale studies were conducted recently with Twitter data to understand social bots' activities during electoral campaigns and their involvement with humans. In this sense, it is one thing for social bots of a specific partisan group to be more numerous during political events; it is quite another to engage with humans successfully. This reasoning led Luceri et al. (2019) to explore the effectiveness of social bots in involving authentic accounts in political conversation regarding the 2018 U.S. midterm elections. From a dataset of 2.6 million tweets from around 1 million users, they were able to measure the engagement of humans with social bots' retweets and replies. With these analyses, Luceri et al. (2019) demonstrated that conservative-leaning social bots were significantly more effective at involving humans in their conversations than their liberal-leaning counterparts. Moreover, their study showed that authentic conservative users interacted through retweets with their social bot counterparts almost twice as much as the liberal group. In their study of social media manipulation during the 2020 U.S.

presidential election, Ferrara and his colleagues (2020) employed a similar approach to Luceri et al. (2019) on their dataset of more than 240 million tweets. Their results also confirmed that right-leaning humans retweeted more right-leaning social bots than left-leaning liberals did for left-leaning social bots.

These findings from the American case indicate that right-leaning social bots are more effective than their liberal counterparts at engaging with human accounts. But what explains this increased capacity of conservative-leaning social bots to interact successfully with humans on social media platforms? The first explanation suggests that this situation is the product of the structure of online social networks. In other words, more conservative individuals possess greater political homophily, which leads them to connect more with accounts that exhibit their preferences and opinions (Boutyline and Willer 2017; Hagen et al. 2022). This political homophily is therefore reflected in their online networks, which are denser. For example, Chen et al. (2021b) conducted a five-month experiment on Twitter to examine the impact of social media content on creating homogenous communities. The authors created five groups of neutral social bots (designated as drifters in the paper) which all followed five distinct accounts with different partisan affiliations. The drifters' accounts were then let loose on Twitter during the remaining duration of the experiment. One of the main conclusions from the paper is that right-wing drifters were "gradually embedded into dense and homogeneous networks where they were constantly exposed to right-leaning content" (Chen et al. 2021b, 6). The authors also found that these drifters began to spread right-leaning content throughout the experiment. Even if the small sample size in this experiment does not allow for causal conclusions, it remains interesting to note that during these five months, the interaction of drifters who followed conservative-leaning accounts tilted towards the right.

Consequently, conservative-leaning social bots could be more effective at interacting with humans, given that conservative networks are denser, more consistent, and more embedded.

A second possible explanation stems from the psychology of conservative users. More specifically, studies focusing on the U.S. have demonstrated that individuals who identify as Republicans are substantially less trusting of mainstream media than Democrats (Pennycook and Rand 2019; van der Linden, Panagopoulos and Roozenbeek 2020). Conservatives' perception of their underrepresentation in mainstream media channels of information makes them more open to acquiring their political news from non-traditional sources and even unknown digital entities such as social bots (Kearney 2021). The anti-mainstream sentiment towards the media is reflected in conservative networks where users exhibit lower standards for interacting with unknown accounts (Kearney 2021). In addition, certain socio-demographic factors, notably age, also help explain why conservative-leaning individuals are more likely to interact with automated accounts. Studies showed that older adults tend to lean more towards the conservative side than younger ones (Kuta 2020; Geys, Heggedal and Sørensen 2022). At the same time, older people generally have lower levels of digital literacy and stronger motivated reasoning and cohort effects (Swire, Ecker and Lewandowsky 2017; Grinberg et al. 2019). The heightened engagement of older people, who are often more conservative-leaning, with social bots could therefore be influenced by their lack of ability to distinguish real from automated accounts. In short, conservative-leaning users do not interact more with social bots because of a lack of intelligence but rather because of reactionary responses guided by patterns of motivated reasoning or a lack of digital literacy. Based on these findings from the American context, I put forward my second hypothesis:

Hypothesis 2: *Social bots who lean on the conservative side should interact more with humans than their liberal counterparts.*

1.3.3. Social Bots' Disinformation

Although disinformation is not a new phenomenon given that the traditional news media have also participated in the propagation of disinformation content, especially during the 20^{th} century, the particularity of social media, which operate without any strong *gate-keeping* forces has considerably accelerated the spread of disinformation content (Pickard 2017; Allcott and Gentzkow 2017). Research on disinformation has increased in recent years, but defining disinformation in the age of social media posed some challenges for scholars (Vraga and Bode 2020). Different definitions have been proposed to describe the same problem, and terminological confusion is also present between the terms disinformation and misinformation (Wu et al. 2019). The ambiguity around these different terms has made it difficult for researchers to consolidate results from various studies (Wu et al. 2019). Based on the definition developed by the independent high-level group of experts set up by the European Commission (2018, 5), this research refers to political disinformation as *all forms of false, inaccurate, or misleading information designed, presented, and promoted to intentionally cause public harm or for profit*. According to this definition, disinformation does not include other forms of deliberate but not misleading distortions of facts such as satire or parody. This concept is distinct from misinformation, which is defined as all forms of misleading or inaccurate information shared by online users (European Commission 2018, 10). A critical distinction between disinformation and misinformation resides in the intention (Wu et al. 2019). In other words, disinformation relates to fabricated or deliberately manipulated online content such as conspiracy theories or "fake news". In contrast, misinformation refers to accidental factual mistakes like incorrect statistics, dates, or photo captions. As this work focuses on broad narratives of false information conveyed during the 2021 Canadian election campaign, the term disinformation will be used throughout this research.

Numerous examples of disinformation stories spread during election periods have been identified in the literature. For instance, during the 2020 U.S. presidential election, one of the most familiar disinformation stories was that the Democrats stole the U.S. election through voter fraud (Chen et al. 2021c). Even if social media did not exclusively relay this disinformation narrative (i.e., many Fox News' hosts echoed this false story), it is interesting to notice that Twitter was flooded with hashtags such as #VoterFraud and #StopTheSteal to propagate this disinformation narrative (Abilov et al. 2021; Pennycook and Rand 2021). Disinformation stories related to Canadian politics were also documented during the most recent Canadian election. For example, a disinformation story claiming that the outgoing Prime Minister Justin Trudeau was planning a "climate lockdown" circulated on social media platforms (Bridgman et al. 2022). Cheryl Gallant, an MP for the Conservative Party of Canada, even posted an online video in which she accused the Liberals of preparing a "climate lockdown" (Taylor 2021).

As presented above, the network structure of conservative users diverges from that of liberals. Some scholars claim that the particular configuration of online networks also makes conservative-leaning users more vulnerable to disinformation (McCright and Dunlap 2017; Tucker et al. 2018). Some studies also explored the relationship between the partisan identity of social media users and their propensity to share political disinformation. For example, Grinberg et al. (2019) examined the exposure and spreading of disinformation on Twitter during the 2016 U.S. presidential election. By identifying more than 300 "fake news" sources and linking a sample of over 16,000 Twitter profiles with their voter registration records, the authors found that right-leaning users shared and were more exposed to disinformation content. However, this study equally showed that disinformation content was highly concentrated among clusters; 1% of users generated

80% of the exposure, and 0.1% were responsible for sharing close to 80% of low-credibility sources (Grinberg et al. 2019, 375).

Similarly, Guess, Nagler, and Tucker (2019) also analyzed the dissemination of disinformation content during the 2016 U.S. election period, but on Facebook. To do so, the authors linked an original survey to the Facebook profile of 3,500 individuals. Guess, Nagler, and Tucker (2019) observed the same trends on Facebook as those of Grinberg et al. (2019) on Twitter. As such, less than 10% of users in their sample shared at least one disinformation article, and among those that did, more identified as Republicans (Guess, Nagler, and Tucker 2019, 2). Moreover, Chang et al. (2021) examined social bots' role in the diffusion of disinformation narratives surrounding the 2020 U.S. presidential campaign. One of their main findings regarding social bots' behavior during this period was that close to 13% of Twitter accounts engaged with disinformation narratives were automated accounts. As a comparison, social bots only represented 5% of the sample of accounts that did not engage with this type of content during the study period (Chang et al. 2021, 318). Thus far, this U.S.-based literature informs us, on the one hand, that right-leaning individuals are more likely to share and be exposed to disinformation and, on the other hand, that social bots are also more involved among online networks that share disinformation content than those who do not. Nonetheless, whether one partisan group of social bots spreads more disinformation than another remains unexplored.

In addition, scholars have also proposed some explanations as to why conservative-leaning users are more susceptible to sharing disinformation stories on different social media platforms. One of the most common explanations to account for the ideological asymmetry in the sharing of disinformation relates to psychological factors. According to a 12-wave panel study on 1,204 participants conducted by Garrett and Bond (2021), conservatives were significantly more susceptible than liberals to believe in “fake news”. The authors mentioned that this is partially explained by the fact that the most common political disinformation stories tend to favor conservative positions. Hence, the increased propensity of conservatives to believe political falsehoods is not due to a significant difference in cognitive abilities across partisan groups but rather a pattern of motivated reasoning (Miller, Saunders and Farhart 2016). Furthermore, the behavior of political elites is a factor that can exacerbate patterns of motivated reasoning within partisan groups. Indeed, the influence of elite cues in the political decision-making process and opinion formation on the mass public is well established in the literature (e.g., Gilens and Murakawa 2002; Van Duyn and Collier 2019). In a climate of increased elite polarization, which has been the case in the United States for the past thirty years (see McCarty, Poole and Rosenthal 2006; Iyengar, Sood and Lelkes 2012), elite cues can exacerbate motivated patterns, as well as polarization between partisan groups (Druckman, Peterson and Slothuus 2013; Miller, Saunders and Farhart 2016). This situation could be problematic when applied to disinformation narratives endorsed and vehiculated by political elites. In fact, Macdonald and Brown (2022) found that from January to July 2020 to the same period in 2022, the percentage of news shared that came from unreliable sources by Republican candidates on their Facebook profiles increased from 8%, on average each day, to 36%. During that same two-year period, the percentage went from 1% to 2% for Democratic candidates. Moreover, in their comparative study of the sharing of unreliable information sources by elected members of the U.S. Congress with those of the German and English parliaments, Lasser et al. (2022) also discovered that Republican politicians propagated more untrustworthy information than Democrats on Twitter. This study included all the tweets posted by elected members of Congress/parliament from January 1st, 2016, to March 16th, 2022. This six-year period allowed the authors to compare over time and across three distinct countries. In the U.S., Lasser et al. (2022) noticed that Republican politicians posted more unreliable domains.

In Germany and the U.K., which operate in multi-party systems, parties on the right shared more untrustworthy sources than those on the left. However, Lasser and colleagues (2022) specified that disparities between left and right parties in Germany and the U.K. are smaller than what was observed in the U.S. and have remained steady throughout the six-year period, which is not the case for the U.S., where the sharing of unreliable information by Republican politicians increased substantially since the election of Joe Biden. In this sense, a distinction in the spreading of false information sources from politicians from left and right parties could also be observed in Canada in proportions similar to those of Germany and the U.K. In sum, disinformation cues sent in increasing proportions by Republican elites, with Donald Trump leading the way, could have stimulated the motivated reasoning of conservative users and ultimately made them more vulnerable to online political disinformation. To adequately mimic the behavior of real users, conservative-leaning social bots could therefore share more disinformation content on social media platforms. Informed by these findings from the literature, I develop the third and last hypothesis:

> **Hypothesis 3:** *Social bots who lean on the conservative side should share more disinformation content than their liberal counterparts.*

1.4. The Case of the 2021 Federal Canadian Election

It is possible to conclude from the studies discussed above that social bot accounts are an integral part of the digital political environment in many Western democracies and that they are used to interfere in public debates taking place on social media platforms, especially Twitter. However, the analysis of partisan differences in the strategies and behaviors of social bots in the specific context of Canada remains relatively unexplored. McKelvey and Dubois (2017) were among the first to study the use of social bots during a Canadian election. Their study of the 2015 Canadian election concluded that social bots did not have "as great an influence on Canadian politics as their international counterparts" (McKelvey and Dubois 2017, 21). However, the essentially qualitative results of this analysis were limited in their reach, given that the authors explored only five suspected automated accounts. With the general election held in 2019, some researchers have been interested in the use of social bots in a Canadian election. For instance, *The Digital Democracy Project* (Owen et al. 2020) produced a detailed report on the digital media ecosystem in Canada in the run-up to the 2019 election. A short section of this report focused on social bots. By testing a sample of around 170,000 users, they found that automated accounts were not prevalent and that hashtags were more homogenous for social bot accounts than humans. In addition, Beskow and Carley's (2020) study, which aimed to evaluate the state-of-the-art graphical and semantic embedding for social media data, also used the 2019 Canadian election as a brief case study. Their analysis showed that accounts associated with the political right and left in Canada were actively involved during the campaign. Automated accounts were most active on Twitter during the release of Prime Minister Justin Trudeau's blackface pictures and the days surrounding the election. In a study to precisely analyze Twitter accounts that used the hashtags #FakeNews and #NotABot during the 2019 federal election, Bellutta, King, and Carley (2021) demonstrated that these two hashtags mainly stemmed from retweets and not from original content. Moreover, the hashtag #NotABot was not indicative of an account not being a social bot and was used to amplify anti-Trudeau rhetoric during the campaign. The most extensive study yet of a Canadian election is Rheault and Musulan (2021), which was previously discussed in the section on social bots' partisan behaviors. In their study, Rheault and Musulan (2021) reported, among other things, that social bot accounts were found within all partisan clusters but that the social bot density was greater for conservative parties, most notably the PPC. These few studies thus show that social bots are an issue of interest in Canada. The concerns surrounding social bots and their propensity to spread political disinformation have been expressed in several media outlets in recent years. Indeed, journalists and experts have even tried to draw this phenomenon to the attention of Canadians (e.g, Kassam 2018; Carvin 2021; Nuttall 2021). Yves Côté, Canada's former chief election watchdog, even stated that disinformation and foreign interference are two of the biggest threats facing Canada's electoral system (Thompson 2022).

The 2021 Canadian federal election thus offers an interesting case to examine in terms of disinformation and partisan differences, as it took place amid the COVID-19 pandemic. Indeed, studies in Canada have shown that all else being equal, right-wing individuals viewed COVID-19 as less severe than their left-wing counterparts (Merkley et al. 2020; Pennycook et al. 2021). Furthermore, this public health crisis has brought new and polarizing issues to the forefront of Canadian politics and fostered disinformation narratives online (Bridgman et al. 2022). Four important disinformation narratives circulating during the election were identified by Bridgman and Lavigne (2022, 8-9): (i) Erin O'Toole, leader of the Conservative Party of Canada (CPC), wanted to privatize the Canadian healthcare system, (ii) the Liberals were preparing to impose a climate lockdown on Canadians, (iii) protestors were disrupting Justin Trudeau's political rallies during the campaign, as part of a centralized operation that aims to prevent the Liberals from campaigning in key ridings, and (iv) Justin Trudeau instructed Jody Wilson-Raybould, the former attorney-general of Canada, to lie about the SNC-Lavalin affair. While the outgoing Prime Minister, Justin Trudeau, was often a central target of these disinformation narratives, the false story about Erin O'Toole wanting to privatize the healthcare system got a lot of online exposure, especially among liberal partisans (Bridgman and Lavigne 2022).

While Canada has for some time resisted the rise of populism and increased polarization, as seen, for example, in the United States, recent evidence suggests that both are gaining momentum (Medeiros 2021; Boxell, Gentzkow and Shapiro 2022; Djuric 2022). In this particular context, on August 15th, 2021, an early election was called at the request of Prime Minister Justin Trudeau. The 44th Canadian general election lasted 36 days, and the vote was held on September 20th, 2021. The outgoing Liberal Prime Minister hoped that his successful handling of the COVID-19 pandemic, according to a majority of Canadians, would secure him a majority government (Argitis and Hagan 2021). A series of different issues, such as foreign policy, climate change, and gun control, marked this election. However, the management of the COVID-19 pandemic emerged as a major and contentious issue of the campaign (Clarke, Scotto and Stewart 2022; Taylor 2022). The Liberal Party proposed the imposition of a vaccination passport on interprovincial trains, commercial flights, and cruise ships, as well as vaccine mandates for federal public servants (Liberal Party of Canada 2021). The New Democratic Party, under the leadership of Jagmeet Singh, also proposed the same policies regarding the management of COVID-19 and proof of vaccination (New Democratic Party 2021). It is also important to highlight that the Bloc Québécois led by Yves-François Blanchet, which is not included in this analysis, nonetheless supported the measures to counter the spread of COVID-19, as did the LPC and the NDP (Bloc Québécois 2021). In contrast, the LPC's most prominent opponent, the Conservative Party led by Erin O'Toole, opposed vaccine requirements, advocating instead for rapid testing of non-vaccinated people (Conservative Party of Canada 2021). The People's Party of Canada, led by Maxime Bernier, promoted the most oppositional policies to COVID-19 measures by objecting to vaccine mandates and passports, mandatory mask mandates, and lockdowns (People's Party of Canada 2021). For its part, the marginal Green Party of Canada, newly led by Annamie Paul, adopted an ambiguous policy regarding the management of the pandemic. The leader of the GPC encouraged Canadians to get vaccinated, but the party platform did not mention vaccine passports or mandates (Green Party of Canada 2021). Experts have even pointed out that the politicization of COVID-19 issues was detrimental to the social climate in Canada, as it exacerbated partisan divisions, polarization, and the reach of disinformation theories (Blouin 2021).

In the end, Justin Trudeau's Liberals were able to win the election with 160 seats in the House of Commons (32.60% of the vote share), but they were unable to form a majority government. The composition of the House of Commons remained almost the same as before the election was called. Nevertheless, the 2021 Canadian federal election is interesting to examine since it occurred in a highly polarizing pandemic context and featured several disinformation narratives. With this information concerning the Canadian case, I amend my three hypotheses formulated previously to situate them precisely in the context of the Canadian federal election of 2021:

> **Hypothesis 1:** *During the 2021 Canadian election, social bots who lean on the conservative side should be more numerous than those leaning on the liberal side.*
>
> **Hypothesis 2:** *During the 2021 Canadian election, social bots who lean on the conservative side should interact more with humans than their liberal counterparts.*
>
> **Hypothesis 3:** *During the 2021 Canadian election, social bots who lean on the conservative side should share more disinformation content than their liberal counterparts.*

1.5. Conclusion

This chapter has shown that social bots are used by various political actors to manipulate public opinion, amplify exposure to certain narratives, and share political disinformation content. They can achieve these objectives through various strategies such as astroturfing, smoke screening or misdirecting. Several studies have also demonstrated the sustained involvement of social bots in electoral campaigns and the partisan differences in their behaviors and interactions with humans. However, previous studies on social bots have focused primarily on the United States and have not addressed partisan differences in the propensity of automated accounts to share disinformation. To address this omission in the literature, this research will examine the activities and amount of disinformation content propagated by social bot accounts based on their partisan leanings during the 2021 Canadian federal election.

Chapter 2. Methodology

This chapter describes the Twitter collection process for the 2021 Canadian election data and includes a brief overview of the dataset that will be used to perform the analyses. In order, the chapter will review the methodology used for (1) detecting social bots, (2) classifying these accounts according to their partisan leaning, (3) measuring social bots' strategies and their interactions with humans, and (4) spotting disinformation content related to the COVID-19 pandemic and the election campaign in users' tweets.

2.1. Data

In this study, Twitter data is leveraged to examine the activities of automated accounts according to their partisan leaning during the 2021 Canadian federal election. To that end, a team at the *Media Ecosystem Observatory* made available their dataset on this most recent election. The data collection was achieved through the Twitter Application Programming Interface (API) for Academic Research on October 23rd, 2021. The tweets contained in the dataset were all created between July 31st, 2021, and October 22nd, 2021, therefore capturing discussions that took place in the days before the beginning of the campaign up to the weeks following the election day.

Table 2.1. Summary of the 2021 Canadian Election English Dataset

Statistics	Count
Total number of unique users	1,114,906
Total number of tweets	11,361,581
Number of retweets	8,438,765

To obtain tweets related to the 2021 Canadian election, a list of relevant keywords and hashtags in both English and French was used as a filter during the data collection process (see Appendix A1). The list was constructed by members of the *Media Ecosystem Observatory* from the most popular terms that circulated on Twitter before the election was launched, in addition to terms related to the main federal political parties and their leaders. As a result, more than 12.3 million tweets from over 1.3 million unique users were collected. As this study focuses on user posts in English, 1,015,628 tweets written in another language and 192,370 users who did not tweet in English were removed from the sample. Over 30 different languages were found in the dataset, but English tweets made up the vast majority, with 92%. After cleaning the original dataset, a total of 11,361,581 tweets from 1,114,906 unique users were retained to form the sample that was used in subsequent analyses. As shown in Table 2.1, retweets are prevalent, as they constitute almost three-quarters of the messages in the sample.

2.2. Bot Detection

To examine the role played by social bots in the 2021 Canadian election, the first task was to identify automated accounts in the dataset. Bot detection is now an essential tool for understanding the dynamics of manipulation within social media networks and, more importantly, for mitigating their negative impacts. However, detecting social bot accounts on online platforms has become increasingly complex, especially since they are more sophisticated, and their operational tactics are constantly evolving (Deb et al. 2019; Cresci et al. 2019). Accordingly, Cresci (2020) proposes a three-stage theoretical framework to understand the evolution of automated accounts and their different characteristics. In the first stage, which lasted until about 2011, automated accounts populating Online Social Networks (OSNs) were simple since they were mainly used to spam content; they showed apparent signs of automation and had little social interaction with other users. In the second stage, social bots became more refined and progressively more credible by increasing their social connections (i.e., following each other) and ceasing to spam the same content continually. Finally, in the third stage, which began around 2016, social bots became capable of mimicking human behaviour by having detailed profiles, extensive networks of followers and friends, and sharing malicious messages.

Researchers have also had to adapt and utilize more precise detection methods as social bots evolved. Consequently, several tools have been developed to identify malicious automated accounts effectively. These bot detection algorithms can be grouped into two broad categories: supervised and unsupervised. Most studies have favored supervised machine-learning approaches to detect automated accounts on Twitter (Rodríguez-Ruiz et al. 2020). Supervised detection tools rely primarily on manually annotated datasets for training purposes. In other words, labeled data from both human and social bot accounts are passed to the learning algorithm, which is then trained to classify "unmarked" accounts. Such classifiers also learn from a wide range of features like text content, profile content, user-based behaviour, network activity, and even temporal activity (Ng, Robertson and Carley 2022). However, it is essential to recognize that there are inherent limitations to this type of method, notably the lack of large, reliable training datasets as well as the absence of ground truth data. These limitations have raised concerns about the validity of the annotating process (Cresci 2020). As previously discussed, the definition of what constitutes a social bot remains somewhat unclear and leaves room for interpretation in the classification process. In addition to inter-coder reliability techniques, scholars have looked at account suspension on Twitter as a validity measure to minimize annotator bias (Chavoshi, Hamooni and Mueen 2016a). Nevertheless, as Yang et al. (2019) mentioned, using annotated labels as proxies for ground truth has proven to be viable and effective in developing supervised bot detection tools.

Unsupervised learning algorithms have also been used in the field of automated accounts detection (e.g., Chavoshi, Hamooni and Mueen 2016b; Chen and Subramanian 2018). This method has been implemented to address the labeling problem that emanates from supervised learning models. Indeed, unsupervised methods focus on detecting groups of coordinated and synchronized accounts instead of classifying them individually (Yang et al. 2019; Cresci 2020). Rather than being trained on labeled data, unsupervised learning algorithms learn connectivity patterns from unlabeled data. Hence, unsupervised approaches do not detect automated accounts individually but identify clusters showing signs of automation. Both supervised and unsupervised methods have their place in the field of bot detection. Each has advantages and limitations, but they complement one another depending on the objective to be reached (i.e., account or group-level detection).

In this work, the supervised machine learning system Botometer[4] (formerly known BotOrNot), which is considered the state-of-the-art model for bot detection (Shao et al. 2018; Vosoughi, Roy and Aral 2018; Sayyadiharikandeh et al. 2020; Yang, Ferrara and Menczer 2022), is used to distinguish between automated and human accounts. This model was employed in several important studies to analyze the behavior and impacts of automated accounts (e.g., Bessi and Ferrara 2016; Ferrara et al. 2016; Varol et al. 2017; Wojcik et al. 2018; Luceri et al. 2019). The system was jointly developed by the Observatory on Social Media (OSoMe) and the Network Science Institute (IUNI) at Indiana University and has been publicly available since 2014. There are two ways to use the model: the web interface or the Botometer Pro API.[5] The present work uses the latest version of the Botometer, which is the fourth version. As mentioned earlier, social bots' characteristics and behaviors change over time. Therefore, the model required updates to remain

[4] https://botometer.osome.iu.edu/
[5] To access the Botometer Pro API, a Rapid API account, and an API subscription plan (free or paid) is required. More details can be found here: https://rapidapi.com/OSoMe/api/botometer-pro

relevant and accurate (Yang, Ferrara and Menczer 2022). Botometer-V4 is trained on several different publicly available datasets.[6] This version has undergone important architectural changes to better capture the unique behavioral patterns of automated accounts and humans. Indeed, Botometer-V4 uses Random Forest classifiers or Ensemble of Specialized Classifiers (ESC), which are trained on various bot classes to identify novel and heterogeneous automated behavior on Twitter. Therefore, ESC results are more reflective of the evolution of social bots' practices and offer more transparency on the scores assigned to users (Sayyadiharikandeh et al. 2020).

Botometer is also trained to calculate and return a "bot-likelihood" score for each account sent through the API. This bot probability score ranges between 0 and 1. A score closer to 0 indicates that the account is more likely to be handled by a human. Conversely, a score closer to 1 indicates that the account shows heavy signs of automation and is more likely to be a bot. However, several steps are performed by the model before predicting the bot score of an account. First, via Twitter's public API, the model can access the public profile of a given account and hundreds of tweets and mentions.[7] Note here that the users' identification numbers were utilized instead of their screen names since Twitter allows users to modify their screen names at any point. In contrast, a user's primary identification remains the same (Rheault and Musulan 2021, 328). With the information provided to the Botometer API, the model extracts and analyzes over 1,200 features that can be grouped into six categories: user-based, friends, network, temporal, content and language, and sentiment (Varol et al. 2017). Afterwards, the ESC results are aggregated using the maximum rule, and the model assigns a bot probability to each user. It is also important to mention that the model returns an "English bot score" and a "universal bot score", which is a language-

[6] The public datasets used to train Botometer-V4 are available here: https://botometer.osome.iu.edu/bot-repository/datasets.html

[7] To access the Botometer Pro API, Twitter API Keys are necessary. It is possible to apply for them here: https://developer.twitter.com/en/docs/twitter-api

independent score best suited for non-English content. This is done for each account. Since accounts that only tweeted in a language other than English were removed from the dataset, the English bot score is considered in this analysis. The changes implemented in this latest version of the model have proven to be more accurate when detecting automated accounts. In their 2020 paper, which focused on testing the performance of Botometer-V4, Sayyadiharikandeh and colleagues demonstrated, using 5-fold cross-validation, that the Botometer-V4 has improved accuracy metrics than previous versions of the model. As such, Sayyadiharikandeh et al. (2020) achieved an F1-score improvement of 56% compared to the Botometer-V3 baseline, which means that the V4 model is a more generalizable predictive model in comparison to earlier versions. Furthermore, Sayyadiharikandeh et al. (2020) also obtained an Area Under the Curve of 0.99 for the Botometer-V4, which suggests that the algorithm can distinguish between automated and human accounts with very high accuracy (Yang, Ferrara and Menczer 2022).

As mentioned above, Botometer returns a bot likelihood score for a given account. The model, therefore, does not explicitly assign an account as either a bot or human. For this study, the bot scores must be dichotomized to assign each account to one of the categories mentioned. There is no consensus on which threshold value should be selected to discriminate between automated and human accounts. Previous studies have used threshold values varying from 0.3 (Luceri et al. 2019), 0.5 (Bessi and Ferrara 2016; Varol et al. 2017; Shao et al. 2018; Rheault and Musulan 2021), 0.7 (Woolley and Guilbeault 2018), 0.75 (Keller and Klinger 2019) up to 0.8 (Broniatowski et al. 2018). The selection of an appropriate threshold value is crucial since a lower threshold value can identify more social bot accounts but can lead to more false positive errors. In contrast, a higher threshold value might reduce the number of false positives but identify fewer actual social bots (Duan et al. 2022). With these considerations in mind and informed by the distribution of bot scores in my sample (see Figure 2.1), a conservative threshold of 0.7 is retained, which given the bot score distribution appears to be the more balanced choice. In other words, each account with a score above 0.7 is treated as an automated account, and each account with a bot score below or equal to 0.7 is labeled as human (see Bessi and Ferrara 2016; Varol et al. 2017; Shao et al. 2018).

Since the Botometer algorithm was previously trained on different datasets than the one used in this research, as a precautionary measure, further manual validation was conducted by the author to ensure the validity of the results obtained. Therefore, 100 accounts (50 from each group) were randomly sampled and validated by closely examining each account's profile and the content of their tweets and retweets.[8] This manual validation method is standard practice and has been used in different papers that employ Botometer (e.g., Sayyadiharikandeh et al. 2020; Chang and Ferrara 2022; Gallwitz and Kreil 2022). As a result, manual validation by the author yields an average agreement of close to 85% with the Botometer algorithm. This correspondence score is approaching the 88% score obtained by the manual validation conducted by Chang and Ferrara (2022, 4), who employed the same manual validation strategy in their study of social bots and humans' online behaviors and interactions regarding the COVID-19 pandemic. However, Chang and Ferrara (2022) randomly sampled 100 accounts from each group on a more than 5.4 million distinct users' dataset. Therefore, the proportion of manually validated accounts in this research is larger than in Chang and Ferrara's study (2022).

[8] The *accountanalysis* app was used to manually analyze Twitter accounts: https://accountanalysis.app/

Figure 2.1. Bot Score Distribution

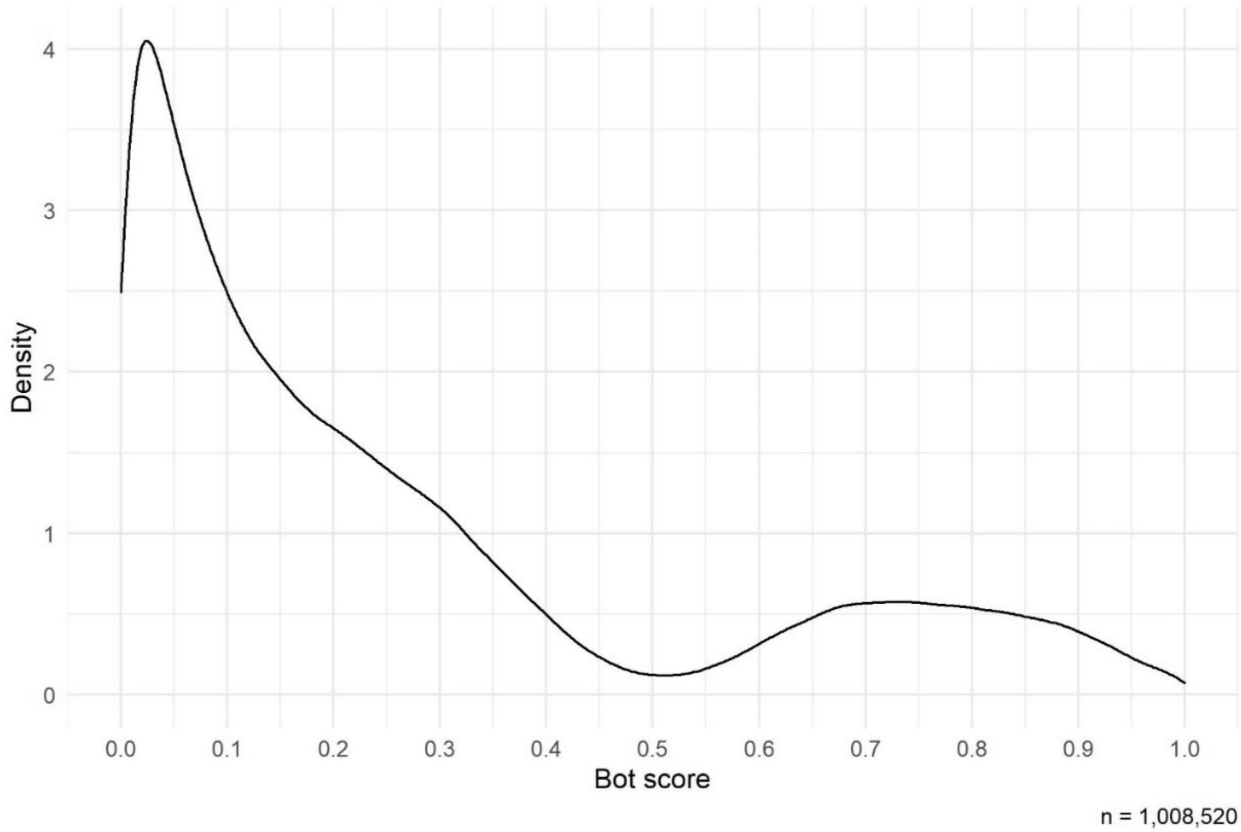

Figure 2.1 shows the distribution of bot scores produced by Botometer-V4. It shows a clear binomial break centered around 0.5. The uptick close to 0.7 also indicates that a significant number of accounts exhibit obvious bot characteristics (Bessi and Ferrara 2016).

In summary, although there is a wide range of bot detection models such as DeBot (Chavoshi, Hamooni, and Mueen 2016b), Tweetbotornot (Kearney 2018), RTbust (Mazza et al. 2019), and Bot-hunter (Beskow and Carley 2018), the supervised learning algorithm Botometer is retained for this analysis. This choice is due to both theoretical and practical considerations. On the one hand, Botometer-V4 has been shown to have the most promising automated account detection results. Recent studies have found that Botometer-V4 is the best-performing bot detection algorithm on recently published Twitter account datasets (Sayyadiharikandeh et al. 2020; De Nicola, Petrocchi and Pratelli 2021). On the other hand, the ability to process 17,280 accounts per day with the Botometer Pro API Ultra plan is another advantage of this model, considering that the sample size in this analysis is greater than one million users.[9]

2.3. Party Leaning Classification

Another essential task in order to assess the involvement of social bots in the 2021 Canadian election is, of course, to determine the partisan leaning of these accounts. Over the past few years, different models to predict the political orientation of social media users have been developed. These models are particularly important for comparing partisan dynamics and measuring partisan polarization on social media (Gruzd and Roy 2014; Yang et al. 2017).

Partisan prediction models for social media users can be grouped into three broad classes. The first approach is content-based and relies on dictionaries to identify political issues in users' messages. Subsequently, specific words and hashtags contained in profile descriptions or posts are used to infer the political orientation of users (e.g., Preoţiuc-Pietro et al. 2017). The models derived from this approach often use word embedding or pre-trained language models to make their classifications. Another part of content-based approaches looks at the URLs of the media outlets they post on their profiles (e.g., Badawy, Ferarra and Lerman 2018; Luceri et al. 2019). The second group of methods is based on users' network connections. This approach relies on the structure of users' networks to predict their political orientation. Some scholars use the information provided by users' network activities (e.g., Stefanov et al. 2020; Rheault and Musulan 2021), such as retweets, replies, mentions, or likes, while others base their prediction on the relationship networks of accounts, like who they follow and who follows them (e.g., Barberá 2015). In addition, some network-based models combine both activity and relation features to classify social media accounts

[9] The Basic and Pro plans are also offered via the Botometer Pro API and can process 500 and 2,000 accounts per day respectively: https://rapidapi.com/OSoMe/api/botometer-pro/pricing

(e.g., Gu et al. 2016; Xiao et al. 2020). Lastly, the third class of models combines features of both the content and network approaches (e.g., Colleoni, Rozza and Arvidsson 2014; Pennacchiotti and Popescu 2011). As such, this hybrid approach considers, on the one hand, how users talk about political issues through keywords and hashtags and, on the other hand, who they follow and interact with online. Interestingly, in their literature review of various party prediction models for Twitter, Pelrine et al. (2022) showed that having more features does not necessarily guarantee higher classification accuracy.

In this work, a model from the content-based approach is retained to predict the partisan leaning of both human and social bot accounts. Most prediction models have focused on classifying U.S. users as either Republican/conservative or Democrat/liberal. The binary classification of users' partisan leanings is also appropriate to use in the Canadian context, given that the political system is primarily structured around the left-right divide, as is the case in the U.S. (Cochrane 2010; Cochrane 2015). Accordingly, this research relies on a binary classification between liberal-leaning and conservative-leaning social media users. The choice of a binary classification process is explained by the fact that large datasets of annotated Canadian users are rare and that Canada operates in a multi-party system where six national parties competed in the last election.[10] Therefore, the multiplication of different political parties makes it difficult to generate reliable and accurate predictive models. Compared to the U.S., where the two major political parties, Republican and Democrat, are closely associated with conservative and liberal ideologies, Canadian parties are essentially brokerage parties, which tend to be less ideological and more flexible (Young and Cross 2002). Modern brokerage politics in Canada is now more defined by

[10] It is important to note that the Bloc Québécois only presents candidates in Québec and does not aspire to form the next Canadian government, as opposed to the five other parties.

parties' increasing abilities to target and aggregate various segments of the electorate (Giasson, Lees-Marshment and Marland 2012). As a result, Canadian parties can be seen as positioned on an ideological continuum ranging from the left to the right. The Liberal Party of Canada (LPC), the New Democratic Party (NDP), and the Green Party of Canada (GPC) are considered left-leaning parties, whereas the Conservative Party of Canada (CPC) and the People's Party of Canada (PPC) are regarded as right-leaning parties (Kevins and Soroka 2018; Merkley 2020; Merkley 2021).[11] This positioning is not arbitrary but stems from the parties' preferences on economic and social issues (Cochrane 2010; Rheault and Cochrane 2020; Vox Pop Labs 2021). It is also interesting to note that following the latest election, the LPC and the NDP reached a "supply-and-confidence" agreement, valid until June 2025, which indicates ideological proximity between the parties (Zimonjic 2022).

In this study, Twitter accounts are labeled ideologically as either liberal-leaning or conservative-leaning using two complementary approaches based on (i) users' profile descriptions and (ii) users' tweet activity. In order to train the two proposed approaches, potential supporters from the five main Canadian political parties (i.e., LPC, CPC, NDP, GPC, and PPC) needed to be identified using their profile descriptions.[12] In other words, potential partisan users were identified using a list of keywords associated with a specific party identification (see Appendix A2). When a user's profile description contained a minimum of one partisan keyword from one of the five political parties, the user was assigned the partisan affiliation associated with this political party.[13]

[11] Note that the Bloc Québécois, a regionalist party from the province of Québec, is not included in this research since it is not a national party. Because of methodological reasons, French tweets were removed from the data sample.
[12] Due to a technical problem in the collection of the 2021 Canadian election dataset, users' profile descriptions were not collected exactly at the same time as their tweets. Their profile descriptions were retrieved a few months after their tweets.
[13] Note that when the user's bio description contained keywords associated with more than one political party, their partisanship was assigned to the party with the most frequent keywords. In cases where two contradicting keywords appeared in the bio description, the first political keyword detected was used to assign the user's party affiliation.

For example, a profile description containing the term "Justin Trudeau" would have been classified as a supporter of the LPC. By using this content-based method, assigning a partisan affiliation to a total of 18,818 users was possible.[14] Afterwards, two groups of 2,000 users flagged as LPC and CPC supporters were randomly sampled. Because of the fewer accounts for the other three parties, the same process was applied, but this time on the totality of flagged users from the NDP, GPC, and PPC. For validation purposes, the author manually labeled all the users from these five distinct samples by examining their profile descriptions and all their available tweets in the dataset. The results of this manual classification process are presented in the confusion matrix in Table 2.2.

Unfortunately, it was impossible to assign a party affiliation to 1,798 users, either because their profile was not explicit enough to infer their party affiliation or because they did not directly discuss issues related to Canadian politics. For example, Table 2.2 shows a high number of NAs for the PPC. This is explained by the fact that hashtag #PPC is also used in online advertising in reference to pay-per-click marketing. Therefore, users who used this hashtag with a non-political connotation could not be assigned a partisan affiliation. Furthermore, the horizontal line of the confusion matrix indicates the actual partisan affiliation of users, while the vertical line indicates the predicted values. As such, the diagonal line of Table 2.2 represents the number of values that were accurately predicted for each of the five political parties. Dividing the number of well-predicted values by the total number of predicted values, this classification process only achieved an overall accuracy of roughly 68%.[15] Even though the overall accuracy of this classification

[14] A total of 7,811 users were initially flagged as LPC supporters, 7,903 as CPC supporters, 1,885 as NDP supporters, 322 as GPC supporters, and 897 as PPC supporters.
[15] See Kulkarni, Chong, and Batarseh (2020, 86-92) for more information on the interpretation of the confusion matrix.

process is low, it is important to reiterate that we only employ the manually validated classifications to train the predictive model.

As such, the manual labels assigned to the 7,099 users are referred to as "weak labels" and were used to train the profile classifier of the predictive model of users' partisan leaning. This classifier can assign a partisan affiliation for each user in the dataset based solely on the content of their Twitter profile descriptions – or Twitter bio descriptions. The profile classifier employed in this study was trained using RoBERTa-large (Liu et al. 2019). RoBERTa-large is a Natural Language Model (NLP) pre-trained on a vast corpus of English text data (Hugging Face 2022). This model has the specificity of having been pre-trained with two specific objectives: masked language modeling and next-sentence prediction.

Table 2.2. Summary of Profile-Based Classification

	LPC	CPC	NDP	GPC	PPC	NO PARTY
Flagged-LPC	1,596 (79.80%)	88	9	2	35	270
Flagged-CPC	57	1,875 (93.75%)	13	0	9	46
Flagged-NDP	142	69	675 (35.85%)	8	29	960
Flagged-GPC	1	1	1	288 (89.72%)	1	29
Flagged-PPC	12	6	6	1	377 (42.12%)	493
Total	1,808	2,039	704	299	451	1,798

Therefore, RoBERTa-large can learn an inner representation of the English language, which is helpful to extract specific language features. In other words, the profile descriptions of the manually labeled accounts, which were assigned a weak label were used as inputs to train the RoBERTA-large model to predict with precision the party affiliation of the remaining users in the dataset. For this model, an 80-20 train-test split was used to create the validation set. This means that 80% of the data was assigned to the training dataset and the remaining 20% was employed as the test dataset. With this 80-20 train-test split, the accuracy of the profile classifier reached an overall score of over 87%. Table 2.3 details the profile classifier results for each Canadian political party. Overall, the profile classifier achieved a good predictive performance across political parties. Indeed, the F1-score, which is the harmonic mean of precision and recall, is over 80% for all parties.[16] Therefore, this indicates that the profile classifier yields precise and robust results. However, one of the main limitations of this classifier is that it can only be applied to users with explicit profile descriptions. Since one of the objectives of this research is to determine the number of active social bots based on their partisan leanings, another type of classifier for users with no political keywords in their Twitter profile description was required.

Table 2.3. Results From the Profile Classifier

Party	Sample Size	Weak Labels	F1-Score (Profile)
LPC	2,000	1,808	88.67
CPC	2,000	2,039	91.67
NDP	1,883	704	80.85
GPC	321	299	95.87
PPC	895	451	83.24
NO PARTY	-	1,798	81.78
Weighted Average	7,099	7,099	87.23

This means that in order to scale up the number of users with an assigned partisan leaning another classifier trained on users' tweets was added to the model. Therefore, this addition made it possible to predict the partisan leaning of users who did not have an explicit political profile based on the content of their tweets. In addition to the weak labels manually assigned to 7,099 users, a sample of more than 1,500 Canadian users with self-declared party affiliations was also added as a validation measure for the activity classifier. The survey data was collected by the *Media Ecosystem Observatory*[17] from September 19 through September 24, 2019, and contains

[16] See Lipton, Elkan and Narayanaswamy (2014) for more information on the F1-score.
[17] The survey data was also used by Dubois and Owen (2020).

information about individuals' party affiliations (LPC, CPC, NDP, GPC, or BQ), as well as their Twitter handles. With this information, it was possible to pair 545 individuals from the survey data with accounts from the 2021 Canadian election dataset. Among the matched users are 214 LPC supporters, 103 CPC supporters, 98 NDP supporters, 19 GPC supporters, 83 users with no party affiliation, and 28 individuals affiliated with marginal Canadian parties that are not included in this research.[18] In other words, these 545 true labels were added to the 7,099 weak labels to test the activity classifier's predictive accuracy. The results of the activity classifier presented in Table 2.4 were therefore obtained using the 7,644 true and weak labels.

The first step for training the activity classifier required us to obtain the users' embeddings. In other words, 1024-dimensional vectors were produced to represent each user's tweets by grouping together all their tweets' embeddings. The users' embeddings were then used to train a random forest classifier for which the profile classifier's labels were employed as weak labels for training purposes. An 80-20 train-test split was subsequently performed, allowing us to find the optimal subset of trees for this predictive model. As a result, the random forest classifier trained on users' embeddings had 1,000 trees with a maximum depth of 50. In other words, each decision tree had a high depth allowing the model to perform more splits and capture more information about the data. Unfortunately, due to the many political parties in the Canadian system, the activity classifier did not perform very well when assigning a party affiliation to the remaining users. Therefore, we switched to a binary classification to increase the activity classifier's predictive capacity. On the one hand, the weak labels from the LPC, NDP, and GPC were aggregated under a new ideological category named "liberal-leaning". On the other hand, the "conservative-leaning" ideological category is the result of the combination of weak labels from the CPC and PPC.

[18] To train the model, the 545 true labels were not included in the random samples of manually labeled users.

Subsequently, the activity classifier was retrained with this new binary classification, which resulted in an improved level of prediction accuracy. As shown in Table 2.4, the activity classifier reached an overall accuracy of more than 85%. The activity classifier's accuracy score is almost the same as the one obtained with the profile classifier. This similar score is primarily explained by the fact that users with no profile description are more difficult to classify, especially since many of them have few tweets on which the model can rely to make predictions. It is relevant to mention that the users classified by the profile classifier were also converted into the binary classification between "liberal-leaning" and "conservative-leaning". Users classified as LPC, NDP, and GPC supporters were grouped into the liberal-leaning category; the same was done for accounts classified as CPC and PPC supporters, which were relabeled as conservative-leaning.

Table 2.4. Results From the Activity Classifier

Partisan Leaning	Precision	Recall	F1-Score	Accuracy
Liberal-Leaning	84.92	85.85	85.39	-
Conservative-Leaning	85.83	84.90	85.36	-
Weighted Average	85.38	85.38	85.37	85.38

If users in the dataset had explicitly political profile descriptions, their partisan leaning was determined based on the profile classifier. In contrast, if the users' profile descriptions did not contain any political keywords, they were then assigned a partisan leaning according to the activity classifier of the model. Weak labels manually assigned to users, and true labels from the survey data were kept as the final partisan leaning for the 7,644 users concerned. Therefore, this method ensures that a partisan leaning was assigned to as many users as possible in the dataset while maintaining a high level of predictive accuracy of more than 85%.

In summary, the results from the Botometer algorithm and the partisan leaning classification model will allow us to establish which accounts are humans/social bots and which lean on the liberal/conservative side. This classification process will result in four distinct groups based on account type and partisan leaning of users. The four groups, whose distribution will be presented in more detail in section 3.2, are the following: liberal-leaning human, conservative-leaning human, liberal-leaning social bot, and conservative-leaning social bot. Knowing the distribution of each of these groups will make it possible to test the first hypothesis and assess whether *social bots who lean on the conservative side were more numerous than those leaning on the liberal side*.

2.3.1 Measuring Social Bots' Strategies and Interactions

The following section aims to introduce the metrics used to identify social bots' techniques and measure their level of engagement with humans. First, I attempt to determine whether social bots employed a particular coordination strategy during the 2021 Canadian election. In this sense, the tweets and retweets associated with the automated accounts in the dataset were checked to find out if it was possible to detect astroturfing. As mentioned by Zhang, Carpenter, and Ko (2013) and Brachten et al. (2017), an astroturfing strategy revolves around artificially amplifying messages around a candidate or a topic through automation. Therefore, to examine the presence of an astroturfing strategy among social bots, the following *Astroturfing Ratios* proposed by Brachten et al. (2017) were employed. While these two metrics may seem simple, they are very valuable as they directly measure the proportion of each of the ten most popular hashtags and keywords based on account types and their partisan leanings. Lastly, these measures are performed separately on the most popular hashtags and keywords for each group (i.e., liberal-leaning human, conservative-leaning human, liberal-leaning social bot, and conservative-leaning social bot), which will allow a comparison of the results between the groups and will also make it possible to determine the presence or not of an astroturfing strategy:

$$\textit{Hashtag Astroturfing Ratio}_{PA}: \frac{\text{hashtag count\_i}}{\text{total no. of tweets}} \qquad (1)$$

$$\textit{Keyword Astroturfing Ratio}_{PA}: \frac{\text{keyword count\_i}}{\text{total no. of tweets}} \qquad (2)$$

P = Partisan leaning (conservative = 0; liberal =1); A = Account type (human = 0; social bot = 1)

These two operations are conducted distinctively on the top ten most popular hashtags, and the top ten most popular keywords in each of the four groups mentioned above. Performing both operations for each partisan leaning group (i.e., conservative and liberal), represented by "P" in the equation, in combination with each account type (i.e., human and social bot) which is represented by "A" in the equation will provide distinct astroturfing ratios for each of the four groups. Therefore, each of the ten most frequently used hashtags and keywords for every group is assigned an *Astroturfing Ratio*. It is also important to specify that the hashtag count, which appears in the numerator of the first operation, represents the number of times a hashtag, which refers to words preceded by the symbol #, was employed by a particular group. The keyword count represents the number of times a group used a keyword. By dividing the number of times the ten most popular hashtags and keywords were used by the number of tweets produced for each of the four groups, it is, therefore, possible to examine if specific topics were artificially pushed by automated accounts (Brachten et al. 2017, 7). Moreover, a ratio equal to one means that a specific hashtag or keyword is included in all the group's tweets. The more the ratio tends towards one, the more likely there was an astroturfing strategy and, thus, artificial amplification.

The *Astroturfing Ratios* are included in this research since they allow, on the one hand, to see the most popular topic in each group and, on the other hand, to determine if automated accounts pushed some topics by comparing the hashtag and keyword frequency between social bots and humans. For example, if the astroturfing ratio of a hashtag like #TrudeauMustGo were significantly higher on the side of conservative social bots than on the side of conservative humans, this would indicate the presence of an astroturfing strategy aimed at amplifying negative messages against the leader of the LPC. Comparing the ratios between human and social bot accounts allows us to determine whether social bots artificially pushed specific topics during the campaign. Misdirecting and smoke screening strategies will also be analyzed by comparing the various groups' most popular hashtags and mentions. For instance, hashtags and mentions unrelated to the Canadian election might signal an attempt by automated accounts to divert users' attention towards other discussion topics.

Another objective of this research is also to quantify the ability of automated accounts to interact with humans. Therefore, to measure the ability of social bots to interact with authentic accounts, the tweets, retweets, and replies from each partisan side are analyzed separately using the following three metrics developed by Luceri et al. (2019). First, *Retweet Pervasiveness*, which measures the intrusiveness of social bot-generated content in human-generated retweets, is retained. This metric is designed to measure the level of interaction between humans and social bots from the same partisan leaning through retweets. Since the interactions between social bots and humans with the same leaning are compared, it is possible to determine the level of the pervasiveness of social bots from each partisan side. Second, the *Reply Rate* is employed to measure the proportion of replies authentic accounts give to social bots. To obtain the *Reply Rate*, the number of human replies to social bots is divided by the total number of human replies for each partisan-leaning group. This measure makes it possible to evaluate the degree of interaction between human and social bot accounts through direct replies. Third, the *Tweet Success Rate* is used. This other metric allows us to measure the percentage of tweets produced by social bots that generated a reaction from humans by looking at posts retweeted at least once from an authentic account (Luceri et al. 2019, 1009). Once again, this metric is useful for determining the interaction level between social bots and humans and for comparing the percentages between liberal-leaning and conservative-leaning automated accounts. In addition, Chang and Ferrara's (2022, 7) modifications, which divide these metrics by the percentage of automated accounts per partisan group, are also adopted. This provides a better representation of the measures within each partisan group, which, as we will see later, do not produce the same number of messages.

$$\textbf{\textit{Retweet Pervasiveness}}_{\boldsymbol{P}}\ \textbf{\textit{(RTP)}}\textbf{:}\ \frac{\text{no. of human retweets from in-group social bot tweets}}{\text{no. of in-group human retweets}} \qquad (3)$$

$$\textbf{\textit{Reply Rate}}_{\boldsymbol{P}}\ \textbf{\textit{(RTP)}}\textbf{:}\ \frac{\text{no. of human replies to in-group social bot tweets}}{\text{no. of in-group human replies}} \qquad (4)$$

$$\textbf{\textit{Tweet Success Rate}}_{\boldsymbol{P}}\ \textbf{\textit{(TSR)}}\textbf{:}\ \frac{\text{no. of social bot tweets retweeted at least once by an in-group human}}{\text{no. of in-group social bots tweets}} \qquad (5)$$

P = Partisan leaning (conservative = 0; liberal =1)

Lastly, these three measures, previously tested in other publications (e.g., Luceri et al. 2019; Chang and Ferrara 2022), make it possible to detect the presence or absence of coordinated strategies for manipulating public opinion, in addition to quantifying the level of engagement of social bots with humans on Twitter. Combined with the results from Botometer and the partisan leaning prediction model, these measures will provide the necessary tools to test the second hypothesis, which argues that *social bots who lean on the conservative side should interact more with humans than their liberal counterparts*.

2.4. Disinformation Detection

As presented in the first chapter, social media have become channels of choice for spreading disinformation content, notably through automated accounts (Shao et al. 2018). Coordinated disinformation campaigns during electoral periods have been found in the United States (Shao et al. 2018; Vosoughi, Roy and Aral 2018), Canada (Rheault and Musulan 2021) and many European countries (Ferrara 2017; Neudert, Kollanyi and Howard 2017). Detecting disinformation on social media has become increasingly important, and many researchers have put extensive efforts into developing efficient and accurate models (Yang et al. 2019b). To determine whether conservative-leaning social bots propagated more disinformation than their liberal-leaning counterparts, it is first necessary to implement an approach that detects disinformation content in users' online posts.

In the literature, disinformation detection models can be grouped into three categories. The first approach is based on the lexical and syntactic features of online messages or news articles (Guo et al. 2021). Therefore, some studies (e.g., Castillo, Mendoza and Poblete 2013; Kwon et al. 2013; Horne and Adali 2017) relied on characteristics like tweets' length, the number of URLs, or the ratio of positive/negative words to discriminate between true and false information. However, lexical features do not fully capture the characteristics of political disinformation content since they vary according to the country under study. Several studies (e.g., Ito et al. 2015; Maleki et al. 2021) have therefore introduced semantic features, such as topics and hashtags, into their disinformation detection models to capture more local disinformation content circulating on social media. By extracting keywords and hashtags associated with disinformation stories, it is possible to identify these messages (Wu et al. 2016).

In their review of the literature on various disinformation detection methods, Guo and colleagues (2021) identified "fusion-based methods" as an approach to detecting false information on social media. Fusion-based methods combine text information contained in users' posts and information diffusion in social networks to examine disinformation content's propagation. In other words, these methods (e.g., Tacchini et al. 2017; Volvoka and Jang 2018) integrate content-based features, as well as users' interactions, such as retweets and likes, to detect disinformation among social networks. This approach allows the identification of the spreaders of disinformation stories and their receivers. Therefore, fusion-based methods are quite helpful when working with users' content and interaction features.

The third group of models relies on artificial intelligence and deep learning methods to detect disinformation content online. These approaches employ deep neural networks to learn the latent textual or semantic representation of disinformation content on social media (Thota et al. 2018). Various deep learning techniques have been implemented to detect disinformation and conspiracy theories, such as Convolutional Neural Networks (CNN), Graph Convolutional Network (GCN), and Recurrent Neural Networks (RNN). Still, their primary purpose is to extract textual information to identify disinformation content correctly (Guo et al. 2021). Some studies (e.g., Monti et al. 2019) even used GCN-based models to analyze the propagation structure of social media messages and then identify posts conveying disinformation. Deep learning approaches are among the most sophisticated, as they can leverage textual, network and visual information contained in messages posted on online platforms (Wang, Yin and Argyris 2021). However, their implementation can be difficult, and the complexity of these methods can ultimately produce "black box" models, which are not fully interpretable (Kariyappa and Qureshi 2020).

Since this research aims to identify which partisan group shared more disinformation content, the content-based approach is chosen. As presented earlier, this type of model has the advantage of targeting country-specific disinformation content. As this work focuses on the 2021 Canadian federal election campaign, a content-based approach can capture disinformation narratives specific to this political event. More specifically, the approach used in this analysis is built on a list of disinformation keywords and hashtags related to the 2021 Canadian election and the COVID-19 pandemic, which was a central topic in this election campaign. Relevant keywords and hashtags associated with COVID-19 were identified in prior studies (i.e., Chen, Lerman and Ferrara 2020; Green et al. 2020; Mahl, Zeng and Schäfer 2021; Moffitt, King and Carley 2021) and then added to the list of COVID-19 disinformation keywords. Afterwards, Canadian disinformation terms were added to the list of election-related disinformation keywords by going through the top 30,000 most popular hashtags found in the original English set of tweets. A manual inspection of the popular hashtags on Twitter was also conducted to expand the list. The final list contains 164 terms associated with COVID-19 disinformation and 165 terms related to the 2021 Canadian election disinformation for a total of 329 distinct keywords and hashtags (see Appendix A3). This approach makes it possible to capture disinformation tweets and classify them by topic of discussion. For example, #Scamdemic, #Plandemic, #CovidHoax, #CovidFraud, #CovidScam, and #CovidDoesNotExist are all hashtags used to convey disinformation narratives about the veracity of the COVID-19 pandemic. Moreover, #PrivateOToole, #ClimateLockdownsAreNext, #TrudeauCrimeMinister, and #TheGreatReset are other examples of hashtags explicitly used to propagate disinformation narratives related to the 2021 Canadian federal election. To ensure that the disinformation list returned disinformation tweets, a manual validation of 2,000 original flagged disinformation tweets from each topic was conducted. The author's manual validation yielded an average accuracy of over 86%. With the terms included in the list, it was possible to correctly identify the presence of disinformation and the topic being discussed. Therefore, if a tweet matched at least one disinformation term from the COVID-19 terms, it was classified as a COVID-19 disinformation tweet. If a tweet had at least one disinformation term related to the Canadian electoral campaign, it was, in turn, classified as an election disinformation tweet.

To quantify the amount of disinformation content propagated per group for each of the two topics previously identified (i.e., the COVID-19 pandemic and the 2021 Canadian election), this study employs two metrics that were calculated individually on each of the four groups. First is the *COVID-19 Disinformation Ratio*, which measures the proportion of COVID-19 disinformation shared relative to all tweets discussing the COVID-19 pandemic for each group. To identify the COVID-19-related tweets in the dataset, which will be used in the denominator of the operation, a list of 178 terms, which also comprise the COVID-19 disinformation terms previously identified, was used (see Appendix A3 and Appendix A4). Therefore, a COVID-19 tweet is defined as a tweet that contains a minimum of one term from the list. The same calculation was also applied to tweets related to the 2021 Canadian election with the *Canadian Election Disinformation Ratio*. However, given that the data was collected using terms related to the election, the sum of Canadian election tweets corresponds to the total number of tweets in the dataset minus the COVID-19 tweets.

$$\textbf{\textit{COVID-19 Disinformation Ratio}}_{\textbf{\textit{PA}}}\textbf{:} \frac{\text{sum of COVID-19 disinformation tweets}}{\text{sum of COVID-19 tweets}} \quad (6)$$

$$\textbf{\textit{Can. Election Disinformation Ratio}}_{\textbf{\textit{PA}}}\textbf{:} \frac{\text{sum of Can. election disinformation tweets}}{\text{sum of Can. election tweets}} \quad (7)$$

P = Partisan leaning (conservative = 0; liberal =1); A = Account type (human = 0; social bot = 1)

These two disinformation ratios are carried out separately for the four groups. They are beneficial as they compare partisan groups and account types. The measures are also interesting since they allow for a comparison of disinformation tweets related to the coronavirus pandemic with those more generally associated with the 2021 Canadian election. Finally, by combining the results of the Botometer algorithm, the partisan leaning prediction model, and the disinformation content identification approach, it will be possible to assess the third hypothesis put forward in this thesis which states that *social bots who lean on the conservative side should share more disinformation content than their liberal counterparts*.

2.5. Conclusion

The goal of the chapter was to outline the methods selected to test the three hypotheses proposed previously. This chapter briefly reviewed the literature on the methods used to detect automated accounts on social media platforms. In this research, the Botometer algorithm proved to be the optimal choice as it is one of the best-performing bot detection models available and can process several thousand users daily. An overview of different approaches to predicting users' partisan leanings was also presented, along with a justification for using a content-based model to determine the partisan leaning of active users during the 2021 Canadian election. Moreover, the five metrics used to measure social bots' strategies and interactions with humans were also introduced. Lastly, the approach retained to detect disinformation content in users' tweets was explained in this chapter.

Chapter 3. Results

This chapter aims to answer the following research questions from the different analyses performed on a Twitter dataset of more than 1.1 million users and over 11.3 million tweets related to the 2021 Canadian election campaign: In the 2021 Canadian election, were conservative-leaning social bots more numerous than their liberal-leaning counterparts? Were conservative-leaning social bots retweeted and replied to more by humans? And finally, did conservative-leaning social bots spread more disinformation? To answer these questions, the analysis of the results will proceed in four steps. The first section presents the results of the Botometer algorithm, which indicates the number of accounts identified as social bots in the dataset. The second section focuses on the results of the partisan-leaning of users obtained through the content-based model. The third section of this chapter discusses social bots' strategies and interactions with human accounts during the campaign. Lastly, the fourth section presents the results regarding the amount of disinformation conveyed by social bots on Twitter.

3.1. Botometer Results

This research sought to detect active social bot accounts during the 2021 Canadian federal election. To achieve this objective, the state-of-the-art Botometer model was used. After running 1,114,906 Twitter accounts through Botometer, the model returned a bot score for 1,008,520 users, corresponding to 90.46% of the accounts. For example, Galgoczy et al.'s (2022) study, which also employed Botometer to detect automated accounts, also obtained a bot score for about 90% of their users. As detailed in Table 3.1, there are three main reasons for the model's inability to return a bot score for certain accounts. The three reasons are italicized in Table 3.1. and the sum of these accounts corresponds to the total of the 106,386 accounts that could not be assigned a bot score. The three percentages in parentheses for these accounts represent their share among the accounts without a bot score (i.e., 9.54% of all the accounts processed through Botometer) and add up to 100%. The most common is that an account was suspended or deleted. This corresponds to just over 60% of users without a bot score. Unfortunately, it is impossible to know why these accounts were suspended or deleted from the platform. It is nonetheless possible to assume that among these users, there could be a significant share of social bots. Indeed, Twitter sometimes conducts verifications of accounts, and those, such as automated accounts that violate the rules regarding the integrity of the platform, or use a fake identity, can have their account suspended.[19] For example, after the election of U.S. President Donald Trump in 2016, *The Washington Post* revealed that Twitter suspended more than 70 million accounts in the months of May and June 2017 (Timberg and Dwoskin 2018). Since the revelations about how Russia used inauthentic social media accounts to interfere in the 2016 U.S. presidential election mainly by launching a large-scale disinformation campaign, Twitter adopted a vast campaign against bot accounts to "promote healthy conversations on the platform" (Timberg and Dwoskin 2018). After the January 6th Capitol Hill Riot, Twitter suspended more than 70,000 users, including Donald Trump, to stop the spread of content related to QAnon conspiracy theories (Romm and Dwoskin 2021). After conducting verifications, Twitter announced that they found many instances where a single individual was operating numerous accounts, showing signs of automation behind the spread of disinformation content on the platform (Sardarizadeh 2021). Regarding the most recent data, Twitter confirmed in July 2022 that it was removing over one million automated accounts per day from its platform (Dang and Paul 2022).

[19] See https://help.twitter.com/en/rules-and-policies/twitter-rules for a complete list of Twitter's terms and rules, as well as possible penalties for users who violate these rules.

Table 3.1. Summary of Bot Scores' Availability

Statistics	Count
Accounts with a bot score	1,008,520 (90.46%)
Accounts without a bot score	106,386 (9.54%)
• *Suspended or deleted account*	64,717 (60.83%)
• *Private account*	41,425 (38.94%)
• *No tweets in the account's timeline*	244 (0.23%)

Furthermore, almost 40% of the users were not assigned a bot score because they turned their Twitter profile into private mode, which blocks the algorithm from reading their tweets. The remaining accounts, which represent less than 1%, deleted their tweets from their profiles' timelines and could not be processed by the model. Since Botometer could not provide bot scores for 106,386 users, they were removed along with their 3,458,847 associated tweets. Therefore, the remainder of this research is based on a sample of 1,008,520 active Twitter accounts with bot scores and the 7,902,734 tweets they produced.

Table 3.2. Summary of Users and Tweets' Distribution

Statistics	User Count	Tweet Count
Humans	881,774 (87.43%)	6,811,859 (86.20%)
Social bots	126,746 (12.57%)	1,090,875 (13.80%)
Total	1,008,520 (100.00%)	7,902,734 (100.00%)

As explained in the literature review on automated accounts' detection models, a threshold of 0.7 was chosen to discriminate between social bot and human accounts. According to this choice, a total of 126,746 accounts were classified as social bots, corresponding to 12.57% of active users in the dataset. As shown in Table 3.2, out of the sample of 7,902,734 tweets, social bots generated 13.80% of the total volume of messages around the electoral campaign discussions. For comparison, Rheault and Musulan (2021) estimated that social bots in the 2019 Canadian election represented roughly 8% of users engaged in online discussions and 13% of the total volume of tweets. It is important to note that although the proportions are slightly different between the two elections, the share of tweets produced by automated accounts remains higher than those classified as social bots. Human accounts, which represented more than 87% of the dataset, generated roughly 86% of the total volume of tweets. In other words, social bots were more active and generated a higher proportion of tweets, considering their number, than authentic accounts.

3.2. Party Leaning Classification

Another essential step in determining whether one partisan group shared more disinformation content than the other during the 2021 Canadian federal election was to identify the partisan leaning of users involved in discussions around the election. To that end, this work relied on a supervised model built and trained to predict the partisan leaning of users in the Canadian context. This content-based predictive model composed of two distinct classifiers (i.e., the profile classifier and the activity classifier) was able to assign a partisan leaning to a total of 770,970 accounts with a bot score, which represents 76.45% of all the users in the dataset. When considering only users who were both given a bot score and a partisan leaning, social bots represented close to 11% of the sample and generated 13.84% of all tweets. It is important to specify here that when the users' true labels were available, they were retained as their partisan leaning. However, when the true labels were unavailable, the model provided the partisan leaning of the users. For those with explicitly political profile descriptions, the activity classifier was applied instead for the rest of the users. On the one hand, the profile classifier made it possible to assign a partisan leaning to 187,447 users, of which 132,010 were classified as liberal-leaning and 55,437 as conservative-leaning.[20] On the other hand, the activity classifier attributed partisan leanings to 582,709 users, among which were 197,974 liberal-leaning and 384,735 conservative-leaning. In the end, combining the true labels

[20] The profile classifier marked around 444,000 users with an explicitly political profile description as "NO PARTY". For these users, the classification results of the activity classifier were retained instead.

from the *Media Ecosystem Observatory* survey and the two classifiers allowed us to assign a partisan leaning to 64% of social bot accounts and 63.24% of human accounts. The model could not attribute a partisan leaning to the entirety of the users in the dataset, primarily because many of them did not have any profile description or had too few tweets for the activity classifier to infer an accurate partisan leaning.

Tables 3.3 and 3.4 show the distribution of users and tweets in each of the four groups, respectively. Looking at Table 3.3, we can see that conservative-leaning accounts were more numerous overall than their liberal-leaning counterparts. This was also the case when looking specifically at social bot accounts. Indeed, social bots who leaned on the conservative side represented more than 54% of the accounts, while this proportion reached a little over 45% for liberal-leaning social bots.

Table 3.3. Summary of Users' Distribution Per Partisan Group

Statistics	Liberal-Leaning	Conservative-Leaning	Total
Humans	293,970 (42.61%)	395,870 (57.39%)	689,840
Social Bots	36,636 (45.16%)	44,494 (54.84%)	81,130
Total	330,606	440,364	770,970

This disparity between the number of conservative and liberal-leaning accounts is consistent with the findings reported by Rheault and Musulan (2021). Indeed, in their study of the 2019 Canadian federal election, they observed a higher percentage of social bots among conservative-leaning clusters, especially the PPC, which the authors claimed was the party that benefited the most from the support of social bot accounts during the campaign (Rheault and Musulan 2021, 333). It is, therefore, interesting to see that both in terms of numbers and proportions, social bots who leaned towards the conservative side were more prevalent than their liberal counterparts, but also that the same trend was noticeable on the side of authentic accounts, where conservative-leaning users made up an even greater proportion of the accounts, at 57.39%. Besides, when examining the proportion of social bots within each partisan group, it was possible to note that social bots represented roughly 11% of the sample of liberal-leaning supporters. In contrast, they represented over 10% of the sample of conservative-leaning supporters. These results are slightly different from those of Bessi and Ferrara (2016), who found during the 2016 U.S. presidential election that social bots made up more than 12% of the sample of Republican supporters, while they made up approximately 9% of the sample of Democrat supporters. Even if the proportions of social bots within each partisan group are reversed from what was observed in the U.S. context, they are still in a similar range. They indicate that automated accounts do not form a large proportion of the supporters of these two groups.

Table 3.4. Summary of Tweets' Distribution Per Partisan Group

Statistics	Liberal-Leaning	Conservative-Leaning	Total
Humans	3,532,482 (52.61%)	3,182,139 (47.39%)	6,714,621 (100.00%)
Social Bots	574,802 (53.28%)	504,131 (46.72%)	1,078,933 (100.00%)
Total	4,107,284 (52.70%)	3,686,270 (47.30%)	7,793,554 (100.00%)

On another note, Table 3.4 shows the distribution of tweets for each group. It is relevant because it complements the results in Table 3.3 by providing information on the number of tweets produced by each group, which is indicative of users' involvement in discussions surrounding the election. One of the most interesting results from Table 3.4 is that although they were more numerous, accounts that leaned on the conservative side, both among humans and social bots, produced a smaller proportion of tweets than those that leaned on the liberal side. In this sense, even though they formed about 45% of the sample of social bots, liberal-leaning users generated more than 53% of all tweets from social bot accounts. In contrast, conservative-leaning accounts represented nearly 55% of all social bots, but they only produced about 46% of all tweets from these types of accounts. Moreover, tweets associated with liberal-leaning social bots represented approximately 14% of all tweets from the liberal-leaning sample. In comparison, tweets from conservative-leaning social bots formed 13.68% of all tweets from the conservative-leaning sample. These results suggest that liberal-leaning social bots were more active and involved in Twitter conversions during the 2021 Canadian election. A higher level of activity in terms of tweets and retweets can therefore explain why fewer accounts were able to generate more messages.

The findings discussed above support the first hypothesis, which argued that *social bots leaning on the conservative side should have been more numerous than liberal-leaning ones*. In absolute and percentage terms, conservative-leaning social bots were found to outnumber liberal-leaning social bots during the 2021 Canadian election. This was also true among human users, where conservative-leaning supporters were more numerous than their liberal counterparts. However, it is important to nuance the conclusions drawn from the respective quantity of users within each partisan group since the analysis of the number of tweets showed that the share of tweets produced by liberal-leaning social bots was greater than that of conservative-leaning social bots. In other words, while the first hypothesis was confirmed, the results also indicate that social bot accounts that leaned towards the liberal side had a higher activity level during the election period. This difference in the number of tweets generated by liberal and conservative-leaning social bots could be explained by Twitter's removal of content that does not comply with the platform's rules regarding, for example, civic integrity or COVID-19 misleading information. For instance, Twitter reported that from July to December 2021, more than 4 million tweets were removed from the platform because they violated Twitter Rules.[21] While it is not possible to know the partisanship of the tweets that were removed from the platform, it can nevertheless be assumed that a significant

[21] See Twitter's Rules Enforcement Report for more details: https://transparency.twitter.com/en/reports/rules-enforcement.html#2021-jul-dec.

share of them was associated with conservative-leaning users, since COVID-19 disinformation was predominantly found on the side of conservatives, especially in the United States (Roozenbeek et al. 2020; Thelwall, Kousha and Thelwall 2021).

3.3. Social Bots' Strategies During The 2021 Campaign

The previous chapters have shown that social bots have used various strategies in past political campaigns, such as astroturfing, smoke screening, and misdirecting in order to increase the visibility of some topics and political actors or influence public opinion in a particular direction. Therefore, one aspect of this work was determining if such strategies could be detected in the 2021 Canadian federal campaign. To ascertain whether an astroturfing strategy conducted by automated accounts took place on Twitter, the abovementioned astroturfing ratios were applied separately to the most popular keywords and hashtags for each group (i.e., liberal-leaning human, conservative-leaning human, liberal-leaning social bot, and conservative-leaning social bot). Tables 3.5 and 3.6 show in detail the astroturfing ratios for human and social bot accounts according to their partisan leanings.[22] We note that the most popular keywords and hashtags between liberal-leaning and conservative-leaning accounts, both for humans and social bots, are quite similar. It is particularly relevant to note in Table 3.5 that the highest astroturfing ratio obtained by humans, both for liberal and conservative-leaning accounts, when it comes to hashtags is 0.084. In comparison, the highest hashtag astroturfing ratio attained by liberal-leaning social bots is 0.074, while this ratio reaches 0.077 for conservative-leaning social bots (see Table 3.6).

[22] For more information on the denominators used in the astroturfing ratio formulas for each of the four groups, see Appendix B.

Table 3.5. Astroturfing Ratios for Humans

Liberal-Leaning			Conservative-Leaning		
Top 10 Hashtags	**Count**	**Ratio**	**Top 10 Hashtags**	**Count**	**Ratio**
#cdnpoli	295,636	0.084	#cdnpoli	267,347	0.084
#elxn44	246,453	0.070	#elxn44	237,567	0.075
#canada	85,384	0.024	#canada	82,270	0.026
#voteppc	60,270	0.017	#voteforhumanrights	28,317	0.009
#ppc	37,052	0.010	#voteppc	18,749	0.006
#onpoli	15,703	0.004	#covid19	17,840	0.006
#ableg	15,362	0.004	#ableg	17,011	0.005
#cpc	14,763	0.004	#ppc	14,647	0.005
#covid19	14,698	0.004	#everychildmatters	12,629	0.004
#lpc	14,176	0.004	#cpc	12,375	0.004
Top 10 Keywords	**Count**	**Ratio**	**Top 10 Keywords**	**Count**	**Ratio**
trudeau	963,197	0.273	trudeau	556,825	0.158
canada	361,311	0.102	reconciliation	440,955	0.125
cdnpoli	324,974	0.092	otoole	302,076	0.086
otoole	314,114	0.089	canada	301300	0.085
elxn44	265,654	0.075	cdnpoli	292271	0.083
reconciliation	251,862	0.071	elxn44	254,839	0.072
justin	242,135	0.069	vote	179,728	0.051
mp	237,530	0.067	bill	178,76	0.051
election	234,569	0.066	election	178,495	0.051
vote	193,052	0.055	day	169,981	0.048

When looking at the keyword astroturfing ratios, once again the highest scores are attributable to humans and not to social bot accounts. Nevertheless, some keyword astroturfing ratios for social bots are higher than those of humans with the same partisan leaning. For instance, the keyword “reconciliation” has a greater ratio among social bots for each partisan-leaning group than their human counterparts. This difference is particularly noteworthy when comparing liberal-leaning humans to liberal-leaning social bots; the ratio value of social bots for this keyword was almost twice that of humans. Nevertheless, the difference between humans’ and social bots’ ratios for all hashtags and keywords except “reconciliation” was not very substantial. It is, therefore, possible to claim that social bot accounts, notably those that leaned on the liberal side, tried to increase the visibility of content related to the issue of reconciliation with Indigenous Peoples by posting more than 250,000 tweets. However, the small overall difference between the astroturfing ratios of humans and social bots does not suggest the presence of a broad, coordinated astroturfing strategy to manipulate the Canadian public online.

In sum, the low ratios, both in terms of keywords and hashtags, do not allow us to conclude that there was an astroturfing strategy on the part of social bots, whether they leaned on the liberal or conservative side. By observing the top ten hashtags and keywords most frequently used by social bots in both partisan-leaning groups and comparing them to their human counterparts, it was impossible to observe elements that could indicate an astroturfing strategy during the 2021 election campaign.

Table 3.6. Astroturfing Ratios for Social Bots

Liberal-Leaning			Conservative-Leaning		
Top 10 Hashtags	**Count**	**Ratio**	**Top 10 Hashtags**	**Count**	**Ratio**
#canada	42,379	0.074	#canada	38,811	0.077
#cdnpoli	27,415	0.048	#cdnpoli	28,189	0.056
#elxn44	21,094	0.037	#elxn44	20,081	0.040
#ppc	14,213	0.025	#americans	6,939	0.014
#voteppc	5,130	0.009	#usa	5,552	0.011
#usa	5,068	0.009	#voteforhumanrights	3,978	0.008
#australia	4,759	0.008	#uk	3,789	0.008
#uk	4,433	0.008	#ppc	2,880	0.007
#seo	4,268	0.007	#voteppc	2,699	0.006
#ontario	3,311	0.006	#news	2,426	0.005
Top 10 Keywords	**Count**	**Ratio**	**Top 10 Keywords**	**Count**	**Ratio**
trudeau	126,985	0.221	reconciliation	83,863	0.166
canada	86,623	0.151	canada	71,984	0.143
reconciliation	68,477	0.119	trudeau	69,047	0.137
otoole	44,991	0.078	bill	40,131	0.080
justin	37,831	0.066	otoole	39,743	0.079
election	32,852	0.057	cdnpoli	31,599	0.063
cdnpoli	30,414	0.053	justin	23,846	0.047
vote	23,496	0.041	elxn44	21,956	0.044
erin	22,704	0.039	erin	21,712	0.043
elxn44	22,626	0.039	vote	20,507	0.041

In addition, the presence of smoke screening or misdirecting strategies was also investigated. As mentioned by Brachten et al. (2017, 7) in their study of the 2017 German state election, the analysis of these two strategies is combined since they are only somewhat distinguishable from each other. As such, to perform the analysis like Brachten et al. (2017), hashtags were examined separately for humans and social bots. This work also included the occurrence of mentions to see if social bots directed their messages toward accounts unrelated to politics or the news media. This strategy allowed to identify the most popular hashtags and mentions within the respective groups and to see if the most frequently used by social bots varied from those of authentic accounts, as well as to examine whether the hashtags and mentions chosen by social bots were indeed related to the Canadian election and the issues associated with it. In this sense, a high occurrence of hashtags and mentions unrelated to Canadian politics on the part of automated accounts could indicate a smoke screening strategy, which involves hijacking popular hashtags to disseminate irrelevant messages. Figure 3.7 displays the twenty hashtags most frequently used by social bots and human accounts, and Figure 3.8 shows the accounts that were mentioned the most by social bots and humans. On the one hand, the top three most popular hashtags are the same for both human and social bot accounts. However, even if there is some overlap in the hashtags most used by the two types of accounts, we note that the most popular hashtags on the human side are all directly related to the Canadian election. In contrast, on the side of social bots, hashtags are much more general by focusing mainly on the names of other countries.

Table 3.7. Top 20 Hashtags Generated by Human and Social Bot Accounts

Human	Social Bot
#cdnpoli	#canada
#elxn44	#cdnpoli
#canada	#elxn44
#voteppc	#ppc
#ppc	#usa
#covid19	#uk
#ableg	#americans
#voteforhumanrights	#voteppc
#onpoli	#autralia
#cpc	#seo
#abpoli	#covid19
#lpc	#ontario
#trudeau	#voteforhumanrights
#cdmedia	#germany
#ndp	#france
#bcpoli	#news
#everychildmatters	#ableg
#forwardforeveryone	#toronto
#elexn44	#alberta
#trudeaumustgo	#sem

A manual inspection of thousands of tweets shared by social bots containing hashtags about other countries (e.g., #usa, #uk, #amerians, #australia, #germany, and #france) found that these were mostly used in combination with the hashtag #canada to share international news related to politics or the latest status of the COVID-19 pandemic. Therefore, social bots did not use hashtags referring to other countries to divert users' attention away from content related to the election or flood Twitter with unrelated and irrelevant messages but to share international policy news with a Canadian audience.

To identify a misdirecting strategy, which consists of using context-related hashtags to convey messages unrelated to the hashtag being used, manual verification by the author of three random samples of 2,000 messages relayed by social bots was conducted.[23] The focus was on the three most common hashtags (e.g., #canada, #cdnpoli, and #elxn44), which could be well suited to be part of a misdirecting strategy given that they are generic hashtags related to the Canadian election and happen to be widely spread by authentic accounts. The assessment of social bots' tweets did not find a consistent and coordinated set of messages to redirect users' attention to other topics. A few marginal tweets unrelated to the hashtags used were found, but these were not linked and constituted isolated incidents. For example, a marketing campaign for the company Dingtone used the hashtags #Canada and #Americans to promote their mobile application (see Figure 3.1). These tweets included the use of the hashtag #Canada. Still, since it was a promotional campaign, it is impossible to speak of a strategy to divert users' attention away from Canadian political issues.

Figure 3.1. Example of Unrelated Tweet Found in the Dataset

[23] The first random sample contained tweets relayed by social bots with the hashtag #canada, the second with the hashtag #cdnpoli, and the third with the hashtag #elxn44.

It is, however, relevant to note that despite the low percentage of votes obtained by the PPC, the hashtags related to this party have nevertheless circulated abundantly during the campaign. Indeed, #VotePPC and #PPC are, respectively, the fourth and fifth most frequently used hashtags by humans, thus surpassing the hashtags of all the other parties and even the #covid19. The overrepresentation of hashtags related to the PPC could be explained by the fact that newer parties now rely heavily on social media platforms to reach new adherents and spread their message and ideas (Auter and Fine 2016). Moreover, the PPC campaigned heavily on its opposition to sanitary measures to counter COVID-19 and ran on a strong anti-immigration platform, two themes that provoked a great deal of reaction and made it possible to mobilize a particular segment of the electorate (Somos 2021).

Table 3.8. Top 20 Mentions Generated by Human and Social Bot Accounts

Human	Social Bot
@JustinTrudeau	@JustinTrudeau
@erinotoole	@erinotoole
@MaximeBernier	@MaximeBernier
@brianlilley	@CPCHQ
@liberalparty	@liberalparty
@CPCHQ	@Dingtone
@TheJagmeetSingh	@brianlilley
@kinsellawarren	@SenSanders
@sunlorrie	@theJagmeetSingh
@WaytowichNeil	@WaytowichNeil
@ezralevant	@RBReich
@NDP	@ThePlumLineGS
@DrJacobsRad	@krismeloche
@TrueNorthCentre	@BernieSanders
@SenSanders	@sunlorrie
@CTVNews	@ezralevant
@BernieSanders	@SenWarren
@CBCNews	@kinsellawarren
@TorontoStar	@kylegriffin1
@krismeloche	@DeanWinnipeg

On the other hand, when looking at the most popular mentions generated by humans and social bots we notice similarities in the accounts most mentioned by each group. On both sides, all the most mentioned Twitter accounts during the electoral period were related to Canadian politics through their role as politicians, journalists or even bloggers. The only exception is the @Dingtone account, which is the Twitter account of a mobile application that offers unlimited calling and texting. The latter is one of the top twenty most mentioned accounts, as its promotional message containing the hashtag #Canada was widely shared on Twitter. Again, a single account unrelated to Canadian politics does not indicate a smoke screening or misdirecting strategy, but in this case, an online marketing campaign. In sum, the analysis and comparison of hashtags and mentions produced by human and social bots accounts did not indicate coordinated smoke screening or misdirecting strategies during the 2021 Canadian election campaign. Therefore, the results presented do not suggest that social bots engaged in a coordinated campaign to influence the Canadian electorate on Twitter during the 2021 election, as no sign of astroturfing, smoke screening and misdirecting was found in the dataset of close to 8 million tweets.

3.4. Interactions With Human Accounts

Another important aspect of this work is to evaluate whether, in the Canadian context, social bots leaning from different partisan groups interacted more with humans via retweets and replies. By applying the *Retweet Pervasiveness*, *Reply Rate*, and *Tweet Success Rate* metrics first introduced by Luceri et al. (2019), we were able to measure the level of interaction of liberal-leaning and conservative-leaning social bots with humans. Table 3.9 summarizes the results of these three interaction measures for the 2021 campaign period. It reveals that liberal-leaning social bots scored higher on RTP and RR than their conservative-leaning counterparts for both absolute and relative metrics. This means, on the one hand, that tweets produced by liberal-leaning social bots were retweeted with greater propensity by human accounts from their group than tweets generated by conservative-leaning social bots. In other words, social bots that leaned on the liberal side were more effective at involving human accounts in their discussions via retweets than what can be observed from those that leaned on the conservative side. On the other hand, liberal-leaning social bots were also more effective than conservative-leaning ones at interacting with their human counterparts through replies. Overall, human engagement with social bot generated-content remained limited during the 2021 Canadian campaign for both groups. Human retweets from social bots' original posts only account for 2.8% of all liberal-leaning humans' retweets, compared to 1.7% for the conservative-leaning side. In addition, the percentage of replies given by humans to social bot accounts represented 0.2% of the entirety of liberal-leaning humans, while this percentage was approximately 0.1% for conservative-leaning ones. In their study of the 2018 U.S. midterm elections, Luceri et al. (2019) found that conservative humans interacted with social bots the most, with an RTP score reaching 25.6% and an RR score of 15.5%. It, therefore, seems that in the highly polarized political and social context of the United States, conservative social bots hold a more central position in the social network than what has been observed in Canada.

Table 3.9. Social Bot Effectiveness in Human Engagement

	Liberal-Leaning		**Conservative-Leaning**	
Metric	Absolute	Relative	Absolute	Relative
RTP	0.028	0.250	0.017	0.166
RR	0.002	0.018	0.001	0.010
TSR	0.188	1.699	0.177	1.750

Absolute stands for the total human-social bot interactions overall human interactions. Relative is normalized over the number of social bots (Chang and Ferrara 2022).

Although the overall interaction of human users with social bots' content was not extensive, it is possible to note via the TSR scores that the percentage of original tweets produced by social bots that got at least one retweet by an authentic user is relatively important with 18.8% for liberal-leaning social bots and 17.7% for conservative-leaning ones. However, when looking at the relative measure of TSR, normalized over the percentage of social bots detected among each partisan group, we find that the conservative-leaning social bots receive a higher score. These results suggest that human users interacted less with conservative-leaning social bots through retweets and replies. Still, the latter were more successful than their liberal counterparts in reaching their target audience. In other words, the content generated by social bots could have been more concentrated among a particular subset of Twitter users and therefore have more original messages retweeted at least once by a human account. This could explain the significant differences between the RTP and the RR scores and the results from the TSR. According to the findings provided by the three different measures of interaction, the second hypothesis cannot be confirmed. Even if social bots who leaned on the conservative side obtained a better *Tweet Success Rate* score than those who leaned on the liberal side, they were not retweeted or replied to more during the Canadian federal election.

Finally, the various interaction measures adopted in this work showed that conservative-leaning social bots did not interact more with their human counterparts during the Canadian election campaign than social bots from the liberal-leaning side. These results are inconsistent with what was found by Luceri et al. (2019) in the 2018 U.S. midterm elections, where conservative social bots were more effective than liberal social bots at interacting with humans, either through retweets or replies. However, the scores obtained in this work are closer to those of Chang and Ferrara's (2022) comparative analysis of social bots' behaviors on Twitter regarding the COVID-19 pandemic between January 21, 2020, and April 1, 2021. The authors found that liberal social bots were much better than conservatives at interacting with humans via retweets. One explanation as to why liberal-leaning social bots were better at involving humans could be because their accounts were more sophisticated (i.e., accounts with a profile image, description, and followers) and perceived by authentic users as humans and, therefore, more credible (Kenny et al. 2022). In other words, accounts which appear more trustworthy could be more easily confused with real users and elicit more interaction from humans (Chang and Ferrara 2022).

3.5. Disinformation Results

In this research, an approach based on hashtags and keywords was employed to detect disinformation associated with the 2021 Canadian election and the COVID-19 pandemic. Before proceeding to detect disinformation content, it was first essential to separate the tweets discussing COVID-19 specifically from those dealing more generally with the election. As such, tweets were filtered and divided according to a list of hashtags and keywords related to the pandemic (see Appendix A4). Table 3.10 summarizes the distribution of tweets per topic and per group. On the one hand, it shows that humans discussed the COVID-19 pandemic more than social bot accounts and that the proportion of tweets focused on COVID-19 was higher among liberal-leaning groups. On the other hand, it reveals that tweets discussing COVID-19 represented a subgroup of more than 11% of the total 7,793,558 tweets about the 2021 Canadian election. These results indicate that the COVID-19 pandemic was an important topic during the election, although it did not occupy the most significant part of the online conversation.

Table 3.10. Summary of Tweets Per Topic

	Liberal-Leaning		**Conservative-Leaning**		**Total**
Topic	Human	Social Bot	Human	Social Bot	
COVID-19	434,622 (12.30%)	59,532 (10.36%)	354,621 (11.14%)	49,482 (9.82%)	898,257 (11.53%)
2021 Canadian Election	3,097,861 (87.70%)	515,271 (89.64%)	2,827,519 (88.86%)	454,650 (90.18%)	6,895,301 (88.47%)
Total	3,532,483 (100.00%)	574,803 (100.00%)	3,182,140 (100.00%)	504,132 (100.00%)	7,793,558 (100.00%)

Table 3.11 presents the results obtained using the content-based approach. Overall, disinformation tweets represented a very small number of the content that circulated on Twitter during the electoral campaign, with a total of 88,675 tweets. For example, disinformation tweets from both topics only formed a little over 1% of all tweets contained in the dataset. As shown in Table 3.11, in absolute terms, disinformation tweets related to the 2021 Canadian election were more frequent than those about COVID-19, which holds for both humans, social bots, and liberal-leaning and conservative-leaning accounts. Overall, 2021 Canadian election disinformation tweets represented about 70% of all disinformation tweets in the dataset. However, when analyzing the proportion of disinformation tweets per topic, it is possible to note that the percentages of disinformation tweets attributable to social bots are higher among the COVID-19 category than for the 2021 Canadian election. In other words, disinformation content produced and spread by social bot accounts for both partisan leanings was more prevalent among the COVID-19 topic. For instance, social bots that leaned on the liberal side were responsible for roughly 7% of all COVID-19 disinformation but were only associated with 5.85% of the 2021 Canadian election disinformation. This observation is the same for conservative-leaning social bots. At the same time, they made up approximately 4% of the Canadian election sample, they were responsible for about 5.5% of COVID-19 disinformation tweets. Overall, it was still human accounts that shared nearly 90% of disinformation tweets during the election period, with respectively 56.26% on the liberal side and 33.07% on the conservative side. In comparison, liberal-leaning social bots spread 6.20% of all disinformation tweets, while conservative-leaning social bots contributed to only 4.47%.

Table 3.11. Summary of Disinformation Tweets Per Topic

	Liberal-Leaning		Conservative-Leaning		Total
Topic	Human	Social Bot	Human	Social Bot	
COVID-19	14,285 (53.66%)	1,872 (7.03%)	9,014 (33.84%)	1,463 (5.49%)	26,634 (100.00%)
2021 Canadian Election	35,603 (57.39%)	3,630 (5.85%)	20,308 (32.73%)	2,500 (4.03%)	62,041 (100.00%)
Total	49,888 (56.26%)	5,502 (6.20%)	29,322 (33.07%)	3,963 (4.47%)	88,675 (100.00%)

Additionally, an analysis focused on users that shared at least one post containing disinformation revealed that this type of content was highly concentrated among a limited number of users. Therefore, only 27,264 distinct users were responsible for the 88,675 disinformation tweets identified in this work. As such, only 3.53% of users in the dataset engaged with disinformation content through direct messages or retweets. Of those users, 2,956 were labeled as social bots; meaning that 10.85% of accounts that propagated disinformation content were social bots. This indicates that the vast majority of political disinformation content produced and shared during the 2021 election campaign is attributable to authentic accounts.

Table 3.12. Summary of Disinformation Ratios

	Liberal-Leaning		Conservative-Leaning	
Metric	Human	Social Bot	Human	Social Bot
COVID-19	0.033	0.031	0.025	0.030
2021 Canadian Election	0.011	0.007	0.007	0.005

In addition, when looking at the disinformation ratio for the COVID-19 pandemic and the 2021 Canadian election respectively, we observe that, surprisingly, liberal-leaning humans are the ones who hold the highest ratios for both topics. It is also interesting to note that liberal-leaning social bots have higher disinformation ratios in each of the two categories when compared to their conservative-leaning counterparts. However, the differences in the disinformation ratios of the two partisan-leaning groups are not very substantial. For example, the difference between liberal and conservative-leaning social bots for the COVID-19 pandemic ratio is 0.001, while it is 0.002 for the Canadian election ratio.

Now that the results regarding the volume of disinformation tweets detected and their distribution within the different partisan-leaning groups have been exposed, it is relevant to explore users' most discussed disinformation narratives further. Figures 3.2 and 3.3 illustrate examples of disinformation tweets that circulated on Twitter during the campaign. Figure 3.2 focuses on disinformation tweets about the 2021 Canadian election, while Figure 3.3 highlights disinformation tweets related to COVID-19. Prevalent narratives regarding the 2021 Canadian election primarily, but not exclusively, involved Justin Trudeau, the outgoing Prime Minister. For example, conspiracy theories such as the "Great Reset", the "New World Order", and "Agenda 2030", claimed that Justin Trudeau and other world leaders were underway to deprive citizens of their civil and economic freedoms in order to install a socialist-communist totalitarian world government (Rectenwald 2022).

Disinformation narratives aimed at undermining the integrity of the Canadian electoral system and the legitimacy of the election results were also spread online during the campaign. Interestingly, these disinformation stories were adapted to the Canadian context, but they are primarily inspired by political disinformation that circulated during the latest U.S. presidential election. Therefore, disinformation stories questioning the integrity of the Canadian election were copied from the U.S. voter fraud narratives (Karadeglija 2021). For instance, messages about voter suppression, the Dominion Voting Systems, as well as tweets directing voters to bring their own pens to decrease the risk of election fraud were first promoted at the time of the 2020 U.S. election, most notably by Donald Trump and high-ranking Republicans officials (Feldman 2020; Bridgman et al. 2022).

Figure 3.2. Examples of Canadian Election Disinformation Tweets

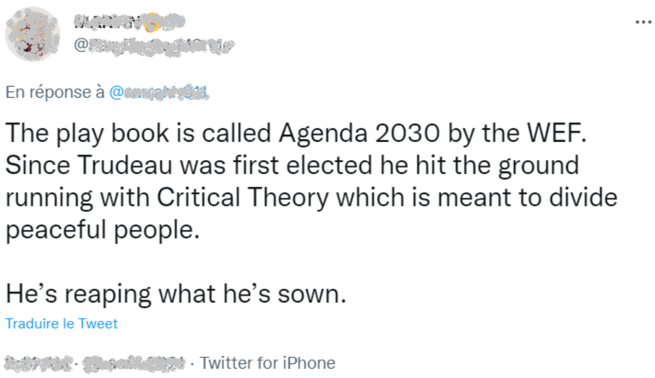

(a) Agenda 2030 disinformation tweet

(b) The Great Reset disinformation tweet

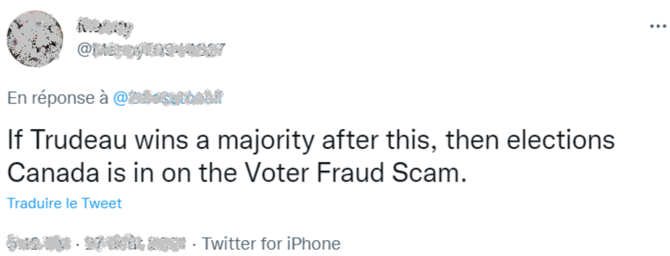

(c) Voter fraud disinformation tweet

This work also found that similar claims of widespread electoral fraud were picked up in Canada and conveyed through Twitter. However, this type of disinformation narrative remained limited to a small number of accounts and tweets within the dataset. Hence, claims of coordinated voter fraud in Canada did not receive as much visibility as in the United States, where known politicians and cable news channels (e.g., Fox News' opinion branch and NewsMax) also propagated such disinformation (Chotiner 2020). In Canada, apart from the leader of the PPC, Maxime Bernier, politicians and traditional news media did not endorse voter fraud conspiracy theories, which may have helped to limit their salience (Lurie 2021).

Figure 3.3. Examples of COVID-19 Disinformation Tweets

(a) False pandemic disinformation tweet

(b) Vaccine disinformation tweet

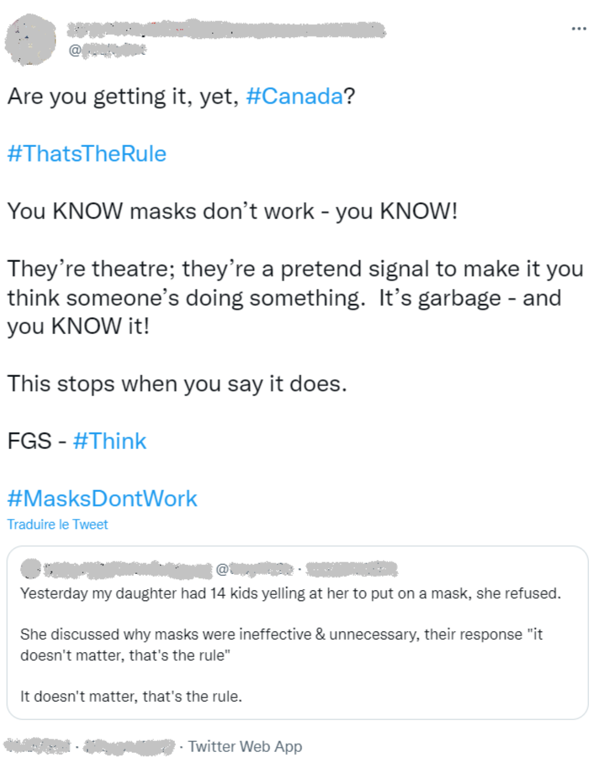

(c) Mask disinformation tweet

Moreover, the disinformation narrative about Erin O'Toole, the leader of the CPC, wanting to privatize the Canadian healthcare system also represented an important part of disinformation tweets that circulated on Twitter during the 2021 electoral campaign. This disinformation story was most prevalent among authentic and liberal-leaning accounts, which may be a factor as to why liberal-leaning accounts shared more disinformation content during the campaign than conservative-leaning ones.[24] In addition, the popularity of this narrative could be explained by the fact it was Chrystia Freeland, a LPC member of Parliament and Canada's Deputy Prime Minister, that originally shared an edited clip of Erin O'Toole answering questions about privatized healthcare (Bridgman et al. 2022). Even though Twitter marked Freeland's tweet as "manipulated media" (see Figure 3.4), the Prime Minister also retweeted the video and attacked his opponent on this subject during different speeches on the campaign trail (Burke 2021). The fact that such high-profile politicians endorsed a story of disinformation could have helped it gain credibility in the eyes of many people - especially left-leaning voters, as they are known in Canada for strongly supporting the free and public healthcare system (Dufresne, Jeram and Pelletier 2014).

Regarding disinformation stories about COVID-19, they mainly focused on claims about the veracity of the pandemic and the sanitary measures to counter the virus, like vaccines and masks. Disinformation stories about the pharmaceutical industry (i.e., Big Pharma) and the so-called "globalist" elite were also among the most important disinformation stories about COVID-19 that were shared on Twitter during the 2021 Canadian electoral campaign. These two disinformation narratives coexisted and were shared in the Canadian context. On the one hand, the first set of narratives conveys the idea that the coronavirus pandemic was an outright hoax or a

[24] For more information on each group's most popular keywords in disinformation tweets, see Appendix C1 to Appendix C4.

fabricated lie by government authorities in order to suppress citizens' freedoms (Grimes 2021). On the other hand, the second set of narratives promoted the idea that the COVID-19 virus was real but was engineered and spread on purpose. Therefore, the global elite would have implemented this scheme to enrich the pharmaceutical industry further and force mandatory vaccination campaigns (Grimes 2021).

Overall, the limited number of tweets relaying disinformation about COVID-19 vaccines could be explained on the one hand by the fact that at the time the dataset used to conduct the analyses was collected, Twitter had in place a strict policy regarding COVID-19 disinformation. For instance, during the period during which the 2021 Canadian campaign took place, Twitter removed tweets conveying "false claims about the virus, or the safety and effectiveness of vaccines" (Klepper 2022). On the other hand, the limited reach of such disinformation stories could also be explained by the fact that very few politicians in Canada questioned the safety and necessity of vaccines to address the pandemic. Even Pierre Poilievre, who replaced Erin O'Toole as head of the CPC in 2022, and who was one of the most prominent critics of the Liberal government's management of the pandemic, encouraged people to get vaccinated and strongly attacked Justin Trudeau for the delays in delivering vaccines to Canadians (Parliament of Canada 2021). However, with the "Freedom Convoy" unfolding at the beginning of 2022, Pierre Poilievre became one of the most popular conservative figures to protest the federal government from imposing vaccine mandates on federal workers and the travelling public (Tasker 2022). Therefore, the amount of COVID-19 disinformation content could have been reduced due to the application of Twitter's policy and the lack of support from Canadian politicians.

Figure 3.4. Chrystia Freeland's Manipulated Tweet

In short, the third hypothesis proposed in this research is rejected. Contrary to what was expected in light of previous studies that looked at disinformation on social media (e.g., Shao et al. 2018; Chang et al. 2021), it was, in fact, liberal-leaning social bot accounts that propagated disinformation stories in larger proportions during the election period although the overall difference between the two groups was modest. This result seems to be mainly attributable to the disinformation narrative surrounding Conservative leader Erin O'Toole's position on the health care system, which originated on Chrystia Freeland's official Twitter account, as well as the involvement of liberal-leaning users in the disinformation discourse surrounding "Big Pharma". Moreover, this work showed that social bots were more predominantly involved in discussions around the COVID-19 pandemic. Their disinformation ratios also indicate that social bot accounts, leaning on both the liberal and the conservative side, propagated more disinformation content related to COVID-19 than to the 2021 Canadian election.

3.6. Conclusion

This chapter revealed the results of the different analyses performed in this research. On the one hand, it presented the results from the Botomoter algorithm and the partisan leaning classification model, which showed that close to 11% of accounts were social bots and that 54.84% of those leaned on the conservative side. In contrast, 45.16% leaned on the liberal side. In addition, analyses of the tactics of social bots across both groups found that there were no indicators of a coordinated astroturfing strategy or even of smoke screening or misdirecting. Overall, the low occurrence of social bots in the dataset and the diversity of topics discussed concerning the 2021 Canadian election explain the absence of a coordinated manipulation strategy by automated accounts. In other words, during the federal campaign, social bots spread tweets about Canadian politics. Still, these posts were not coordinated and did not amplify the visibility of political content in a meaningful way on Twitter. The measures of *Retweet Pervasiveness*, *Reply Rate*, and *Tweet Success Rate* were used to determine the level of interaction between social bots and humans based on their partisan leanings. These three metrics led to the conclusion that liberal-leaning social bots were the most effective at interacting with humans through retweets and replies during the campaign, even if the overall interactions between social bots and humans remained limited. Finally, the content-based approach used to detect disinformation related to the COVID-19 pandemic and the 2021 Canadian election demonstrated that disinformation tweets were not numerous during the campaign and were limited to a small number of users. By looking at the *Disinformation ratios* for each group, it was possible to show that liberal-leaning social bots shared more disinformation content than their conservative-leaning counterparts. The analysis also revealed that the disinformation relayed during the federal campaign was mainly focused on the COVID-19 issue, which holds for both social bots and humans and for liberal-leaning and conservative-leaning accounts.

Chapter 4. Discussion and Conclusion

In recent years, research on automated accounts and their implications for the democratic process has generated a great deal of interest among scholars in the social and computer sciences fields. The presence and engagement of social bots during election campaigns have been well documented by researchers (e.g., Bessi and Ferrara 2016; Howard and Kollanyi 2016; Varol et al. 2017; Luceri et al. 2019; Rheault and Musulan 2021). Prior to this study, scholars mainly focused on the U.S. They explored the number of social bots active on social media platforms, their partisan affiliation, the strategies they employ, and their involvement in disinformation campaigns. Nevertheless, it remains challenging to determine the influence of automated accounts on citizens' attitudes and voting behaviors as we still do not know if social bots' partisan leanings contribute to the volume of disinformation content spread online. This is especially the case for Canada, where very few studies have addressed the issue of social bots around the time of elections. Therefore, this master's thesis aimed to investigate three dimensions related to the partisan affiliation of social bot accounts: their presence, their strategies and involvement in discussions with humans, and the amount of disinformation content propagated during the 2021 Canadian federal election. To evaluate the impact of social bots' partisan leaning on their behaviors and strategies, this research relied on a large English dataset of close to 11.4 million tweets from more than 1.1 million distinct users involved in Twitter discussions around the 2021 Canadian election. Ultimately, the models made it possible to assign a bot score and a partisan leaning to 770,970 unique users who generated close to 7,800,000 tweets. We identified 330,606 liberal and 440,364 conservative leaners, of which 36,636 were liberal social bots, and 44,494 were conservative social bots. By using Twitter data focused on the 2021 Canadian election, this master's thesis adds to the literature on the role of partisan leanings on the behaviors and attitudes of automated accounts in election campaigns.

4.1. General Discussion

Drawing on the social science literature focused on social media, it was possible to derive three distinct hypotheses to explore in this research. The first hypothesis claimed that conservative-leaning social bots should be more numerous than their liberal-leaning counterparts, due to higher incentives for the conservative side to require the use of social bots, notably to create the impression of widespread popular support and to truly speak on behalf of ordinary citizens (Mudde 2019; Frost 2020; Silva and Proksch 2021). In coherence with previous studies (e.g., Bessi and Ferrara 2016; Chang et al. 2021; Rheault and Musulan 2021), this research demonstrated that conservative-leaning social bots were more numerous than liberal-leaning ones during the period of the 2021 Canadian election campaign. Of all the accounts identified as social bots, conservative-leaning accounts made up the largest share, with almost 55%. As such, the first hypothesis is confirmed. Although they generated fewer tweets than their liberal-leaning counterparts, conservative-leaning social bots were more prevalent on the Twitter platform during the 2021 Canadian election campaign.

The second hypothesis assumed that conservative-leaning social bots should be retweeted and replied to more by human accounts because of the structure of their online networks, which exhibits more signs of political homophily (e.g., Boutyline and Willer 2017; Hagen et al. 2022) and because of stronger patterns of motivated reasoning among conservative-leaning social media users (e.g., Swire, Ecker and Lewandowsky 2017; Grinberg et al. 2019). The results from the *Retweet Pervasiveness*, *Reply Rate*, and *Tweet Success Rate* measurements did not support the second hypothesis. Instead, the liberal-leaning social bots generated more interaction with their human counterparts throughout the 2021 Canadian election. Indeed, liberal-leaning social bots were more efficient at interacting with authentic accounts through retweets and replies. It is nevertheless important to note that humans' overall engagement with social bots was not overly pronounced during the 2021 campaign.

Lastly, the third hypothesis stated that social bots who leaned on the conservative side should propagate more disinformation content than those who leaned on the liberal side. This hypothesis stemmed from the fact that the most prevalent political disinformation stories tend to support conservative positions (e.g., Garrett and Bond 2021), as well as studies that showed that right-wing politicians tend to propagate more untrustworthy news sources on social media platforms, which can send powerful cues to party supporters (Lasser et al. 2022; Macdonald and Brown 2022). As opposed to what was expected, liberal-leaning social bots are the ones which shared the greatest amount of disinformation messages on Twitter. The disinformation narrative surrounding the privatization of the healthcare system by CPC leader Erin O'Toole and discussions around Big Pharma's involvement in the pandemic both seem to explain the volume of disinformation content associated with liberal-leaning social bots. Therefore, the third hypothesis of this master's thesis is not confirmed since the disinformation ratios of liberal-leaning social bots are higher than those of the conservative-leaning ones, both for the topic of COVID-19 and the 2021 Canadian election.

4.1.1. Limits and Future Works

The results found in this research have important implications in the first place, for our understanding of the behaviors and activities of social bots in the Canadian context and in the second place, for our comprehension of the role of their partisan leaning on the propagation of political disinformation. Although this research provided necessary evidence on the behaviors and strategies of social bots in the context of a Canadian federal election campaign, it nevertheless has important limitations. A significant limitation is the selection of the threshold to distinguish between real accounts and social bots. Although the threshold selection is based on previous studies (e.g., Woolley and Guilbeault 2018) and the distribution of bot scores in the sample, the choice of the threshold entails an element of arbitrariness. For instance, a threshold of 0.7 with the Botometer algorithm has been considered a conservative choice in the literature. In other words, by opting for a conservative threshold to discriminate between human and social bot accounts, this research could have underestimated the quantity of active social bots during the most recent federal campaign (Duan et al. 2022). For example, Appendix D shows that using thresholds of 0.3 and 0.5 can slightly alter each group's partisan division and increase the number of social bots. The threshold of 0.7 was nevertheless preferred in this research to reduce the amount of false positive social bot accounts.

Another aspect to consider is the approach based on keywords and hashtags employed in users' tweets to detect disinformation around the election and the COVID-19 pandemic. This approach could have slightly underestimated the amount of disinformation content propagated by users during the campaign since other techniques could have also been used to promote disinformation on social media, such as videos, images, memes, and URLs (Basch et al. 2021; Garimella and Eckles 2020). Nonetheless, the analysis of the disinformation content shared by social bots from each partisan leaning remains a good proxy for patterns in the dissemination of disinformation content by social bot accounts.

Despite the vast corpus of millions of tweets employed in this research, removing messages in French is an additional limitation. This choice is mainly explained by methodological considerations related to the implementation of the various detection models used throughout this work. Since French-speaking Canadians represent a significant portion of the electorate, especially in Québec, future works should incorporate French tweets and the partisanship of Bloc Québécois supporters. It would then be interesting to see if the dynamics introduced by these additions would modify the tendencies observed in a strictly English sample. In addition, a more fine-grained analysis capable of classifying users according to their party affiliation, not their partisan leaning, could be an interesting avenue to explore in future works focusing on the role of social bots in electoral campaigns in multi-party systems.

This research relied on a multi-step approach to better understand the role played by social bots' partisan leaning in their interactions with humans and in the amount of political disinformation they spread. While the results showed little differences between the two partisan-leaning groups in Canada, future works should continue to monitor and investigate the involvement of social bots in the spread of disinformation. This is particularly true since Elon Musk purchased the social network Twitter in October 2022. Elon Musk's restructuring of Twitter aims to increase freedom of speech and remove "censorship" on the platform (Zakrzewski, Siddiqui and Menn 2022; Conger and Hirsch 2022). For instance, Mr. Musk even offered "amnesty" to suspended Twitter accounts, such as Donald Trump, Steve Bannon, and David Duke, which leaves many experts concerned about a resurgence of hate speech and disinformation on the platform (Milmo 2022).

Moreover, liberal-leaning accounts might have been more engaged with political content and spread more disinformation on Twitter during the Canadian election campaign because of an underlying process of mass migration from mainstream social media platforms (e.g., Twitter, Facebook, Instagram) to new platforms (e.g., Parler, Gab, MeWe) (Otala et al. 2021). For instance, Parler and Gab respectively marketed themselves as the "free speech alternative" and as "championing free speech" in order to attract conservative users to their platforms (Šipka, Hannák and Urman 2022). Alternative social media platforms like Parler, Gab and MeWe are particularly appealing to conservative audiences because they perceive traditional platforms' efforts to label and remove misleading posts as anti-conservative bias and censorship (Greenhalgh, Krutka and Oltmann 2021). Therefore, the political left in Canada is not necessarily more gullible than the right regarding disinformation content. The overall higher engagement of liberal-leaning accounts and their greater spread of disinformation could be explained by the fact that the most active and extreme conservative-leaning users have migrated to less restrictive social media platforms or were suspended by Twitter (Greenhalgh, Krutka and Oltmann 2021). This migration process of conservative-leaning users towards alternative social media platforms could have left more room for liberal-leaning social bots and authentic users. It is also essential to mention that alternative platforms existed during the 2019 Canadian federal election, where Rheault and Musulan (2021) found an important cluster of active conservative users. However, it was estimated that following election fraud claims during the 2020 U.S. presidential campaign and the January 6th attack on the Capitol, Parler experienced a significant spike in the number of new users joining the platform (Aliapoulios et al. 2021). Brandom (2020) reported that Parler nearly had one million downloads on its mobile application after Election Day, on November 3rd, 2020. These events seem to have amplified the popularity of these platforms and accelerated the desertion of conservative-leaning users of Twitter in favor of less-moderated platforms, which could have had an impact on the number of actively engaged conservative-leaning users on Twitter during the 2021 Canadian election.

Lastly, the contributions of this work are threefold. First, by performing a large-scale analysis of millions of tweets, this research provides insights into the role of partisanship in the dissemination of social bots' disinformation on Twitter, which had never been examined before. It also adds specifically to the Canadian literature by presenting how different partisan groups discussed issues related to the 2021 federal election and the COVID-19 pandemic. Second, by putting forward a two-step approach to identify users' partisan leanings, this work also provides a contribution to computer science methodology. Labeling users based on their profile description and the content of their tweets is a promising methodology, which could potentially be a way to identify users not only by their partisan leaning but their party affiliation in multi-party systems in future works. Third, the findings presented regarding social bots' tactics and engagement with humans update the political science literature on modern political communication about the use of automated accounts as campaigning tools in Canada.

References

Abdine, Hadi, Yazhu Guo, Virgile Rennard, and Michalis Vazirgiannis. 2022. "Political Communities on Twitter: Case Study of the 2022 French Presidential Election." *arXiv preprint arXiv:2204.07436*: https://doi.org/10.48550/arXiv.2204.07436

Abilov, Anton, Yiqing Hua, Hana Matatov, Ofra Amir, and Mor Naaman. 2021. "*VoterFraud2020*: a Multi-modal Dataset of Election Fraud Claims on Twitter." *Proceedings of the Fifteenth International AAAI Conference on Web and Social Media*: 901-912.

Abokhodair, Norah, Daisy Yoo, and David W. McDonald. 2015. "Dissecting a Social Botnet: Growth, Content and Influence in Twitter." *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*: 839-851.

Abramowitz, Alan I. 2010. *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. New Haven: Yale University Press.

Aliapoulios, Max, Emmi Bevensee, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Savvas Zannettou. 2021. "An Early Look at the Parler Online Social Network." *Proceedings of the International AAAI Conference on Web and Social Media* 15(1): 943-951.

Allcott, Hunt et Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31(2): 211-236.

Allem, Jon-Patrick, and Emilio Ferrara. 2018. "Could Social Bots Pose a Threat to Public Health?" *American Journal of Public Health* 108(8): 1005-1006.

Al-Rawi, Ahmed, Oumar Kane, and Aimé-Jules Bizimana. 2021. "Topic Modelling of Public Twitter Discourses, Part Bot, Part Active Human User, on Climate Change and Global Warming." *Journal of Environmental Media* 2(1): 31-53.

Argitis, Theophilos, and Shelly Hagan. 2021, "Trudeau Powers Into Canada Vote With High Marks on Covid-19." *Bloomberg*. https://www.bloomberg.com/news/articles/2021-08-12/trudeau-expected-to-call-election-sunday-seeking-a-third-term#xj4y7vzkg (August 28, 2022).

Auter, Zachary J., and Jeffrey A. Fine. 2016. "Negative Campaigning in the Social Media Age: Attack Advertising on Facebook." *Political Behavior* 38(4): 999-1020.

Azzimonti, Marina, and Marcos Fernandes. 2022. "Social media networks, fake news, and polarization." *European Journal of Political Economy*: 1-25. https://doi.org/10.1016/j.ejpoleco.2022.10225

Badawy, Adam, Emilio Ferarra, and Kristina Lerman. 2018. "Analyzing the Digital Traces of Political Manipulation: The 2016 Russian Interference Twitter Campaign." *The 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*: 258-265.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23(1): 76-91.

Barnidge, Matthew, Homero Gil de Zúñiga, and Trevor Diehl. 2017. "Second Screening and Political Persuasion on Social Media." *Journal of Broadcasting & Electronic Media* 67(2): 309-331.

Basch, Corey H., Zoe Meleo-Erwin, Joseph Fera, Christie Jaime, and Charles E. Basch. 2021. "A global pandemic in the time of viral memes: COVID-19 vaccine misinformation and disinformation on TikTok." *Human Vaccines & Immunotherapeutics* 17(8): 2373-2377.

Bastos, Marco T., and Dan Mercea. 2019. "The Brexit Botnet and User-Generated Hyperpartisan News." *Social Science Computer Review* 37(1): 38-54.

Bellutta, Daniele, Catherine King, and Kathleen M. Carley. 2021. "Deceptive Accusations and Concealed Identities as Misinformation Campaign Strategies." *Computational and Mathematical Organization Theory* 27: 302-323.

Bensinger, Greg. 2020. "Take a Social Media Break Until You've Voted." *The New York Times*. https://www.nytimes.com/2020/10/06/opinion/facebook-twitter-election-misinformation.html (August 23, 2022).

Beskow M., David, and Kathleen M. Carley. 2018. "Introducing Bothunter: A Tiered Approach to Detection and Characterizing Automated Activity on Twitter." *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, Springer.

Beskow M., David, and Kathleen M. Carley. 2020. "Bot-Match: Social Bot Detection with Recursive Nearest Neighbors Search." *arXiv:2007.07636*. https://doi.org/10.48550/arXiv.2007.07636

Bessi, Alessandro, and Emilio Ferrara. 2016. "Social bots distort the 2016 US Presidential election online discussion." *First Monday* 21(11): 1-14.

Bloc Québécois. 2021. *Québécois – La plateforme politique, Bloc 2021*. https://www.blocquebecois.org/wp-content/uploads/2021/08/blocqc-Plateforme-2021.pdf (August 29, 2022).

Blouin, Louis. 2021. "Quels sont les risques de politiser la vaccination obligatoire?" *Ici Radio-Canada*. https://ici.radio-canada.ca/nouvelle/1817436/vaccination-election-federal-trudeau-pandemie (August 29, 2022).

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415): 295-298.

Boulianne, Shelley. 2015. "Social Media Use and Participation: A Meta-Analysis of Current Research." *Information, Communication & Society* 18(5): 524-538.

Boutyline, Andrei, and Robb Willer. 2017. "The Social Structure of Political Echo Chambers: Variation in Ideological Homophily in Online Networks." *Political Psychology* 38(3): 551-569

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. 2022. "Cross-country trends in affective polarization." *The Review of Economics and Statistics*: 1-60.

Brachten, Florian, Stefan Stieglitz, Lennart Hofeditz, Katharina Kloppenborg, and Annette Reimann. 2017. "Strategies and Influence of Social Bots in a 2017 German state election – A case study on Twitter." *Australasian Conference on Information Systems*, Hobart, Australia: 1-12.

Brandom, Russell. 2020. "Parler, a conservative Twitter clone, has seen nearly 1 million downloads since Election Day." *The Verge*. https://www.theverge.com/2020/11/9/21557219/parler-conservative-app-download-new-users-moderation-bias (November 30, 2022).

Bridgman, Aengus, and Mathieu Lavigne. 2022. "Reach, perceived accuracy, and engagement with real and fake stories during the 2021 Canadian federal election." The 2022 Canadian Political Science Association Annual Conference. Montréal, Canada.

Bridgman, Aengus, Mathieu Lavigne, Melissa Baker, Thomas Bergeron, Danielle Bohonos, Anthony Burton, Kathrine McCoy, Mackenzie Hart, Mathieu Lavault, Rupinder Liddar, Pangying Peng, Christopher Ross, Jaclyn Victor, Taylor Owen, and Peter Loewen. 2022. *Mis- and Disinformation During the 2021 Canadian Federal Election*. Media Ecosystem Observatory.

Broniatowski, David A., Amelia M. Jamison, SiHua Q, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C. Quinn, and Mark Dredze. 2018. "Weaponized Health Communication: Twitter Bots and Russian Trolls Amplify the Vaccine Debate." *American Journal of Public Health* 108(10): 1378-1384.

Burke, Ashley. 2021. "Twitter adds 'manipulated' warning label to tweet from Liberal candidate Chrystia Freeland." *CBC News*. https://www.cbc.ca/news/politics/twitter-labels-freeland-tweet-manipulated-media-1.6149734 (November 2, 2022).

Butt, Craig, and Thomas Hounslow. 2013. "Spambots target tweeting pollies." *The Sydney Morning Herald*. https://www.smh.com.au/politics/federal/spambots-target-tweeting-pollies-20130427-2ilwg.html (September 17, 2022).

Canada. House of Commons. 2018. *Democracy Under Threat: Risks and Solutions in the Era of Disinformation and Data Monopoly*. Ottawa: Report of the Standing Committee on Access to Information, Privacy and Ethics. https://www.ourcommons.ca/DocumentViewer/en/42-1/ETHI/report-17/page-ToC

Cantini, Riccardo, Fabrizio Marozzo, Domenico Talia, and Paolo Trunfio. 2022. "Analyzing Political Polarization on Social Media by Deleting Bot Spamming." *Big Data and Cognitive Computing* 6(1): 1-16.

Carvin, Stephanie. 2021. "The Big Lie: Is Canada's Election 44 at Risk from Foreign Interference?" *Centre for International Governance Innovation*. https://www.cigionline.org/articles/the-big-lie-is-canadas-election-44-at-risk-from-foreign-interference/ (September 30, 2022).

Castillo, Carlos, Marcelo Mendoza, and Barbara Poblete. 2013. "Predicting Information Credibility in Time-Sensitive Social Media." *Internet Research* 23(5): 560-588.

Chang, Ho-Chun Herbert, Emily Chen, Meiqing Zhang, Goran Muric, and Emilio Ferrara. 2021. "Social Bots and Social Media Manipulation in 2020 – The Year in Review." In *Handbook of Computational Social Science, Volume 1*, eds. Uwe Engel, Anabel Quan-Haase, Sunny Liu and Lars E. Lyberg, 304-323. London: Routledge.

Chang, Ho-Chun Herbert, and Emilio Ferarra. 2022. "Comparative Analysis of Social Bots and Humans During the COVID-19 Pandemic." *Journal of Computational Social Science*: 1-17.

Chavoshi, Nikan, Hossein Hamooni, and Abdullah Mueen. 2016a. "Identifying Correlated Bots in Twitter." *International Conference on Social Informatics*: 14-21.

Chavoshi, Nikan, Hossein Hamooni, and Abdullah Mueen. 2016b. "DeBot: Twitter Bot Detection via Warped Correlation." Paper presented at ICDM, Barcelona (Spain): 817-822.

Chen, Emily, Kristina Lerman, and Emilio Ferrara. 2020. "Tracking Social Media Discourse About the COVID-19 Pandemic: Development of a Public Coronavirus Twitter Data Set." *JMIR Public Health and Surveillance* 6(2): 1-12.

Chen, Chang-Feng, Wen Shi, Jing Yang, and Hao-Huan Fu. 2021a. "Social Bots' Role in Climate Change Discussion on Twitter: Measuring Standpoints, Topics, and Interaction Strategies." *Advances in Climate Change Research* 12(6): 913-923.

Chen, Wen, Diogo Pacheco, Kai-Cheng Yang, and Filippo Menczer. 2021b. "Neutral bots probe political bias on social media." *Nature Communications* 12(5580): 1-10.

Chen, Emily, Herbert Chang, Ashwin Rao, Kristina Lerman, Geoffrey Cowan, and Emilio Ferrara. 2021c. "COVID-19 Misinformation and the 2020 US Presidential Election." *The Harvard Kennedy School Misinformation Review* 1: 1-17.

Chen, Zhouhan, and Devika Subramanian. 2018. "An Unsupervised Approach to Detect Spam Campaigns that Use Botnets on Twitter." *arXiv preprint arXi*: 1804.05232.

Chotiner, Isaac. 2020. "Why Newsmax Supports Trump's False Voter-Fraud Claims." *The New Yorker*. https://www.newyorker.com/news/q-and-a/why-newsmax-supports-trumps-false-voter-fraud-claims (October 15, 2022).

Clarke, Harold D., Thomas J. Scotto, and Marianne C. Stewart. 2022. "Canada Votes in a COVID Election." In *The Canadian Federal Election Of 2021*, eds. Jon H. Pammett and Christopher Dornan, 295-321. Montréal: McGill-Queen's University Press.

Cochrane, Christopher. 2010. "Left-Right Ideology and Canadian Politics." *Canadian Journal of Political Science* 43(3): 583-605.

Cochrane, Christopher. 2015. *Left and Right: The Small World of Political Ideas*. Montréal: McGill-Queen's University Press.

Coldewey, Devin. 2012. "Romney Twitter account gets upsurge in fake followers, but from where?" *NBC News*. https://www.nbcnews.com/tech/tech-news/romney-twitter-account-gets-upsurge-fake-followers-where-flna928605 (September 7, 2022).

Colleoni, Elanor, Alessandro Rozza, and Adam Arvidsson. 2014. "Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data." *Journal of Communication* 64(2): 317-332.

Conger, Kate, and Lauren Hirsch. 2022. “Elon Musk Completes $44 Billion Deal to Own Twitter.” *The New York Times*. https://www.nytimes.com/2022/10/27/technology/elon-musk-twitter-deal-complete.html (November 26, 2022).

Conover, Michael D., Bruno Goncalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. “Predicting the Political Alignment of Twitter Users.” *IEEE International Conference on Privacy, Security, Risk, and Trust*: 192-199.

Conservative Party of Canada. 2021. *Canada’s Recovery Plan*. https://cpcassets.conservative.ca/wp-content/uploads/2021/08/25132033/5ea53c19b2e3597.pdf (August 29, 2022).

Cresci, Stefano, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2015. “Fame for sale: Efficient detection of fake Twitter followers.” *Decision Support Systems* 80: 56-71.

Cresci, Stefano, Marinella Petrocchi, Angelo Spognardi, and Stefano Tognazzi. 2019. “Better Safe Than Sorry: An Adversarial Approach to improve Social Bot Detection.” *Proceedings of the 10th ACM Conference on Web Science*: 47-56.

Cresci, Stefano. 2020. “A Decade of Social Bot Detection.” *Communications of the ACM* 63(10): 72-83.

Dang, Sheila, and Katie Paul. 2022. “Twitter says it removes over 1 million spam accounts each day.” *Reuters*. https://www.reuters.com/technology/twitter-says-it-removes-over-1-million-spam-accounts-each-day-2022-07-07/ (October 3, 2022).

Davis, Clayton, A., Onur Varol, Emilio Ferarra, Alessandro Flammini, and Filippo Menczer. 2016. “BotOrNot: A System to Evaluate Social Bots.” *Proceedings of the 25th International Conference Companion on World Wide Web*: 273-274.

Deb, Ashok, Luca Luceri, Adam Badawy, and Emilio Ferrara. 2019. “Perils and Challenges of Social Media and Election Manipulation Analysis: The 2018 US Midterms.” *Companion of the Web Conference 2019*: 1-10.

de Lima Salge, Carolina Alves, and Nicholas Berente. 2017. “Is That Social Bot Behaving Unethically?” *Communications of the ACM* 60(9): 29-31.

De Nicola, Rocco, Marinella Petrocchi, and Manuel Pratelli. 2021. “On the Efficacy of Old Features for the Detection of New Bots.” *Information Processing and Management* 58(6): 1-15.

Deryugina, O. V. 2010. “Chatterbots.” *Scientific and Technical Information Processing* 37(2): 143-147.

Des Mesnards, Nicolas Guenon, David Scott Hunter, Zakaria el Hjouji, and Tauhid Zaman. 2022. “Detecting Bots and Assessing Their Impact in Social Networks.” *Operations Research* 70(1): 1-22.

Dimitrova, Daniela V., and Jörg Matthes. 2018. “Social Media in Political Campaigning Around the World: Theoretical and Methodological Challenges.” *Journalism & Mass Communication Quarterly* 95(2): 333-342.

Djuric, Mickey. 2022. “Canadians are becoming more divided over COVID-19, politics: survey.” *The Canadian Press*. https://globalnews.ca/news/8733411/covid-politics-dividing-canadians-survey/ (August 28, 2022).

Downes, Steven. 2012. “Jasper admits to using Twitter bots to drive election bid.” *Inside Croydon*. https://insidecroydon.com/2012/11/26/jasper-admits-to-using-twitter-bots-to-drive-election-bid/ (September 7, 2022).

Druckman, James N., Erik Peterson, and Rune Slothuus. 2013. “How Elite Partisan Polarization Affects Public Opinion Formation.” American Political Science Review 107(1): 57-79.

Duan, Zening, Jianing Li, Josephine Lukito, Kai-Cheng Yang, Fan Chen, Dhavan V. Shah, and Sijia Yang. 2022. “Algorithmic Agents in the Hybrid Media System: Social Bots, Selective Amplification, and Partisan News about COVID-19.” *Human Communication Research* 48(3): 516-542.

Dubois, Elizabeth, and Fenwick McKelvey. 2019. “Political Bots: Disrupting Canada's Democracy.” *Canadian Journal of Communication* 44(2): 28-32.

Dubois, Elizabeth, and Taylor Owen. 2020. “Understanding the Digital Ecosystem: Findings from the 2019 Federal Election.” *The Digital Democracy Project*. https://www.mediatechdemocracy.com/media-ecosystem-observatory

Dufresne, Yannick, Sanjay Jeram, and Alexandre Pelletier. 2014. “The True North Strong and Free Healthcare? Nationalism and Attitudes Towards Private Healthcare Options in Canada.” *Canadian Journal of Political Science* 47(3): 569-595.

Dunham, Ken, and Jim Melnick. 2009. *Malicious Bots: An Inside Look into the Cyber-Criminal Underground of the Internet*. New York: CRC Press.

Enli, Gunn. 2017. “Twitter as Arena for the Authentic Outsider: Exploring the Social Media Campaigns of Trump and Clinton in the 2016 US Presidential Election.” *European Journal of Communication* 32(1): 50-61.

European Commission. 2018. *A Multi-Dimensional Approach to Disinformation: Report of the Independent High Level Group on Fake News and Online Disinformation*. Luxembourg: Publications Office of the European Union. https://digital-strategy.ec.europa.eu/en/library/final-report-high-level-expert-group-fake-news-and-online-disinformation (September 29, 2022).

European Parliament. 2021. *Key Social Media Risks to Democracy: Risks from Surveillance, Personalisation, Disinformation, Moderation and Microtargeting*. Brussels: European Parliamentary Research Service. https://www.europarl.europa.eu/RegData/etudes/IDAN/2021/698845/EPRS_IDA(2021)698845_EN.pdf

Farneti, Roberto. 2012. “Cleavage Lines in Global Politics: Left and Right, East and West, Earth and Heaven.” *Journal of Political Ideologies* 17(2): 127-145.

Feldman, Max. 2020. “Dirty Tricks: 9 Falsehoods that Could Undermine the 2020 Election.” *The Brennan Center*. https://www.brennancenter.org/our-work/research-reports/dirty-tricks-9-falsehoods-could-undermine-2020-election (October 17, 2022).

Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. “The Rise of Social Bots.” *Communications of the ACM* 59(7) 96-104.

Ferrara, Emilio. 2017. “Disinformation and Social Bot Operations in the Run Up to the French Presidential Election.” *arXiv preprint arXiv:1707.00086*: 1-33.

Ferrara, Emilio. 2019. “The History of Digital Spam.” *Communications of the ACM* 62(8): 82-91.

Ferrara, Emilio. 2020a. "What Types of COVID-19 Conspiracies Are Populated by Twitter Bots?" *First Monday* 25(6).

Ferrara, Emilio. 2020b. "Bots, Elections, and Social Media: A Brief Overview." In *Disinformation, Misinformation, and Fake News in Social Media: Emerging Research Challenges and Opportunities*, eds. Kai Shu, Suhang Wang, Dongwon Lee, and Huan Liu, 95-114. Switzerland: Springer.

Ferrara, Emilio, Herbert Chang, Emily Chen, Goran Muric, and Jaime Patel. 2020. "Characterizing social media manipulation in the 2020 U.S. presidential election." *First Monday* 25(11). https://doi.org/10.5210/fm.v25i11.11431

Forelle, Michelle, Phil Howard, Andrés Monroy-Hernández, and Saiph Savage. 2015. "Political Bots and the Manipulation of Public Opinion in Venezuela." *arXiv preprint arXiv:1507.07109*. https://doi.org/10.48550/arXiv.1507.07109

Freelon, Deen. 2019. "Tweeting Left, Right & Center: How users and attention are distributed across Twitter." *Reports of The Knight Foundation*. https://knightfoundation.org/reports/tweeting-left-right-center-how-users-and-attention-are-distributed-across-twitter/ (December 5, 2022).

Fridman, Ariel, Rachel Gershon, and Ayelet Gneezy. 2021. "COVID-19 and vaccine hesitancy: A longitudinal study." *PLoS ONE* 16(4): 1-12.

Frost, Catherine. 2020. "The Power of Voice: Bots, Democracy and the Problem of Political Ventriloquism." *Journal of Political Power* 13(1): 6-21.

Galgoczy, Michael C., Atharva Phatak, Danielle Vinson, Vijay K. Mago, and Philippe J. Giabbanelli. 2022. "(Re)shaping online narratives: when bots promote the message of President Trump during his first impeachment." *PeerJ Computer Science* 8: 1-26.

Gallwitz, Florian, and Michael Kreil. 2022. "Investigating the Validity of Botometer-based Social Bot Studies." *arXiv preprint arXiv*. https://doi.org/10.48550/arXiv.2207.11474

Garimella, Kiran, and Dean Eckles. 2020. "Images and Misinformation in Political Groups: Evidence from WhatsApp in India." *The Harvard Kennedy School Misinformation Review* 1(5): 1-12.

Garrett, Kelly R. 2019. "Social media's contribution to political misperceptions in U.S. Presidential elections." *PLoS ONE* 14(3).

Garrett, Kelly R., and Robert M. Bond. 2021. "Conservatives' susceptibility to political misperceptions." *Science Advances* 7(23): 1-9.

Gebelhoff, Robert. 2019. "Here's how social media could threaten democracy — even without the help of Russians." *The Washington Post*. https://www.washingtonpost.com/opinions/2019/09/11/heres-how-social-media-could-threaten-democracy-even-without-help-russians/ (August 17, 2022).

Geer, David. 2005. "Malicious Bots Threaten Network Security." *Computer* 38(1): 18-20.

Geys, Benny, Tom-Reiel Heggedal, and Rune J.Sørensen. 2022. "Age and vote choice: Is there a conservative shift among older voters?" *Electoral Studies* 78(102485): 1-8.

Giasson, Thierry, Jennifer Lees-Marshment, and Alexander Marland. 2012. “Introducing Political Marketing.” In *Political Marketing in Canada*, eds. Alexander Marland, Thierry Giasson and Jennifer Lees-Marshment, 3-21. Vancouver: University of British Colombia Press.

Gilens, Martin, and Naomi Murakawa. 2002. “Elite Cues and Political Decision-Making.” *Research in Micropolitics* 6(1): 15-49.

Goga, Oana, Giridhari Venkatadri, and Krishna P. Gummadi. 2015. “The Doppelgänger Bot Attack: Exploring Identity Impersonation in Online Social Networks.” *Proceedings of the 2015 Internet Measurement Conference*: 141-153.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2021. “Social Media, Sentiment and Public Opinions: Evidence from #Brexit and #USElection.” *European Economic Review* 136(103772): 1-22.

Gorwa, Robert, and Douglas Guilbeault. 2020. “Unpacking the Social Media Bot: A Typology to Guide Research and Policy.” *Policy and Internet* 12(2): 225-248.

Green, Jon, Jared Edgerton, Daniel Naftel, Kelsey Shoub, and Skyler J. Cranmer. 2020. “Elusive consensus: Polarization in elite communication on the COVID-19 pandemic.” *Science Advances* 6(28): 1-5.

Green Party of Canada. 2021. *Platform 2021 - Be Daring* https://www.greenparty.ca/sites/default/files/gpc_platform_en_v-02.pdf (August 29, 2022).

Greenhalgh, Spencer, Daniel G. Krutka, and Shannon M. Oltmann. 2021. “Gab, Parler, and (Mis)educational Technologies: Reconsidering Informal Learning on Social Media Platforms.” *The Journal of Applied Instructional Design* 10(3): 1-14.

Grimes, David Robert. 2021. “Medical Disinformation and the Unviable Nature of COVID-19 Conspiracy Theories.” *PLoS ONE* 16(3).

Gruzd, Anatoliy, and Jeffrey Roy. 2014. “Investigating Political Polarization on Twitter: A Canadian Perspective.” *Policy & Internet* 6(1): 28-45.

Gu, Yupeng, Ting Chen, Yizhou Sun, Bingyu Wang. 2016. “Ideology Detection for Twitter Users with Heterogeneous Types of Links.” *arXiv preprint arXiv:1612.08207*: 1-9. https://doi.org/10.48550/arXiv.1612.08207

Guess, Andrew, Jonathan Nagler, and Joshua Tucker. 2019. “Less than you think: Prevalence and predictors of fake news dissemination on Facebook.” *Science Advances* 5(1): 1-8.

Guilbeault, Douglas, and Samuel C. Woolley. 2016. “How Twitter Bots Are Shaping the Election.” *The Atlantic*. https://www.theatlantic.com/technology/archive/2016/11/election-bots/506072/ (September 16, 2022).

Gunther, Richard, Paul A. Beck, and Erik C. Nisbet. 2019. “"Fake News" and the Defection of 2012 Obama Voters in the 2016 Presidential Election.” *Electoral Studies* 61(102030): 1-17.

Guo, Bin, Yasan Ding, Lina Yao, Yunji Liang, and Zhiwen Yu. 2021. “The Future of False Information Detection on Social Media: New Perspectives and Trends.” *ACM Computing Surveys* 53(4): 1-36.

Hagen, Loni, Stephen Neely, Thomas E. Keller, Ryan Scharf, and Fatima Espinoza Vasquez. 2022. "Rise of the Machines? Examining the Influence of Social Bots on a Political Discussion Network." *Social Science Computer Review* 40(2): 264-287.

Hall, Wendy, Ramine Tinati, and Will Jennings. 2018. "From Brexit to Trump: Social Media's Role in Democracy." *Computer* 51(1): 18-27.

Hare, Christopher, and Keith T. Poole. 2014. "The Polarization of Contemporary American Politics." *Polity* 46(3): 411-429.

Hasert, Tommy, and Gabriele Hooffacker. 2019. "Overrated bots?" *Journalism Research* 2(2): 135-147.

Hayati, Pedram, Kevin Chai, Vidyasagar Potdar, and Alex Talevski. 2009. "HoneySpam 2.0: Profiling Web Spambot Behaviour." In *Principles and Practice of Multi-Agent Systems*, eds. Jung-Jin Yang, Makoto Yokoo, Takayuki Ito, Zhi Kin and Paul Scerri, 335-344. Berlin: Springer.

Himelein-Wachowiak, McKenzie, Salvatore Giorgi, Amanda Devoto, Muhammad Rahman, Lyle Ungar, Andrew H. Schwartz, David H, Epstein, Lorenzo Leggio, and Brenda Curtis. 2021. "Bots and Misinformation Spread on Social Media: Implications for COVID-19." *Journal of Medical Internet Research* 23(5).

Horne, Benjamin D., and Sibel Adali. 2017. "This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News." *Eleventh International AAAI Conference on Web and Social Media*: 759-766.

Howard, Philip N., and Bence Kollanyi. 2016. "Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum." *arXiv:1606.06356*: 1-6.

Howard, Philip N., Samuel Woolley, and Ryan Calo. 2018. "Algorithms, bots, and political communication in the US 2016 election: The challenge of automated political communication for election law and administration." *Journal of Information Technology & Politics* 15(2): 81-93.

Hugging Face. 2022. "RoBERTa large model – Model Card." *Hugging Face.co*. https://huggingface.co/roberta-large?text=Paris+is+the+%3Cmask%3E+of+France (October 27, 2022).

Imperva. 2022. *2022 Imperva Bad Bot Report. Evasive Bots Drive Online Fraud*. https://www.imperva.com/resources/resource-library/reports/bad-bot-report/?utm_medium=PR&utm_source=Organic-Referral (September 2, 2022).

Ito, Jun, Hiroyuki Toda, Yoshimasa Koike, Jing Song, and Satoshi Oyama. 2015. "Assessment of Tweet Credibility with LDA Features." *Proceedings of the 24th International Conference on World Wide Web*: 953-958.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. "Affect, Not Ideologya Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76(3): 405-431.

Jamieson, Kathleen Hall. 2018. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President: What We Don't, Can't, and Do Know*. Oxford: Oxford University Press.

Jeangène Vilmer, Jean-Baptiste. 2019. "The "Macron Leaks" Operation: A Post-Mortem." *Atlantic Council*. https://www.atlanticcouncil.org/in-depth-research-reports/report/the-macron-leaks-operation-a-post-mortem/ (September 5, 2022).

Jones, Jason J., Robert M. Bond, Eytan Bakshy, Dean Eckles, and James H. Fowler. 2017. "Social Influence and Political Mobilization: Further Evidence from a Randomized Experiment in the 2012 U.S. Presidential Election." *PLoS ONE* 12(4).

Jones, Keenan, Jason R. C. Nurse, and Shujun Li. 2022. "Out of the Shadows: Analyzing Anonymous' Twitter Resurgence during the 2020 Black Lives Matter Protests." *Proceedings of the Sixteenth International AAAI Conference on Web and Social Media* 16(1): 417-428.

Karadeglija, Anja. 2021. "From COVID conspiracies to rigged voting machines, misinformation plagues federal election." *National Post*. https://nationalpost.com/news/politics/election-2021/from-covid-conspiracies-to-rigged-voting-machines-misinformation-plagues-the-election (October 14, 2022).

Kariyappa, Sanjay, and Moinuddin K. Qureshi. 2020. "Defending Against Model Stealing Attacks with Adaptive Misinformation." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*: 770-778.

Kassam, Ashifa. 2018. "Canada fails to grasp scale of social media 'bot' use in politics, academics warn." *The Guardian*. https://www.theguardian.com/world/2018/apr/12/canada-bots-twitter-politics-social-media (September 30, 2022).

Kearney, Michael W. 2018. "Tweetbotornot: Detecting Twitter Bots." https://tweetbotornot.mikewk.com/ (August 16, 2022).

Kearney, Michael W. 2021. "Conservatives and Twitter Bots." In *Conservative Political Communication: How Right-Wing Media and Messaging (Re)Made American Politics*, eds. Sharon E. Jarvis, 155-165. New York: Routledge.

Keller, Tobias R., and Ulrike Klinger. 2019. "Social Bots in Election Campaigns: Theoretical, Empirical, and Methodological Implications." *Political Communication* 36(1): 171-189.

Keller, Franziska B., David Schoch, Sebastian Stier, and JungHwan Yang. 2020. "Political Astroturfing on Twitter: How to Coordinate a Disinformation Campaign." *Political Communication* 37(2): 256-280.

Kenny, Ryan, Baruch Fischhoff, Alex Davis, Kathleen M. Carley, and Casey Canfield. 2022. "Duped by Bots: Why Some are Better than Others at Detecting Fake Social Media Personas." *Human Factors: SAGE Journals*: 1-15. doi: 10.1177/00187208211072642.

Kevins, Anthony, and Stuart N. Soroka. 2018. "Growing Apart? Partisan Sorting in Canada, 1992–2015." *Canadian Journal of Political Science* 51(1): 103-133.

Kim, Young Mie, Jordan Hsu, David Neiman, Colin Kou, Levi Bankston, Soo Yun Kim, Richard Heinrich, Robyn Baragwanath, and Garvesh Raskutti. 2018. "The Stealth Media? Groups and Targets behind Divisive Issue Campaigns on Facebook." *Political Communication* 35(4): 515-541.

King, Gary, Jennifer Pan, and Margaret E. Roberts. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review* 111(3): 484-501.

Klepper, David. 2022. "Twitter ends enforcement of COVID misinformation policy." *Ap News – The Associated Press*. https://apnews.com/article/twitter-ends-covid-misinformation-policy-cc232c9ce0f193c505bbc63bf57ecad6 (November 30, 2022).

Kollanyi, Bence. 2016. "Where Do Bots Come From? An Analysis of Bot Codes Shared on GitHub." *International Journal of Communication* 10: 4932-4951.

Kollanyi, Bence, Philip N. Howard, and Samuel C. Woolley. 2016. "Bots and Automation over Twitter during the Third U.S. Presidential Debate." *Data Memo Oxford Internet Institute, UK: Project on Computational Propaganda*. https://demtech.oii.ox.ac.uk/research/posts/bots-and-automation-over-twitter-during-the-third-u-s-presidential-debate/

Kulkarni, Ajay, Deri Chong, and Feras A. Batarseh. 2020. "Foundations of data imbalance and solutions for a data democracy." In *Data Democracy: At the Nexus of Artificial Intelligence, Software Development, and Knowledge Engineering*, eds. Feras A. Batarseh and Ruixin Yang, 83-106. London: Elsevier.

Kushin, Matthew James, and Masahiro Yamamoto. 2011. "Did Social Media Really Matter? College Students' Use of Online Media and Political Decision Making in the 2008 Election." In *New Media, Campaigning and the 2008 Facebook Election*, eds. Thomas J. Johnson and David D. Perlmutter, 55-78. Oxfordshire: Routledge.

Kuta, Sarah. 2020. "There Are Two Americas, and Age Is the Divider." *Chicago Book Review*. https://www.chicagobooth.edu/review/there-are-two-americas-and-age-divider (September 20, 2022).

Kwon, Sejeong, Meeyoung Cha, Kyomin Jung, Wei Chen, and Yajun Wang. 2013. "Prominent Features of Rumor Propagation in Online Social Media." *IEEE 13th International Conference on Data Mining*: 1103-1108.

Larson, Jennifer M., Jonathan Nagler, Jonathan Ronen, and Joshua A. Tucker. 2019. "Social Networks and Protest Participation: Evidence from 130 Million Twitter Users." *American Journal of Political Science* 63(3): 690-705.

Lasser, Jana, Segun Taofeek Aroyehun, Almog Simchon, Fabio Carrella, David Garcia, and Stephen Lewandowsky. 2022. "Social Media Sharing by Political Elites: An Asymmetric American Exceptionalism." *Computers and Society*: 1-6.

Lavoie, Kim, Vincent Gosselin-Boucher, Jovana Stojanovic, Samir Gupta, Myriam Gagné, Keven Joyal-Desmarais, Katherine Séguin, Sherri Sheinfeld Gorin, Paula Ribeiro, Brigitte Voisard, Michael Vallis, Kimberly Corace, Justin Presseau, and Simon Bacon. 2021. "Understanding national trends in COVID-19 vaccine hesitancy in Canada: results from five sequential cross-sectional representative surveys spanning April 2020–March 2021." *British Medical Journal* 12(4): 1-12.

Liberal Party of Canada. 2021. *Forward For Everyone*. https://liberal.ca/wp-content/uploads/sites/292/2021/09/Platform-Forward-For-Everyone.pdf (August 29, 2022).

Linvill, Darren L., Brandon C. Boatwright, Will J. Grant, and Patrick L. Warren. 2019. "'THE RUSSIANS ARE HACKING MY BRAIN!' Investigating Russia's Internet Research Agency Twitter Tactics During the 2016 United States Presidential Campaign." *Computers in Human Behavior* 99: 292-300.

Lipton, Zachary Chase, Charles Elkan, and Balakrishnan Narayanaswamy. 2014. "Thresholding Classifiers to Maximize F1 Score." *arXiv:1402.1892*: 1-16. https://doi.org/10.48550/arXiv.1402.1892

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. "RoBERTa: A Robustly Optimized BERT Pretraining Approach." *arXiv:1907.11692*: https://doi.org/10.48550/arXiv.1907.11692

Luceri, Luca, Ashok Deb, Adam Badawy, and Emilio Ferrara. 2019. "Red Bots Do It Better: Comparative Analysis of Social Bot Partisan Behavior." *Companion Proceedings of the 2019 World Wide Web Conference*: 1007-1012.

Lupton, Robert N., Steven M. Smallpage, and Adam M. Enders. 2020. "Values and Political Predispositions in the Age of Polarization: Examining the Relationship between Partisanship and Ideology in the United States, 1988–2012." *British Journal of Political Science* 50(1): 241-260.

Lurie, Shira. 2021. "After an ugly campaign, Canada must heed the lessons of January 6." *Rabble*. https://rabble.ca/politics/canadian-politics/after-ugly-campaign-canada-must-heed-lessons-january-6/ (October 14, 2022).

Macdonald, Maggie, and Megan A. Brown. 2022. "Republicans are increasingly sharing misinformation, research finds." *The Washington Post*. https://www.washingtonpost.com/politics/2022/08/29/republicans-democrats-misinformation-falsehoods/ (September 12, 2022).

Mahl, Daniela, Jing Zeng, and Mike S. Schäfer. 2021. "From "Nasa Lies" to "Reptilian Eyes": Mapping Communication About 10 Conspiracy Theories, Their Communities, and Main Propagators on Twitter." *Social Media + Society* 7(2): 1-12.

Maleki, Maryam, Mohammad Arani, Erik Buchholz, Esther Mead, and Nitin Agarwal. 2021. "Applying an Epidemiological Model to Evaluate the Propagation of Misinformation and Legitimate COVID-19-Related Information on Twitter." *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, Springer: 24-34.

Marcellino, William, Madeline Magnuson, Anne Stickells, Benjamin Boudreaux, Todd C. Helmus, Edward Geist, and Zev Winkelman. 2020. *Counter-Radicalization Bot Research: Using Social Bots to Fight Violent Extremism*. RAND Corporation, California.

Marietta, Morgan. 2012. *A Citizen's Guide to American Ideology: Conservatism and Liberalism in Contemporary Politics*. New York: Routledge.

Marlow, Thomas, Sean Miller, and J. Timmons Roberts. 2021. "Bots and Online Climate Discourses: Twitter Discourse on President Trump's Announcement of U.S. Withdrawal from the Paris Agreement." *Climate Policy* 21(6): 765-777.

Martini, Franziska, Paul Samula, Tobias R. Keller, and Ulrike Klinger. 2021. "Bot, or not? Comparing three methods for detecting social bots in five political discourses." *Big Data & Society* 8(2): 1-13.

Mazza, Michele, Stefano Cresci, Marco Avvenuti, Walter Quattrociocchi, and Maurizio Tesconi. 2019. "RTbust: Exploiting Temporal Patterns for Botnet Detection on Twitter." *Proceedings of the 10th ACM Conference on Web Science*: 183-192.

McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge: The MIT Press.

McCright, Aaron M., and Riley E. Dunlap. 2017. "Combatting Misinformation Requires Recognizing Its Types and the Factors that Facilitate Its Spread and Resonance." *Journal of Applied Research in Memory and Cognition* 6(4): 389-396.

McKay, Spencer, and Chris Tenove. 2021. "Disinformation as a Threat to Deliberative Democracy." *Political Research Quarterly* 74(3): 703-717.

Medeiros, Mike. 2021. "Demand without Supply: Populist Attitudes without Salient Supply-Side Factors of Populism." *Canadian Journal of Political Science* 54(4): 918-938.

Mellon, Jonathan, and Christopher Prosser. 2017. "Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users." *Research & Politics* 4(3): 1-9.

Merkley, Eric, Aengus Bridgman, Peter John Loewen, Taylor Owen, Derek Ruths, and Oleg Zhilin. 2020. "A Rare Moment of Cross-Partisan Consensus: Elite and Public Response to the COVID-19 Pandemic in Canada." *Canadian Journal of Political Science* 53(2): 311-318.

Merkley, Eric. 2020. "Polarization Eh? Ideological Divergence and Partisan Sorting in the Canadian Mass Public." *OSF Preprints*: 1-30. doi:10.31219/osf.io/cnzer.

Merkley, Eric. 2021. "Ideological and Partisan Bias in the Canadian Public." *Canadian Journal of Political Science* 54(2): 267-291.

Messing, Solomon, and Sean J. Westwood. 2012. "Selective Exposure in the Age of Social Media: Endorsements Trump Partisan Source Affiliation When Selecting News Online." *Communication Research* 41(8): 1042-1063.

Metaxas, Panagiotis Takis, and Eni Mustafaraj. 2010. "From Obscurity to Prominence in Minutes: Political Speech and Real-Time Search." *Web Science Conference*, Raleigh (United States): 1-7.

Metaxas, Panagiotis Takis, and Eni Mustafaraj. 2012. "Social Media and the Elections." *Science* 338(6106): 472-473.

Miller, Blake. 2016. "Automated Detection of Chinese Government Astroturfers Using Network and Social Metadata." *Social Science Research Network* 2738325: 1-35. http://dx.doi.org/10.2139/ssrn.2738325

Miller, Joanne M, Kyle L. Saunders, and Christina E. Farhart. 2016. "Conspiracy Endorsement as Motivated Reasoning: The Moderating Roles of Political Knowledge and Trust." *American Journal of Political Science* 60(4): 824-844.

Milmo, Dan. 2022. "Elon Musk offers general amnesty to suspended Twitter accounts." *The Guardian*. https://www.theguardian.com/technology/2022/nov/24/elon-musk-offers-general-amnesty-to-suspended-twitter-accounts (November 26, 2022).

Moffitt, J.D., Catherine King, and Kathleen M. Carley. 2021. "Hunting Conspiracy Theories During the COVID-19 Pandemic." *Social Media + Society* 7(3): 1-17.

Mudde, Cas. 2019. *The Far Right Today*. Medford: Polity Press.

Müller, Jan-Werner. 2016. *What Is Populism?* Philadelphia: University of Pennsylvania Press.

Murthy, Dhiraj, Alison B. Powell, Ramine Tinatim Nick Anstead, Leslie Carr, Susan J. Halford, and Mark Weal. 2016. "Bots and Political Influence: A Sociotechnical Investigation of Social Network Capital." *International Journal of Communication* 10: 4952-4971.

Neudert, Lisa-Maria, Bence Kollanyi, and Philip N. Howard. 2017. *Junk News and Bots during the German Federal Presidency Election: What Were German Voters Sharing Over Twitter?* Project on Computational Propaganda: Oxford, UK. https://comprop.oii.ox.ac.uk/

New Democratic Party of Canada. 2021. *Ready for Better: New Democrats' Commitments to You*. https://xfer.ndp.ca/2021/Commitments/Ready%20for%20Better%20-%20NDP%202021%20commitments.pdf (August 29, 2022).

Ng, Lynnette Hui Xian, Dawn C. Robertson, and Kathleen M. Carley. 2022. "Stabilizing a supervised bot detection algorithm: How much data is needed for consistent predictions?" *Online Social Networks and Media* 28(100198): 1-8.

Nuttall, Jeremy. 2021. "Canada must safeguard against foreign interference in future elections: expert." *The Toronto Star*. https://www.thestar.com/politics/federal-election/2021/09/17/canada-must-safeguard-against-foreign-interference-in-future-elections-expert.html (September 30, 2022).

Orabi, Mariam, Djedjiga Mouheb, Zaher Al Aghbari, and Ibrahim Kamel. 2020. "Detection of Bots in Social Media: A Systematic Review." *Information Processing and Management* 57(4): 1-23.

Otala, Jacqueline M., Gillian Kurtic, Isabella Grasso, Yu Liu, Jeanna Matthews, and Golshan Madraki. 2021. "Political Polarization and Platform Migration: A Study of Parler and Twitter Usage by United States of America Congress Members." *Companion Proceedings of the Web Conference 2021*. Ljubljana, Slovenia: 224-231.

Owen, Taylor, Peter Loewen, Derek Ruths, Aengus Bridgman, Robert Gorwa, Stephanie MacLellan, Eric Merkley, and Oleg Zhilin. 2020. "Lessons in Resilience: Canada's Digital Media Ecosystem and the 2019 Election." *Public Policy Forum*.

Ozer, Mert, Mehmet Yigit Yildirim, and Hasan Davulcu. 2019. "Measuring the Polarization Effects of Bot Accounts in the U.S. Gun Control Debate on Social Media." *Proceedings of ACM Conference (Conference'17)*: 1-10.

Paradise, Abigail, Rami Puzis, and Asaf Shabtai. 2014. "Anti-Reconnaissance Tools: Detecting Targeted Socialbots." *IEEE Internet Computing* 18(5): 11-19.

Parliament of Canada. House of Commons. *Debates of the House of Commons*, Vol. 150, 2nd Session, 43rd Parliament, April 21st, 2021, p. 5918.

Pastor-Galindo, Javier, Mattia Zago, Pantaleone Nespoli, Sergio López Bernal, Alberto Huertas Celdrán, Manuel Gil Pérez, José A. Ruipérez-Valiente, Gregorio Martínez Pérez, and Félix Gómez Mármol. 2020. "Spotting Political Social Bots in Twitter: A Use Case of the 2019 Spanish General Election." *IEEE Transactions on Network and Service Management* 17(4): 2156-2170.

Pelrine, Kellin, Anne Imouza, Gabrielle Desrosiers-Brisebois, Sacha Lévy, Jacob-Junqi Tian, Zachary Yang, Aarash Feizi, André Blais, Jean-François Godbout, and Reihaneh Rabbany. 2022. "Party Prediction for Twitter." *Conference of the American Political Science Association Meeting and Exhibition*, Montréal, Canada.

Pennacchiotti, Marco, and Ana-Maria Popescu. 2011. "A Machine Learning Approach to Twitter User Classification." *Proceedings of the International AAAI Conference on Web and Social Media* 5(1): 281-288.

Pennycook, Gordon, and David G. Rand. 2019. "Fighting Misinformation on Social Media Using Crowdsourced Judgments of News Source Quality." *Proceedings of the National Academy of Sciences* 16(7): 2521-2526.

Pennycook, Gordon, and David G. Rand. 2021. "Examining False Beliefs About Voter Fraud in the Wake of the 2020 Presidential Election." *The Harvard Kennedy School Misinformation Review* 2(1): 1-19.

Pennycook, Gordon, Jonathon McPhetres, Bence Bago, and David G. Rand. 2021. "Beliefs About COVID-19 in Canada, the United Kingdom, and the United States: A Novel Test of Political Polarization and Motivated Reasoning." *Personality and Social Psychology Bulletin* 48(5): 750-765.

People's Party of Canada. 2021. *Our Platform*. https://www.peoplespartyofcanada.ca/platform (August 29, 2022).

Pescetelli, Niccolo, Daniel Barkoczi, and Manuel Cebrian. 2022. "Bots Influence Opinion Dynamics Without Direct Human-Bot Interaction: The Mediating Role of Recommender Systems." *Applied Network Science* 7(46): 1-19.

Petrocik, John Richard. 2009. "Measuring party support: Leaners are not independents." *Electoral Studies* 28(4): 562-572.

Pickard, Victor. 2017. "Media Failures in the Age of Trump." *The Political Economy of Communication* 4(2): 118-122.

Plattner, Marc F. 2010. "Democracy's Past and Future: Populism, Pluralism, and Liberal Democracy." *Journal of Democracy* 21(1): 81-92.

Pozzana, Iacopo, and Emilio Ferrara. 2020. "Measuring Bot and Human Behavioral Dynamics." *Frontiers in Physics*. https://doi.org/10.3389/fphy.2020.00125 (September 11, 2022).

Preoţiuc-Pietro, Daniel, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. "Beyond Binary Labels: Political Ideology Prediction of Twitter Users." *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers): 729-740.

Ratkiewicz, Jacob, Michael Conover, Mark Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011. "Detecting and Tracking Political Abuse in Social Media." *Proceedings of the International AAAI Conference on Web and Social Media* 5(1): 297-304.

Rectenwald, Michael D. 2022. "What Is the Great Reset?" *The Independent Institute*. https://www.independent.org/publications/article.asp?id=14039 (October 13, 2022).

Rheault, Ludovic, and Christopher Cochrane. 2020. "Word Embeddings for the Analysis of Ideological Placement in Parliamentary Corpora." *Political Analysis* 28(1): 112-133.

Rheault, Ludovic, and Andreea Musulan. 2021. "Efficient detection of online communities and social bot activity during electoral campaigns." *Journal of Information Technology & Politics* 18(3): 324-337.

Rodríguez-Ruiz, Jorge, Javier Israel Mata-Sánchez, Raúl Monroy, Octavio Loyola-González, and Armando López-Cuevas. 2020. "A one-class classification approach for bot detection on Twitter." *Computers & Security* 101715: 1:14.

Romm, Tony, and Elizabeth Dwoskin. 2021. "Twitter purged more than 70,000 accounts affiliated with QAnon following Capitol riot." *The Washington Post*. https://www.washingtonpost.com/technology/2021/01/11/trump-twitter-ban/ (October 3, 2022).

Roozenbeek, Jon, Claudia R. Schneider, Sarah Dryhurst, John Kerr, Alexandra L. J. Freeman, Gabriel Recchia, Anne Marthe van der Bles, and Sander van der Linden. 2020. "Susceptibility to Misinformation about COVID-19 Around the World." *The Royal Society Publishing* 7(10): 1-15.

Sardarizadeh, Shayan. 2021. "Twitter suspends 70,000 accounts linked to QAnon." *BBC News*. https://www.bbc.com/news/technology-55638558 (October 3, 2022).

Sayyadiharikandeh, Mohsen, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2020. "Detection of Novel Social Bots by Ensembles of Specialized Classifiers." *Proceedings of the 29th ACM International Conference on Information & Knowledge*: 2725-2532.

Schindler, Paulina, Marek Opuszko, and Meena Stöbesand. 2021. "Dimensions of Fake News." *DIGITAL 2021, Advances on Societal Digital Transformation*: Athens, Greece.

Schuchard, Ross, Andrew Crooks, Anthony Stefanidis, and Arie Croitoru. 2019. "Bots Fired: Examining Social Bot Evidence in Online Mass Shooting Conversations." *Palgrave Communications* 5(1): 1-12.

Shah, Sono, Emma Remy, Aaron Smith, Dennis Quinn, and Regina Widjaya. 2020. "Differences in How Democrats and Republicans Behave on Twitter." *Pew Research Center*. https://www.pewresearch.org/politics/2020/10/15/differences-in-how-democrats-and-republicans-behave-on-twitter/ (December 5, 2022).

Shao, Chengcheng, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. "The Spread of Low-Credibility Content by Social Bots." *Nature Communications*: 1-9.

Silva, Bruno Castanho, and Sven-Oliver Proksch. 2021. "Fake It 'Til You Make It: A Natural Experiment to Identify European Politicians' Benefit from Twitter Bots." *American Political Science Review* 115(1): 316-322.

Somos, Christy. 2021. "What the rise of the PPC says about Canada in 2021." *CTV News*. https://www.ctvnews.ca/politics/federal-election-2021/what-the-rise-of-the-ppc-says-about-canada-in-2021-1.5596859 (November 3, 2022).

Šipka, Andrea, Anikó Hannák, and Aleksandra Urman. 2022. "Comparing the Language of QAnon-Related Content on Parler, Gab, and Twitter." *The 14th ACM Web Science Conference 2022*. Barcelona, Spain: 411-421.

Stefanov, Peter, Kareem Darwish, Atanas Atanasov, Preslav Nakov. 2020. "Predicting the Topical Stance and Political Leaning of Media using Tweets." *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*: 527-537.

Stella, Massimo, Emilio Ferrara, and Manlio De Domenico. 2018. "Bots Increase Exposure to Negative and Inflammatory Content in Online Social Systems." *Proceedings of the National Academy of Sciences* 115(49): 12435-12440.

Stieglitz, Stefan, Florian Brachten, Björn Ross, and Anne-Katharina Jung. 2017. "Do Social Bots Dream of Electric Sheep? A Categorisation of Social Media Bot Accounts." *Australasian Conference on Information Systems*. https://doi.org/10.48550/arXiv.1710.04044

Stieglitz, Stefan, and Florian Brachten. 2018. "How Powerful Are Social Bots? Understanding the types, purposes and impacts of bots in social media." *Academic Society for Management & Communication, Communication Snapshots: Social Bots in Social Media*. https://www.akademische-gesellschaft.com/en/service/publikationen/communication_snapshots/ (December 5, 2022).

Stocking, Galen, and Nami Sumida. 2018. "Social Media Bots Draw Public's Attention and Concern." *Pew Research Center*. https://www.pewresearch.org/journalism/2018/10/15/social-media-bots-draw-publics-attention-and-concern/ (September 6, 2022).

Swire, Briony, Ullrich K. H. Ecker, and Stephan Lewandowsky. 2017. "The role of familiarity in correcting inaccurate information." *Journal of Experimental Psychology: Learning, Memory and Cognition* 43(12): 1948-1961.

Tacchini, Eugenio, Gabriele Ballarin, Marco L. Della Vedova, Stefano Moret, and Luca de Alfaro. 2017. "Some Like it Hoax: Automated Fake News Detection in Social Networks." *Arxiv Preprint Arxiv:1704.07506*: 1-12. https://doi.org/10.48550/arXiv.1704.07506

Tasker, John Paul. 2022. "Conservative leadership hopeful Pierre Poilievre tables bill to ban COVID-19 vaccine mandates." *CBC News*. https://www.cbc.ca/news/politics/pierre-poilievre-vaccine-mandate-ban-1.6476419 (November 29, 2022).

Taylor, Stephanie. 2021. "Tories tell candidate who raised concern over 'climate lockdown' to take down videos." *The Canadian Press*. https://globalnews.ca/news/8150044/otoole-climate-lockdown-question/ (September 25, 2022).

Taylor, Stephanie. 2022. “Mandatory COVID-19 Vaccines a Hot-Button Issue as Election Race Kicks Off.” *The Canadian Press*. https://globalnews.ca/news/8117939/canada-election-mandatory-covid-vaccine/ (August 29, 2022).

Thelwall, Mike, Kayvan Kousha, and Saheeda Thelwall. 2021. “Covid-19 Vaccine Hesitancy on English-language Twitter.” *Profesional de la información* 30(2): 1-13.

Thompson, Elizabeth. 2022. “Disinformation, foreign interference threatening Canada’s electoral system, elections watchdog warns.” *CBC News*. https://www.cbc.ca/news/politics/canada-elections-watchdog-cote-1.6497875 (September 30, 2022).

Thota, Aswini, Priyanka Tilak, Simeratjeet Ahluwalia, and Nibhrat Lohia. 2018. “Fake News Detection: A Deep Learning Approach.” *SMU Data Science Review* 1(3): 1-20.

Timberg, Craig and Elizabeth Dwoskin. 2018. “Twitter is sweeping out fake accounts like never before, putting user growth at risk.” *The Washington Post*. https://www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/ (October 3, 2022).

Tucker, Joshua A., Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. *Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature*, 1-95. http://dx.doi.org/10.2139/ssrn.3144139

van der Linden, Sander, Costas Panagopoulos, and Jon Roozenbeek. 2020. “You are fake news: political bias in perceptions of fake news.” *Media, Culture & Society* 42(3): 460-470.

van Ditmars, Mathilde M. 2022. “Political Socialization, Political Gender Gaps and the Intergenerational Transmission of Left-Right Ideology.” *European Journal of Political Research*: 1-22. doi: 10.1111/1475-6765.12517

Van Duyn, Emily, and Jessica Collier. 2019. “Priming and Fake News: The Effects of Elite Discourse on Evaluations of News Media.” *Mass Communication and Society* 22(1): 29-48.

Varol, Onur, Emilio Ferrara, Clayton A. Davis, Filippo Menczer, and Alessandro Flammini. 2017. “Online Human-Bot Interactions: Detection, Estimation, and Characterization.” *Proceedings of the International AAAI Conference on Web and Social Media* 11(1): 280-289.

Volvoka, Svitlana, and Jin Yea Jang. 2018. “Misleading or Falsification: Inferring Deceptive Strategies and Types in Online News and Social Media.” *Companion Proceedings of the Web Conference*: 575-583.

Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. “The Spread of True and False News Online.” *Science* 359(6380): 1146-1151.

Vox Pop Labs. 2021. “Vote Compass 2021 – Canadian Federal Election.” *Vox Pop Labs & CBC/Radio-Canada*. https://votecompass.cbc.ca/canada (September 16, 2022).

Vraga, Emily K., and Leticia Bode. 2020. “Defining Misinformation and Understanding its Bounded Nature: Using Expertise and Evidence for Describing Misinformation.” *Political Communication* 37(1): 136-144.

Wagner, Claudia, Silvia Mitter, Christian Körner, and Markus Strohmaier. 2012. "When Social Bots Attack: Modeling Susceptibility of Users in Online Social Networks." *Proceedings of 2nd Workshop on Making Sense of Microposts*: 41-48.

Wang, Zuhui, Zhaozheng Yin, and Young Anna Argyris. 2021. "Detecting Medical Misinformation on Social Media Using Multimodal Deep Learning." *IEEE Journal of Biomedical and Health Informatics* 25(6): 2193-2203.

Wojcik, Stefan, Solomon Messing, Aaron Smith, Lee Rainie, and Paul Hitlin. 2018. "Bot in the Twittershere." *Pew Research Center*. https://www.pewresearch.org/internet/2018/04/09/bots-in-the-twittersphere/ (August 16, 2022).

Wojcik, Stefan, and Adam Hughes. 2019. "Sizing Up Twitter Users." *Pew Research Center*. https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/ (December 5, 2022).

Woolley, Samuel C. 2016. "Automating Power: Social Bot Inference in Global Politics." *First Monday* 21(4).

Woolley, Samuel C., and Philip N. Howard. 2016. "Political Communication, Computational Propaganda, and Autonomous Agents – Introduction." *International Journal of Communication* 10: 4882-4890.

Woolley, Samuel, C., and Douglas Guilbeault. 2018. "United States: Manufacturing Consensus Online." In *Computational Propaganda: Political Parties, Politicians, and Political Manipulation on Social Media*, eds. Samuel C. Woolley and Philip N. Howard, 186-211. New York: Oxford University Press.

Wu, Liang, Fred Morstatter, Xia Hu, and Huan Liu. 2016. "Mining Misinformation in Social Media." In *Big Data in Complex and Social Networks*, eds. My T. Thai, Weili Wu and Hui Xiong, 125-152. London: Chapman and Hall/CRC.

Wu, Liang, Fred Morstatter, Kathleen M. Carley, and Huan Liu. 2019. "Misinformation in Social Media: Definition, Manipulation, and Detection." *ACM SIGKDD Explorations Newsletter* 21(2): 80-90.

Xiao, Zhiping, Weiping Song, Haoyan Xu, Zhicheng Ren, and Yizhou Sun. 2020. "TIMME: Twitter Ideology-detection via Multi-task Multi-relational Embedding." *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*: 2258-2268.

Xu, Anbang, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. "A New Chatbot for Customer Service on Social Media." *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*: 3506-3510.

Yan, Harry Yaojun, Kai-Cheng Yang, Filippo Menczer, and James Shanahan. 2021. "Asymmetrical Perceptions of Partisan Political Bots." *New Media & Society* 23(10): 3016-3037.

Yang, Muheng, Xidao Wen, Yu-Ru Lin, and Lingjia Deng. 2017. "Quantifying Content Polarization on Twitter." *IEEE 3rd International Conference on Collaboration and Internet Computing*: 299-308.

Yang, Kai-Cheng, Onur Varol, Clayton A. Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019a. "Arming the public with artificial intelligence to counter social bots." *Human Behavior and Emerging Technologies* 1(1): 1-27.

Yang, Shuo, Kai Shu, Suhang Wang, Renjie Gu, Fan Wu, and Huan Liu. 2019b. "Unsupervised Fake News Detection on Social Media: A Generative Approach." *Proceedings of the AAAI Conference on Artificial Intelligence* 33(1): 5644-5651.

Yang, Kai-Cheng, Christopher Torres-Lugo, and Filippo Menczer. 2020. "Prevalence of Low-Credibility Information on Twitter During the COVID-19 Outbreak." *arXiv:2004.14484*: https://doi.org/10.36190/2020.16

Yang, Kai-Cheng, Emilio Ferrara, and Filippo Menczer. 2022. "Botometer 101: Social Bot Practicum for Computational Social Scientists." *Journal of Computer Social Science*: 1-18.

York, Jillian C. 2011. "Syria's Twitter Spambots." *The Guardian*. https://www.theguardian.com/commentisfree/2011/apr/21/syria-twitter-spambots-pro-revolution (September 5, 2022).

Young, Lisa, and William Cross. 2002. "Incentives to Membership in Canadian Political Parties." *Political Research Quarterly* 55(3): 547-569.

Yuan, Xiaoyi, Ross J. Schuchard, and Andrew T. Crooks. 2019. "Examining Emergent Communities and Social Bots Within the Polarized Online Vaccination Debate in Twitter." *Social Media + Society* 5(3): 1-12.

Zakrzewski, Cat, Faiz Siddiqui, and Joseph Menn. 2022. "Musk's 'free speech' agenda dismantles safety work at Twitter, insiders say." *The Washington Post*. https://www.washingtonpost.com/technology/2022/11/22/elon-musk-twitter-content-moderations/ (November 26, 2022).

Zhang, Jerry, Darrell Carpenter, and Myung Ko. 2013. "Online Astroturfing: A Theoretical Perspective." *Proceedings of the 19th Americas Conference on Information Systems* 4(1): 2559-2565.

Zhang, Menghan, Ze Chen, Xue Qi, and Jun Liu. 2022. "Could Social Bots' Sentiment Engagement Shape Humans' Sentiment on COVID-19 Vaccine Discussion on Twitter?" *Sustainability* 14(9): 1-16.

Zimonjic, Peter. 2022. "How the Liberal-NDP agreement will work and what it might mean for Canadians." *CBC News*. https://www.cbc.ca/news/politics/trudeau-singh-how-it-will-work-1.6393710 (September 18, 2022).

Appendix

Appendix A. Lists of Keywords and Hashtags

Table A1. List of Keywords and Hashtags Used to Collect the Dataset

English: trudeau, freeland, o'toole, bernier, blanchet, jagmeet singh, annamie, debate commission, reconciliation, elxn44, cdnvotes, canvotes, canelection, cdnelection, cdnpoli, canadianpolitics, canada, forwardforeveryone, readyforbetter, securethefuture, NDP2021, votendp, orangewave2021, teamjagmeet, UpRiSingh, singhupswing, singhsurge, VotePPC, PPC, peoplesparty, bernierorbust, mcga, saveCanada, takebackcanada, maxwillspeak, LetMaxSpeak, FirstDebate, frenchdebate, GovernmentJournalists, JustinJournos, everychildmatters, votesplitting, ruralcanada
French: debatdeschefs, électioncanadienne, polican, bloc, jevotebloc

Table A2. List of Keywords Used to Infer Partisan Leaning

LPC: Justin Trudeau, Liberal, Liberal Party of Canada, LPC, LPC2021
CPC: Erin O'Toole, Conservative, Conservative Party of Canada, CPC, CPC2021, Secure the Future
NDP: Jagmeet Singh, New Democrat, New Democratic Party, NDP, NDP2021
GPC: Annamie Paul, Green Party of Canada, GPC, GPC2021
PPC: Maxime Bernier, People's Party, People's Party of Canada, PPC, PPC2021

Table A3. List of Keywords and Hashtags Used to Identify Disinformation

COVID-19 Disinformation Keywords: ccpvirus, ccp virus, kungflu, kung flu, kongflu, kong flu, covidflu, covid flu, sheepnomore, sheep no more, maskhoax, mask hoax, maskshoax, masks hoax, facefreedom, face freedom, sheepwearmasks, sheep wear masks, masksareforsheep, masks are for sheep, masksdontwork, masks dont work, healthypeopledontwearmasks, healthy people dont wear masks, vaccinesdontwork, vaccines dont work, knowtherisk, know the risk, bewaretheneedle, beware the needle, vaccinedeath, vaccine death, vaxxofdeath, vaxc of death, deathshot, death shot, bigpharma, big pharma**,** saynobillgatesvaccine, say no bill gates vaccine, gatesfoundation, billgatesbioterrorist, bill gates bio terrorist, scamdemic, scamdemic2020, scamdemic 2020, scamdemic2021, scamdemic 2021, scamdemicisover, scamdemic is over, covidhoax, covid hoax, covid19hoax, covid19 hoax, covidfraud, covid fraud, plandemic, fakepandemic, fake pandemic, pandemichoax, pandemic hoax, coronahoax, corona hoax, covid1984, covid 1984, covid plot, controlavirus, whohoax, who hoax, fakecovid19, fake covid19, fake covid 19, whatcovid, what covid, covidisfake, covid is fake, covidisalie, covid is a lie, covidscam, covid scam, covidfascism, covid fascism, covid19fascism, covid19 fascism, medicalfascism, medical fascism, medicalcoercion, medical coercion, medicalapartheid, medicalapartheid, medicaltyranny, medical tyranny, insidiouscovidvaccines, insidious covid vaccines, killervaccine, killer vaccine, dictatorshipcovidvaccine, dictatorship covid vaccine, filmyourhospital, film your hospital, covidvaxexposed, ivermectinworks, ivermectin works, nurembergtrialsforcovid, nuremberg trials for covid, nuremburgtrials, nuremburg trials, nuremberg2, nuremberg 2, nuremburg2, nuremburg 2, nuremberg2021, nuremberg2021, nuremburg2021, nuremburg2021, nurembergcode, nuremberg code, nuremburgcode, nuremburg code, magiccovidvax, coviddoesnotexist, covid does not exist, virusesdonotexist, viruses do not exist, liberalismistherealpandemic, pfizerexposed, pfizer exposed, exposepfizer, expose pfizer, modernaexposed, moderna exposed, exposemoderna, expose moderna, wearelivingalie, we are living a lie, wakeup, wake up, covidenabler, covid enabler, covid19enabler, covid19 enabler, itwasneveraboutyourhealth, it was never about your health, clotshots, clot shots, covidwasaninsidejob, covid was an inside job, covidplot, covid plot, casedemic, vaccineinjured, vaccine injured, vaccineinjuries, vaccine injuries, vaccineinjury, vaccine injury, covidslavery, covid slavery, vaccineslavery, vaccine slavery, medicalslavery, medical slavery, winnipeglab, winnipeg lab
Canadian Election Disinformation Keywords: privatehealthcare, private healthcare, private o toole, privateotoole, private otoole, 2tierotoole, climatehoax, climate hoax, climatechangehoax, climate change hoax, climatescam, climate scam, globalwarminghoax, global warming hoax, globalwarmingfraud, global warming fraud, awgfraud, awg fraud, fakeclimatecrisis, fake climate crisis, climatelockdown, climate lockdown, climatelockdowns, climate lockdowns, climatelockdownsarenext, climate lockdowns are next, stolenelection, stolen election, electionfraud, election fraud, voterfraud, voter fraud, voterfraud2021, voter fraud 2021, votermanipulation, voter manipulation, electionmanipulation, election manipulation, electionintegrity, election integrity, riggedelection, rigged election, rigged, trudeaucrimes, trudeau crimes, crimeministertrudeau, crime minister trudeau, trudeaucrimeminister, trudeau crime

minister, trudeaucrimeministers, trudeau crime ministers, crimeminister, crime minister, traitortrudeau, traitor trudeau, trudeautraitor, trudeautraitor, trudeaudictatorship, trudeau dictatorship, trudeaudictator, trudeau dictator, trudeautreason, trudeau treason, talibantrudeau, taliban trudeau, trudeaucorruption, trudeau corruption, liberalcorruption, liberal corruption, conservativecorruption, conservative corruption, cpccabal, cpc cabal, liberalcabal, liberal cabal, cabal, governmentlies, government lies, trudeaucastro, depopulation, depopulationagenda, depopulation agenda, agenda21, agenda 21, agenda2030, agenda 2030, unagenda, un agenda, globalist, newworldorder, new world order, oneworldorder, one world order, communistcanada, communist canada, communisttrudeau, communist trudeau, communistliberalparty, communist liberal party, liberalfascism, liberal fascism, neoliberalfascist, neoliberal fascist, neo liberal fascist**,** holdtheline, hold the line, donotcomply, do not comply, savethechildren, save the children, trudeaupedo, trudeau pedo, pedotrudeau, pedo trudeau, manipulatedmedia, manipulated media, fakenewsmedia, fake news media, cbcfakenews, cbc fake news, cbcisfakenews, cbc is fake news, greatreset, great reset, thegreatreset, the great reset, globalreset, global reset, deepstate, deep state, deepstatecabal, deep state cabal, cpcdeathcult, trudeauhoax, trudeau hoax, trudeauscam, trudeau scam, lavscam, lavscam, beijingpuppettrudeau, beijing puppet trudeau, truanon, trudeaushiddenagenda, trudeaus hidden agenda, wefagenda, wef agenda, wefglobalist, wef globalist, coverup, cover up, steal and cheat, votingmachine, votingmachine, dominionvotingmachines, dominion voting machines, dominionvotingsystems, dominion voting systems, bringyourownpen, bring your own pen, voteinperson, vote in person, stopthesteal, stop the steal, stolethevote, stole the vote

Table A4. List of Added Keywords and Hashtags to Identify COVID-19 Tweets

Terms: covid, coronavirus, pandemic, mask, vaccine, vaccination, vaccinated, antivax, covidiot, ppe, pfizer, moderna, stayhome, stay home

Appendix B. Number of Tweet for Each Partisan Group

Table B1. Total Number of Tweets Per Group

Account Type	Partisan Leaning	Number of Tweets
Human	Liberal-Leaning	3,532,482
	Conservative-Leaning	3,182,139
Social Bots	Liberal-Leaning	574,802
	Conservative-Leaning	504,131

Appendix C. Most Popular Keywords in Disinformation Tweets Per Group

Table C1. Distribution of the Most Popular Keywords in the 2021 Canadian Election Disinformation Tweets by Humans

Liberal-Leaning		Conservative-Leaning	
Top 10 Keywords	**Count**	**Top 10 Keywords**	**Count**
trudeau	11,969	trudeau	5,788
canada	4,655	healthcare	2,732
media	3,554	canada	2,619
vote	3,424	private	2,579
globalist	3,355	vote	2,577
otoole	3,316	otoole	2,394
manipulated	3,059	erin	1,724
healthcare	2,785	conservative	1,574
election	2,630	election	1,551
private	2,578	manipulated	1,385

Table C2. Distribution of the Most Popular Keywords in the 2021 Canadian Election Disinformation Tweets by Social Bots

Liberal-Leaning		Conservative-Leaning	
Top 10 Keywords	**Count**	**Top 10 Keywords**	**Count**
trudeau	1,238	trudeau	568
canada	545	healthcare	381
media	444	private	334
vote	426	canada	414
globalist	414	otoole	303
otoole	412	erin	243
manipulated	400	vote	200
healthcare	346	globalist	196
twitter	300	liberal	182
private	293	climate	175

Table C3. Distribution of the Most Popular Keywords in the COVID-19 Disinformation Tweets by Humans

Liberal-Leaning		Conservative-Leaning	
Top 10 Keywords	**Count**	**Top 10 Keywords**	**Count**
up	6,349	up	5,061
wake	5,490	wake	4,394
trudeau	4,252	trudeau	2,339
canada	3,549	big	2,153
big	1,717	pharma	1,684
pharma	1,440	canada	1,675
medical	1,406	people	1,110
winnipeg	1,294	otoole	1,110
lab	1,267	vote	1,090
people	1,213	cpc	837

Table C4. Distribution of the Most Popular Keywords in the COVID-19 Disinformation Tweets by Social Bots

Liberal-Leaning		Conservative-Leaning	
Top 10 Keywords	**Count**	**Top 10 Keywords**	**Count**
up	906	up	758
wake	827	wake	646
trudeau	426	big	394
big	350	pharma	309
canada	286	trudeau	185
pharma	266	canada	179
people	197	people	145
medical	145	drugs	130
vaccine	141	vaccine	119
otoole	135	otoole	111

Appendix D. Partisan Divisions with Different Thresholds

Table D1. Partisan Divisions with a Threshold of 0.3

Account Type	Partisan Leaning	Number of Users	Number of Tweets
Human	Liberal-Leaning	245,959	2,854,633
	Conservative-Leaning	326,795	2,574,260
Social Bots	Liberal-Leaning	84,646	1,252,651
	Conservative-Leaning	113,568	1,112,010

Table D2. Partisan Divisions with a Threshold of 0.5

Account Type	Partisan Leaning	Number of Users	Number of Tweets
Human	Liberal-Leaning	275,363	3,225,491
	Conservative-Leaning	369,524	2,939,210
Social Bots	Liberal-Leaning	55,242	881,793
	Conservative-Leaning	70,839	747,060