#HASHTAGS AS A MEANS FOR POLITICAL ACTION

EXPLORING THE IMPACT OF HASHTAG ACTIVISM ON POLITICAL PROCESSES USING GEOLOCATED SOCIAL MEDIA DATA

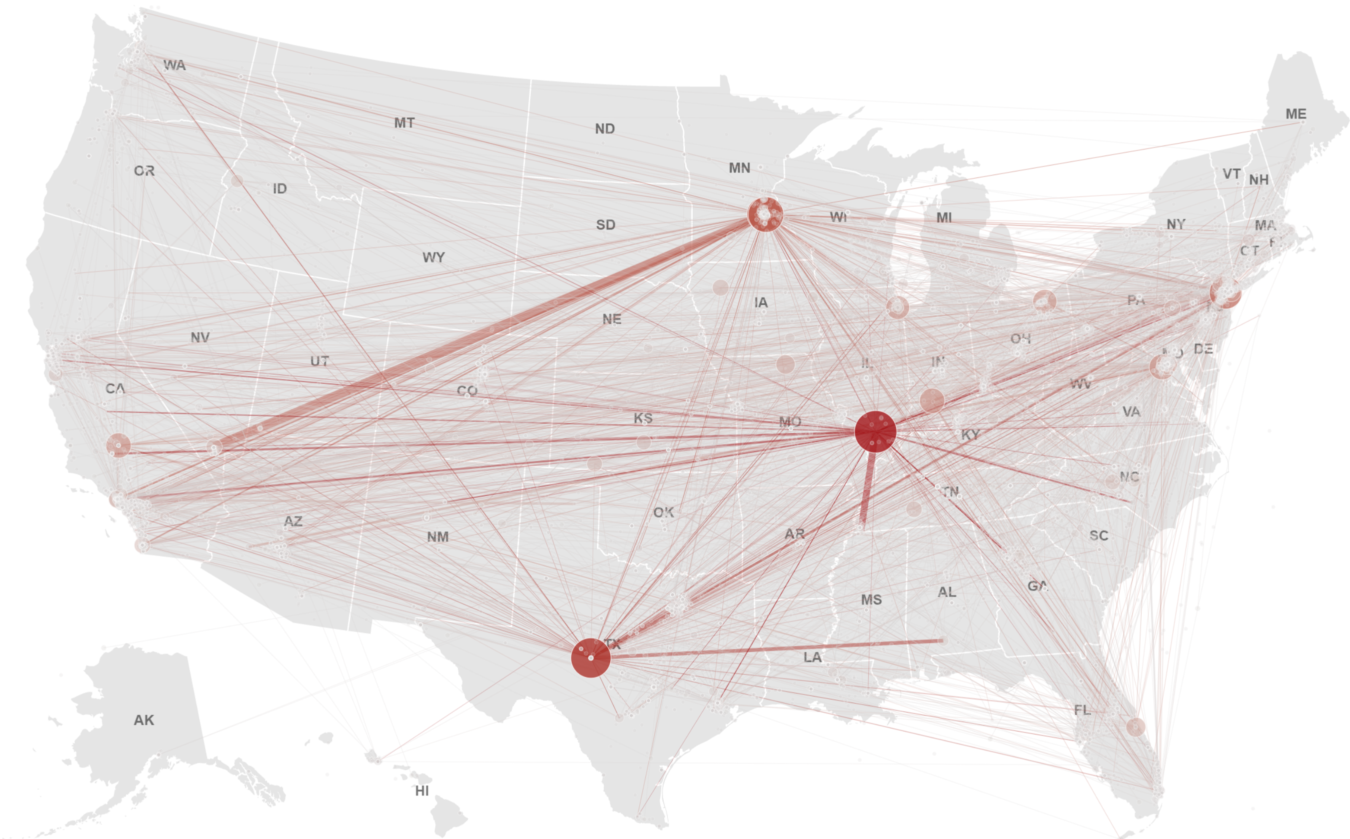

Master Thesis

as partial requirement for a
Master of Science in Engineering and Policy Analysis,
to be defended publicly on Friday 20th of August, 2021.

by

Leonardo Antonio NICOLETTI

Faculty of Technology, Policy, and Management,
Delft University of Technology, The Hague and Delft, Netherlands.

Thesis committee:

Dr. T. Verma,	Delft University of Technology, first supervisor
Prof. dr. M.E. Warnier,	Delft University of Technology, chair and second supervisor
Dr. P. Santi,	Massachussets Institute of Technology, external advisor

Keywords: social movements, big data, physical protest, social change, public policy, machine learning, policing

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

Acknowledgements

This project could not have been possible without the support and advice of many people. First, I would like to thank my thesis committee. To my supervisor, Dr. Trivik Verma, thank you for not only providing me with crucial perspectives and knowledge relating to many of the theoretical and statistical concepts employed in this research, but also for believing in my ideas and sometimes unfledged visions. Your continued mentorship during this masters degree as a whole has been central to my development as a professional, and to the discovery of many of my professional interests. To my chair, Dr. Martijn Warnier, thank you for the precise and constructive feedback you provided me with during our meetings. To my second supervisor, Dr. Paolo Santi, thank you for finding interest and value in this project and for putting me in the right direction when I encountered roadblocks along the way.

I would like to thank the Computational Urban Science and Policy (CUSP) lab at TU Delft and all of its members, where I received the professional and emotional support I needed to carry this project to fruition. More generally, I would like to thank my peers and professors from the Engineering and Policy Analysis masters program, from whom I accumulated a wealth of knowledge and formative experiences.

To the MIT Senseable City Lab, thank you for opening your doors to me, for sharing your resources with me, and for allowing me to be part of your incredible research team, where experts from so many disciplines come together and continuously innovate in the field urban science.

To Mikhail Sirenko, thank you for being a friend, teacher, and mentor to me during the past two years. Your guidance has been fundamental for my personal and professional development. I would not have challenged myself the way I did, had I never met you.

Last, I would like to thank my parents, Giuseppe and Camilla, and my brother, Orlando. The unconditional love and support you have given me throughout this project and my life in general has been my greatest source of strength, stability, and inspiration. Mamma, Papa, Orlando, grazie.

Leonardo Nicoletti
Vancouver, BC, July, 2021

Preface

Personal data, code, models, and advanced computation are exciting tools for us scientists. In so little time, their increased accessibility has empowered us to study the world in revolutionary ways. In particular, these tools have created a bridge across seemingly distant scientific fields, such as computer science and sociology, and have allowed us to deepen our understanding of complex human behaviours and social structures. Based on the analysis and modelling of geo-located social media data, the research presented in this thesis finds that the way in which we use social media in times of crisis is a strong indicator of where, when, and with what intensity we choose to take the streets and protest against injustice.

While this is a fascinating research outcome, it is also worry-some because we live in a world where powerful institutions such as law enforcement agencies often use their power to oppress the most vulnerable and disenfranchised groups. In the United States, research has demonstrated that the use of force by police is among the leading causes of death for Black and Brown men, and that racially-rooted police violence is on the rise. Similar bias in police practices has been observed in many other countries around the world. Additionally, as paradoxical as it sounds, state-sanctioned violence of this kind is widespread during physical protests, where the very people that demand justice and accountability for people like George Floyd are oppressed further. In this light, 2020 and 2021 were no exceptions, as it was documented during the #JusticeForGeorgeFloyd social movement in the United States, the case study of this research, and as it is being documented during the #ColombiaResiste social movement which is happening as this paragraph is being written.

In this sense, the social impact of the research presented here may not turn out to be as exciting as its scientific impact. As proposed in the discussion of this study, our findings may serve to empower activist groups with tools to better allocate their resources and to communicate with local communities more efficiently. At the same time, however, the impact of this research could turn out to be more grim, where the models and insights of this research could equip law enforcement agencies to respond to physical protests more rapidly and to use more oppressive tactics against protesters. As scientists, it is our responsibility to conduct and communicate novel scientific research, but we have little power at our disposal to ensure that such research is used solely for the betterment of society. Such task pertains to policy makers and (to a lesser extent) private companies, who do have regulatory and veto power at their disposal.

Fortunately, there are effective ways in which policy makers and social media companies can avoid unethical applications of our findings. With regards to law and policy making, regulatory policies can be put in place for law enforcement agencies to respect the personal privacy of individuals. For example, such policies could limit the use of location data from social media by law enforcement agencies for the purpose of tracking demonstrations and ongoing physical protests. With regards to social media companies such as Twitter, usage policies can be put in place to limit access of specific data attributes to law enforcement agencies. For example, just like Twitter has recently increased access to its data for academic purposes, it may limit access to its data for law enforcement purposes (i.e. by restricting access to location information to law enforcement developers). By creating boundaries to access, the use of data for social good can be ensured and activist groups can safeguard their opportunity to lawfully voice their concerns for more progressive policies through physical protest.

Leonardo Nicoletti
Vancouver, BC, July 2021

Executive Summary

Throughout history, social movements have often been catalysts for radical societal change. In the past two decades, hashtag activism, the use of social media platforms for internet activism, has become a driving force behind the development of social movements across the world. From #MeToo to #IdleNoMore and most recently #JusticeForGeorgeFloyd, social media has been used strategically by activists to mobilize communities to come together and protest against different forms of injustice. In the field of social movements science, a large body of research has studied the role of hashtag activism for the formation of social movements, but less efforts have been allocated towards the study of the spatio-temporal relationship that exists between hashtag activism and political processes. In other words, in spite of the a-spatial nature of social media, ***can the study of hashtag activism help us understand human behaviours and societal processes that occur off-line, in the physical space?***

Such conundrum is the basis of the research in this master thesis, where a data-driven framework is implemented to investigate the spatio-temporal relationship between hashtag activism and two important political processes: physical protest activity and legislative action. Through a combination of time series analysis, regression, geo-spatial analysis, and machine learning, and the use of a large Twitter data-set of geo-located social media posts, such relationships are quantified, modelled and visualized for the 2020 #JusticeForGeorgeFloyd social movement. First, time series analysis is used to measure the temporal relationship between hashtag activism, physical protest, and legislative action. Second, geo-spatial analysis is used to compare the nature of this relationship across different spatial resolutions. Finally, regression analysis and deep learning are used to model, generalize and forecast such spatio-temporal relationships.

The outcomes of this research are two-fold. First, in the case of #JusticeForGeorgeFloyd, it was found that at the national, state, and county scale, hashtag activism bears a very strong positive temporal relationship with physical protest activity. Through the statistical modelling of this relationship, it is possible to predict the intensity of national, state, and county level protest activity on a given day based on the intensity of state and county level hashtag activism the day before. At the county level, it was found that in some "outlier" counties the temporal relationship between hashtag activism and physical protest is characterized by a disproportional amount of physical protest compared to the intensity of hashtag activism, while in other counties this relationship is characterized by a disproportional amount of hashtag activism compared to the intensity of physical protests. In order to make sense of this tendency this research proposes the *Mobilization Synergy Index (MSI)*, which makes it possible to statistically quantify this trade-off, and to visualize its spatial distribution in an intuitive way. Through this index, both local policy makers and activists could gain valuable insight into their communities and as a result adopt strategies that better reflect community needs and concerns.

Second, the analysis of instances of hashtag activism in relation to legislative responses revealed that states in which people engaged more in hashtag activism on average during #JusticeForGeorgeFloyd also experienced more legislative responses related to policing. Additionally, results from temporally disaggregated national level and state level data suggests that the number of legislative responses on a given day bears a very similar positive relationship with each of hashtag activism activity and physical protest activity taken at a specific time lag prior to the day in which the legislative responses occurred. At the state level, it was found that while both hashtag activism and physical protest activity result in future legislative responses, this temporal relationship is highly variable by state. Such tendency may serve as an indicator of state-level political responsiveness during times of crisis.

To the best of our knowledge, this is the first national level quantitative study aimed at measuring the temporal relationships between hashtag activism, physical protest activity, and legislative action, at various spatial resolutions. Therefore, the study provides several contributions to the field of social movements research and, more broadly, the computational social sciences.

CONTENTS

List of Figures

LIST OF TABLES

1

INTRODUCTION

1.1. BACKGROUND

MAY 25th 2020 represented a key moment for American society. The death of George Floyd in circumstances that decisively pointed to police brutality and ethnic discrimination ignited widespread protests in Minneapolis and many other US cities (The New York Times, 2020). During the months of June, July, and August, more than 20 million Americans marched across the United States in the names of hashtags #JusticeForGeorgeFloyd and #BlackLivesMatter, and inspired people in more than 2000 cities and towns worldwide to march in solidarity (The New York Times, 2020). At the end of June and to this day, a bold policy proposal emerged from the movement and started circulating on social media: #DefundThePolice. What began as a general call for justice quickly evolved into a demand to policy makers at the federal, state, and local level to take clear action against systematic discrimination and questionable police practices in the country. This demand, fueled by the #DefundThePolice hashtag, has already progressed in a number states and cities, where policies resulting in the reallocation of public funds away from police forces and towards other uses have been or are being implemented (Jacobs *et al.*, 2021). This powerful wave of activism has not only challenged American society to rethink the role and utility of deeply ingrained institutions such as law enforcement, but it had also fueled Joe Biden's presidential campaign, who not only supported the #BlackLivesMatter movement, but who also committed to tackling racial discrimination and police violence during his term (Biden, 2020; Biden, 2020). Biden was elected in November, 2020.

1.2. KNOWLEDGE GAP AND RESEARCH QUESTIONS

Throughout the past two decades, hashtag activism, "the creation and proliferation of online activism stamped with a hashtag" (Jackson *et al.*, 2020), has become a crucial strategy for activists trying to influence social change (Jackson *et al.*, 2020). Scholars within this area of research have shown the role of social media in acting as: a tool for some demographics (i.e. youth, underrepresented groups) to shape political discourses (Carney, 2016; M. Li *et al.*, 2021); a tool to empower and mobilize publics (Xiong *et al.*, 2019); a tool to frame and re-frame social movements (Ince *et al.*, 2017); a tool to maintain the vitality of a social movement (Tan *et al.*, 2013); an environment that can contribute to the polarization of political opinions (Grover *et al.*, 2019). Other studies have found that social media can be studied to understand the structure of social movements, key players and the strategies responsible for the effectiveness of social movements (Burns and Eltham, 2009). However, research that focuses on the political consequences of hashtag activism is severely limited as (1) scholars have barely explored the relationship between hashtag activism and concrete political action such as the introduction of new legislation and (2) existing quantitative research on the relationship between hashtag activism and other forms of political action (i.e. physical protest) is inconclusive. Finally, a majority of the literature has been limited by the use of small-to-medium samples of social media data carrying little to no information regarding the geographical processes involved in social media movements.

In this context, scholars who study social movements have stressed that the geographical characteristics of a social movement can be important determinants of its political effectiveness (Amenta *et al.*, 2010). In a review of 54 social movements, of which 43 took place in the United States, Amenta *et al.* (2010) note that "for a movement to achieve political influence" the mobilization of large numbers of people is sometimes not enough, as "state actors need to see it as potentially facilitating or disrupting their own goals". In district-based electoral systems such as the U.S. Congress, social movements have to fit this context to have long-term political implications, which means "gaining a wide geographical presence" (Amenta *et al.*, 2010). Nevertheless, empirical evidence demonstrating the role of spatial patterns involved in social movements in enabling political action is lacking, and scholars have barely explored the role of location based hashtag activism in contributing to a social movement's political effectiveness (Ince *et al.*, 2017). For instance, while there is some evidence that hashtag activism may facilitate physical protest activity (Bonilla and Rosa, 2015; Boling, 2020; Juris, 2012), and that the mobilization of people for protests during social movements is effective in influencing political action (Enos *et al.*, 2019; Weldon *et al.*, 2011; Lipsky, 1968), whether increased social media activity in a specific geographical location is likely to trigger political responses in that specific location is unclear.

In this context, this study aims to explore the political impact of hashtag-activism at different administrative levels in the United States (i.e. city, county, and state). By making use of a large collection of geo-located tweets collected during the 2020 #JusticeForGeorgeFloyd social movement, data on legislative responses to policing from the National Conference of State Legislatures (NCSL) (NCSL, 2020), and data on physical protest activity from The Armed Conflict Location & Event Data Project (ACLED) (Raleigh *et al.*, 2010), the aim is to quantify the spatial and non-spatial relationships that may exist between hashtag activism and different forms of political action towards police brutality in the United States. In doing so, this research uncovers general statistical trends regarding the relationship between social media activity and place-specific political processes (e.g. physical protest, legislative change). The relevant research question can be formulated as follows:

What are the political impacts of location-based hashtag activism during and after social movements?

Throughout this research two sub-questions are considered in order to address this main question:

1. *What is the relationship between location-based hashtag activism and physical protest activity?*
2. *What is the relationship between location-based hashtag activism and legislative action?*

1.3. RESEARCH APPROACH

The empirical investigation around the main research question follows a mixed methods approach, involving a case study and quantitative analysis.

Throughout (recent) history, there have been a number of social media movements that have influenced public policy (Amenta *et al.*, 2010). In this research, we choose to focus on the #JusticeForGeorgeFloyd movement as a case study to investigate the role of location-based social media activity in contributing to the political mobilization of people through physical protest, and to state-level legislative responses aimed at reforming the U.S. police institution for reducing discriminatory practices within the activity of policing. The recency, political relevance, and vast amount of data associated with this movement make this case study highly relevant for addressing our research question.

Based on this case study, the research is implemented via a data-driven quantitative approach. By analyzing a large collection of geo-located tweets from the #JusticeForGeorgeFloyd movement through quantitative methods, insight can be gained into specific aspects of the research question. More specifically, the use of time series analysis and forecasting provides insight into the temporal relationships that may exist between hashtag activism and different forms of political processes; the use of regression analysis provides insight into the statistical significance of those relationships; and the use of geo-spatial analysis provides insight into the spatial relationships that may exist between hashtag activism and different forms of political processes at different spatial scales.

The use of a mixed methods research approach involving both a case study and quantitative methods is advantageous for several reasons. First, there are large amounts of data available for the case study chosen, thus providing a strong empirical foundation for this research. Second, the case study chosen is highly relevant for the current socio-political climate of many countries globally. Finally, the use of Twitter data provides a myriad of attributes that are relevant for investigating the research sub questions (i.e. network information, spatial information, lexical information, temporal information). Figure section 1.3 summarizes the structure and logical flow of this thesis research.

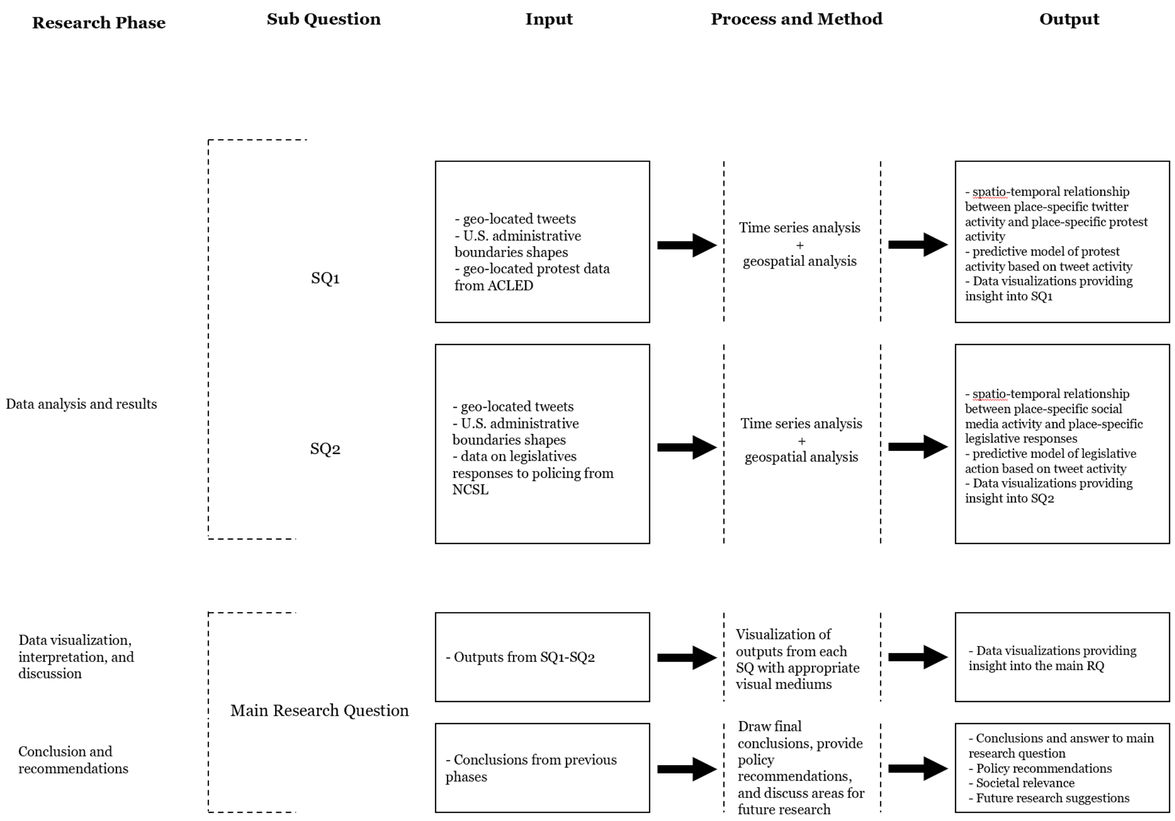

Figure 1.1: Research Flow Diagram

1.4. Report Structure

This master thesis report is structured according to six chapters. Chapter 1 provides contextual background to the research, outlines the knowledge gaps that drive it, and presents the chosen methodological and empirical research approach. The chapter concludes with a description of the structure and flow of this report.

Chapter 2 provides the theoretical foundation for this thesis research, introduces the case study, and discusses the knowledge gaps that this research addresses. The chapter first discusses the research landscape surrounding key topics directly related to this research, namely (a) hashtag activism and social movements (b) hashtag activism and physical protest and (c) hashtag activism legislative action. In doing so, important scientific contributions are described and relevant knowledge gaps are discussed. Second, the chapter introduces the case study of #JusticeForGeorgeFloyd and discusses its relevance for use in this research and for answering the main research question. Finally, the chapter describes the scientific and theoretical links between the highlighted knowledge gaps and the case study, and introduces the theoretical framework that drives this research.

Chapter 3 describes the data and methods utilized to leverage the #JusticeForGeorgeFloyd case study and address each sub-question of this research. In the first section of the chapter, the data required to conduct the research is introduced. For each data source the collection process is highlighted, the data structure is illustrated, and the pre-processing steps used to prepare them for analysis are described. In the second section, the computational procedures and modelling techniques used in the research are introduced. As such, specific methods for time series modelling, regression analysis, machine learning, and geo-spatial analysis are described and their implementation in relation to the research problem is clarified.

Chapter 4 presents the results of the analysis conducted to investigate the two sub-questions of this research. In applying each step of the proposed experimental framework presented in chapter 3 to the case study of #JusticeForGeorgeFloyd, insight into the relationship between hashtag activism and two forms of political processes - physical protest and legislative action - are presented. First, the spatio-temporal relationship between hashtag activism and physical protest is explored. For this, time-series modelling, regression analysis and geo-spatial analysis steps are applied to explore how hashtag activism and physical protest are related in time and space. This section concludes with the development of a novel index, aimed at providing an accurate and granular picture of where hashtag activism is most likely to translate into physical protest activity and what that may mean for local policy makers. Second, the spatio-temporal relationship between hashtag activism and legislative action is explored. For this, time-series modelling, regression analysis and geo-spatial analysis steps are applied to explore how hashtag activism and legislative action are related in time and space.

Chapter 5 discusses the main results presented in chapter 4 and the possible implications of this research within the realms of academia, policy, and activism. The chapter first discusses each result in (a) the scientific context by outlining if and how it constitutes a contribution to the broader research landscape; (b) the socio-political context by outlining its societal meaning and implications for policy making. Second, this chapter discusses the scientific limitations of the study and describes areas of future research to address these limitations.

Chapter 6 summarizes answers to the research sub-questions and the main research question, and concludes the study with a reflection on the ethical implications of this research, and in which ways these can be addressed.

2

Literature Overview

Hashtag activism is a relatively recent societal phenomenon. Starting in 2008, adoption of social media platforms such as Twitter and Facebook became increasingly widespread for private and public entities, non-profit organizations, and the public as a whole (Curtis *et al.*, 2010). In 2011, the use of Twitter and the hashtag #IranElection as a strategic tool for Iranians to protest a contested election and mobilize politically exemplified the political power and implications of hashtag activism (Jackson *et al.*, 2020). Since then, scholars have defined the concept of hashtag activism as the “act of fighting for or supporting a cause with the use of hashtags as the primary channel to raise awareness of an issue and encourage debate via social media” (Tombleson and Wolf, 2017) and more broadly the “the creation and proliferation of online activism stamped with a hashtag” (Jackson *et al.*, 2020). Throughout the past decade, hashtag activism has played a role in numerous social movements across the world, and scholars have used the increasingly available data generated by social media platforms to study the mechanisms of hashtag activism and the extent of its socio-political implications.

The focus of this research is on the relationship between hashtag activism and two concrete forms of political processes: physical protest and legislative action. Physical protest is defined as a group of strategies for "relatively powerless groups" (Lipsky, 1968) to take collective action through physical means and obtain bargaining power in the political process (Lipsky, 1968; Eisinger, 1973). Physical protest has been, and continues to be, an important political resource for minority and marginalized groups to disrupt routine activity and apply political pressure on public authorities and policy-makers (Lipsky, 1968; Eisinger, 1973). Legislative action is defined as any official change in the legislative process such a the introduction of new legislation by representatives, the implementation of new legislation into law, or the reform of an existing piece of legislation.

In this context, this chapter lays the theoretical foundations that guide this research and discusses state of the art literature regarding three relevant areas of study:

1. *Hashtag Activism and Social Movements*
2. *Hashtag Activism and Physical Protest*
3. *Hashtag Activism and Legislative Action*

2.1. Hashtag Activism and Social Movements

According to Della Porta and Mattoni (2015), social movements are formed by networks of informal relations between numerous groups and individuals with shared beliefs and solidarity. During a social movement, these networks are mobilized around contentious themes through the frequent use of various forms of protest with the objective of enabling different forms of social change (Della Porta and Mattoni, 2015). Thus, the functioning of social movements is two fold, where (1) internal communication across social networks serves to strengthen collective identities and solidify value systems and (2) external communication with outside actors such as potential allies, bystanders and opponents serves to gather further support and challenge institutions (Della Porta and Mattoni, 2015).

Historically, mass media has played a central role in the diffusion of information during social movements. Activists have leveraged mass media platforms to engage in political discourse and attempt to influence public opinions and gather support across the public and political spheres (Della Porta and Mattoni, 2015). However, because of the constraints associated with journalism (i.e. selection and descriptive biases), the influence that activists could exert through the use of mainstream media has been, in the past, limited (Della Porta and Mattoni, 2015). In this context, through the recent adoption of social media, hashtag activism has become a crucial strategy for activists to shape social movements in drastically new ways and influence social change (Jackson *et al.*, 2020). In greatly facilitating interactions between organizations and the public, social media platforms foster environments of political discussion and participation where rapid information diffusion beyond spatial boundaries facilitates the creation of communities that share values, characteristics, and interests (Kent, 2013; Fuchs and Sandoval, 2014). In addition, through the use of social media data, scholars have demonstrated that hashtag activism is an effective tool for underrepresented social groups (e.g. ethnic minorities, youth, women) to shape political discourses and advance social movements (Carney, 2016; M. Li *et al.*, 2021; Xiong *et al.*, 2019). For example, in an analysis of social media posts data from Twitter collected during the #MeToo social movement, Xiong *et al.* (2019) found that hashtags functioned as a tool to empower and mobilize publics around the issue of sexual assault, and as a mechanism for Social Movement Organizations (SMOs) to effectively engage with the public. Xiong *et al.* (2019) concludes that such mechanisms allow for a more efficient organization of social movements. Through leveraging other case studies, scholars have demonstrated that hashtag activism is also effective for framing and re-framing social movements, and to maintain the vitality of social movements over time (Ince *et al.*, 2017; Tan *et al.*, 2013). Finally, there is consensus that, in the context of social justice and social change, decentralized networks such as those found on social media have allowed activist groups to position themselves as key actors to influence public discourse and call for collective action (Benford and Snow, 2000).

In short, by creating novel opportunities to influence the structure and reach of social movements, hashtag activism is found to play a central role in the development of social movement characteristics (i.e. networks of informal relations, shared values, mobilization around contentious themes, calls for protest). However, most research surrounding hashtag activism during social movements has placed little focus on the spatial relationships that may exist between hashtag activism and social movements. Thus, the question of whether hashtag activism can explain the geographical patterns of social movements remains largely disputed.

2.2. Hashtag Activism and Physical Protest

Physical protest is a key characteristic of social movements. Activists use physical protest as a non-conventional form of political action to disrupt daily routines, attract public attention and apply political pressure on protest targets (e.g. policy-makers) (Della Porta and Mattoni, 2015). While scholars agree that physical protest is a political resource for social movements (Enos *et al.*, 2019; Weldon *et al.*, 2011; Lipsky, 1968; Heaney, 2020), there is no consensus regarding the potential relationship that may exist between hashtag activism and physical protest (Brantly, 2019; Howard *et al.*, 2011; Shirky, 2011). This lack of consensus partly originates from the dearth of empirical research on this topic, and partly from the fact that the little empirical research that exists is limited to isolated case studies.

The hypothesis that hashtag activism may bear a relationship with physical protest stems from the general consensus from social movements research that the “best predictor of attendance of a protest is not ideology or other substantive factors but whether an individual already knows someone else going to the protest” (Schussman and Soule, 2005; Tufekci and C. Wilson, 2012). In the digital era, such tendency could suggest that protest organizers coordinate protests on social media and that protest attenders learn about the logistics of a physical protest through their social media networks. In a survey of participants from the Tahrir Square protests in Egypt Tufekci and C. Wilson (2012) discovered patterns that confirm this hypothesis. During the 2011 Egyptian revolution, not only was social media the primary source of information regarding physical protests and their associated logistics, it also offered pieces of information that authorities could not control and that were key to individuals’ decisions to participate in physical protests (Tufekci and C. Wilson, 2012). The authors conclude that, during the Tahrir Square protests, use of social media increased the probability that an individual participated in physical protests (Tufekci and C. Wilson, 2012). Similarly, through leveraging big data from multiple social media platforms generated during Ukraine's 2013-2014 Euromedian protests, Brantly (2019) found evidence suggesting that social media plays a role in shifting protests from the digital to the physical space. The strongest correlations were found with social media activity taken at time t-1 and physical protest at time t. Although statistically significant, such evidence was associated with a wide margin of error (Brantly, 2019). Following a similar methodology Bastos *et al.* (2015) found that, for the case of the Indignados, Occupy, and Vinegar social movements, Twitter and Facebook activity could be informative of next day's protest activity. This study, however, found this relationship to be bi-directional, in that physical protest could also be predictive of next day's social media activity. Due to its limited geographical scope, it also could not test whether these relationships exist across different geographical locations of a country. Moreover, similar research conducted by S. L. Wilson (2017) with the use of geo-located tweets collected for the same case study generated conflicting results: while social media activity increased nationwide during mass protests, it decreased at sites of mass protest. More recently, a study of twitter data related to the 2016 Black Lives Matter (BLM) protests in Charlotte, NC, found a unidirectional causal relationship where physical protests appear to drive tweet activity (Karduni and Sauda, 2020). In addition, the authors identified two distinct communities who engage in social media during crisis events in the same way; 1) those who use both social media and (physical) urban space and 2) those who are connected to the movement only through social media (Karduni and Sauda, 2020).

Overall, the study of the relationship between hashtag activism and physical protest is inconclusive because of little empirical research and limited case studies. Additionally, most of the literature within this area of research studies this relationship through a broad lens and does not consider the geo-spatial relationships that may exist between digital and physical protests at sub-national levels (Brantly, 2019).

2.3. Hashtag Activism and Legislative Action

Several scholars have dedicated their research efforts towards understanding how social movements have influenced politics in terms of legislative action, the framing of political discourse, institutional change and reform. Scholars within this area of research have demonstrated that social movements in general almost always result in some sort of political response (Amenta *et al.*, 2010); that politicians often leverage social movements to frame their policy discourses and connect to voters (Zhuravskaya *et al.*, 2020); and that the mobilization of people for protests during social movements is effective in influencing political action (Enos *et al.*, 2019; Weldon *et al.*, 2011; Lipsky, 1968). Recently, more literature on the role of networked social movements and hashtag activism in enabling political change has become available. For example, Ouassini (2021) examined the case of Amina Fali, a young woman in Morocco who committed suicide after being forced to marry her rapist under Article 475 of the Moroccan penal code. The author presents qualitative and anecdotal evidence that, through hashtags such as #RIPAmina, social media channels Twitter and Facebook were instrumental in pressuring the Moroccan government to repeal Article 475 (Ouassini, 2021).

As such, while scholars generally agree that hashtag activism can have many political consequences including legislative action (Shirky, 2011; Earl and Kimport, 2011; Bennett and Segerberg, 2012), supporting evidence through quantitative research is mostly lacking. In 2016, a study of millions of social media posts collected from twitter during the 2014 BLM social movement in Ferguson, MO, measured the impact of hashtag activism on "elite responses" in the form of social media acknowledgements by political leaders accounts (e.g. president, members of congress etc.) (Freelon *et al.*, 2018). By providing empirical evidence that repeated social media participation rates by BLM users was accurate in predicting elite responses, this study was the first to demonstrate that social movements can attract elite attention through social media (Freelon *et al.*, 2018).

Overall, empirical research aimed at quantifying the potential relationship between hashtag activism and legislative action is very scarce. In this light, scholars have stressed that researchers should leverage larger sources of data and geo-information to explore the role of hashtag activism in contributing to the political outcomes of social movements' (Ince *et al.*, 2017).

2

2.4. CASE STUDY: #JUSTICEFORGEORGEFLOYD

2.4.1. THE PROBLEM OF POLICE BRUTALITY IN THE UNITED STATES AND THE BIRTH OF #BLACKLIVESMATTER

In the past decade, police brutality has harnessed much public attention in the United States. With both the increasingly common use of smartphones and social media across the U.S. population, the world has witnessed a constant flow of heartbreaking videos documenting the killing of unarmed Black and Brown people at the hands of U.S. police forces. At the same time, U.S. public authorities have remained systematically opaque with regards to the publicly available information regarding police involved deaths (Edwards *et al.*, 2019; Burghart, 2020). As a result, the past decade has also seen the development of a serious research effort dedicated to (1) building reliable and accurate databases on all U.S. police involved deaths since the year 2000 and (2) using such databases to generate robust statistical evidence of trends in police violence across U.S. cities and demographics (Edwards *et al.*, 2019; Burghart, 2020). This research effort has been crucial in the framing of police brutality as a grand challenge for U.S. society and in increasing the legitimacy of civil rights activist groups such as #BlackLivesMatter. Indeed, by leveraging the Fatal Encounters database (Burghart, 2020), a landmark study published in the Proceedings of the National Academy of Sciences of the United States of America (Edwards *et al.*, 2019) has demonstrated that police violence is among the leading causes of death for young men of color in the United States. Over the course of their life time, Black American men have a 1 in 1000 risk of being killed by the police, more than double the risk that white American men face (Edwards *et al.*, 2019). Compared to white Americans, Native Americans and Latinx people also face a higher risk of being killed by the police (Edwards *et al.*, 2019). Edwards *et al.* (2019) also demonstrate that police involved deaths per 100,000 have increased for all demographics since the year 2000, and with this increase the risk gap between white Americans and Americans of color has grown wider.

In the wake of this grim reality, a powerful wave of activism dedicated to ending racial inequality and state sanctioned violence was born. In 2013, after the acquittal of 17 year old Trayvon Martin's murderer, #BlackLivesMatter, "a global organization in the US, UK, and Canada, whose mission is to eradicate white supremacy and build local power to intervene in violence inflicted on Black communities by the state and vigilantes" (Matter, n.d.), was created. The placement of a hashtag at the beginning of its name reflects an crucial characteristic of the #BlackLivesMatter social movement: the use of hashtag activism as a primary strategy for political mobilization (Ray *et al.*, 2017). Indeed, the #BlackLivesMatter social movement has continuously leveraged social media (and mostly Twitter) to empower youth of color in the United States and elsewhere to engage in political discourse and play an active role in debates that initiated after the deaths of unarmed black men at the hands of the police (Carney, 2016). This has been the case since the killing of Trayvon Martin in 2012 to today, with the killing of George Floyd in 2020. In this context, another key characteristic of the #BlackLivesMatter movement is the use of names and grievance as a way to build and strengthen collective identities (Ray *et al.*, 2017). Through hashtags such as #TrayvonMartin, #TamirRice or #EriceGartner the #BlackLivesMatter movement humanises victims of police brutality and thus amplifies the severity of the problem and the urgency through which political change must happen (Ray *et al.*, 2017; Lebron, 2017).

2.4.2. BACKGROUND ON #JUSTICEFORGEORGEFLOYD AND RELEVANCE FOR RESEARCH

The murder of George Floyd by officer Derek Chauvin on the 25th of May 2020 and the events that followed is the most recent example of the use of hashtag activism as a key strategy for activists to call for social change. Indeed, daily counts of #BlackLivesMatter tweets exploded shortly after the murder of George Floyd and remained high for numerous months after (Priniski *et al.*, 2021). At the same time, thousands of physical protests were organized across the country for people to march in solidarity of George Floyd and demand immediate change regarding police reform, police accountability, and racial injustice (The New York Times, 2021). At the end of 2020 and over the past year, several U.S. government bodies at the state and local level have addressed such demands, developing legislation that ranged from curbing specific types of use of force by police officers, to reallocating a portion of public funds away from police forces and towards other uses (Jacobs *et al.*, 2021). While it is intuitive that physical and hashtag activism may have facilitated such political changes, the extent to which hashtag activism may have played a role in the political mobilization of people across the United States and the development of legislation is unclear.

In this context, the social significance of the #JusticeForGeorgeFloyd social movement, the intense use of social media during the movement, its geographic coverage in terms of physical protest, and the still-ongoing political activity related to the movement make this case highly relevant for the study of the political implications of hashtag activism. Additionally, the availability of high quality data regarding social media activity, physical protest, and legislative change represents an opportunity to leverage quantitative methods in this study and contribute to the growing body of evidence-based literature in the field social movement research.

2.5. CONCEPTUALIZATION

From the literature review, it is clear that scientific understanding of the relationship between hashtag activism and political processes is still in its infancy. On a theoretical level, scholars agree that social movements and, more recently, networked social movements have real political consequences. On an empirical level, however, there is a lack of evidence-based research to support this claim, and the set of evidence-based research that does exist has been inconclusive. More specifically, such research is in disagreement with regards to the statistical relationship between hashtag activism and physical protest, and is completely lacking with regards to the relationship between hashtag activism and legislative change. Finally, few scholars have incorporated a geo-spatial dimension in their study of the socio-political implications of hashtag activism and social movements in general (Brantly, 2019; S. L. Wilson, 2017).

In building on top of the theoretical understanding that (networked) social movements can have political consequences, this research aims to leverage a relevant case study through a quantitative approach to provide an evidence based understanding of the political implications of hashtag activism. As such, fig. 2.1 illustrates the theoretical framework that guides this research. Through this framework, and the use of the #JusticeForGeorgeFloyd social movement as a case study, the relationship between hashtag activism, physical protest, and legislative change are modelled and quantified.

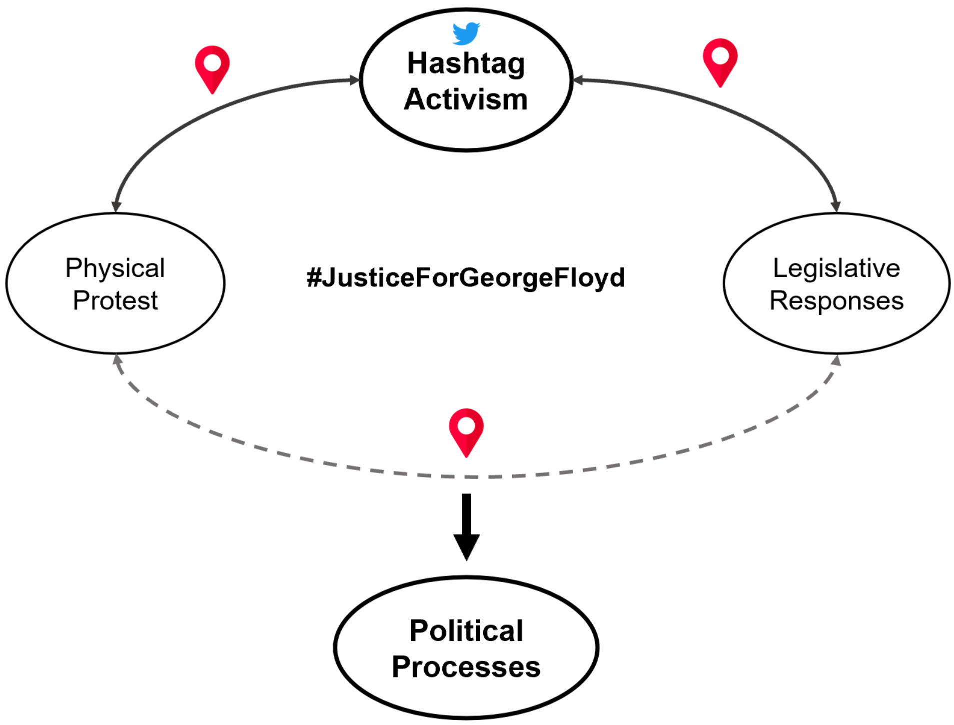

Figure 2.1: Theoretical Framework

3

METHODOLOGY

QUANTIFYING the impacts of location-based hashtag activism on physical protest and legislative action requires an experimental design that takes into consideration key problem-specific dimensions. Each of hashtag activism, physical protest, and legislative action describe complex societal mechanisms that are (a) time-dependent (i.e. they may occur differently at specific times), (b) location-dependent (i.e. they may occur differently at specific geographical locations) and (c) dynamic (i.e. their occurrence may evolve across time and space). Thus, appropriate quantitative methods must be selected to account for the temporal, spatial, and evolving nature of these processes.

First, we must consider the specific temporal properties of hashtag activism, physical protest, and legislative action, and quantify how these may or may not relate to each other through time. In order to address sub-question 1 (i.e. What is the relationship between location-based hashtag activism and physical protest activity?) and sub-question 2 (i.e. What is the relationship between location-based hashtag activism and legislative action?) one needs to understand how instances of hashtag activism over time relate to past, present, and future rates of physical protest activity and legislative action. Additionally, we need to understand if temporal information relating to hashtag activism can provide useful insight regarding future instances of physical protest and legislative action. Thus, the method chosen to address these analytical problems is *Time-Series Analysis*, to extract temporal patterns, combined with *Machine Learning* techniques for *Time-Series Forecasting*.

Second, we must consider the geo-spatial characteristics of hashtag activism, physical protest, and legislative action, and quantify how these may or may not relate to each other through space. Each of these processes, and especially hashtag activism and physical protest, is not only dynamic through time: rather as time goes by it is likely to also evolve across space. For example, a series of physical protests may take place across New York City, NY, today and another series of physical protests may take place in Pittsburgh, PA, tomorrow. Similarly, thousands of social media users may be engaging in hashtag activism from different geographical locations today, and they may continue to do so from new geo-graphical locations tomorrow. Thus, in order to address the sub-questions of this research one must combine *Time-Series Analysis* with traditional *Geo-spatial Analysis* in order to accurately study such interactivity that exists across both spatial and temporal dimensions.

Third, in order to quantify the impacts of hashtag activism on physical protest activity and legislative action, we must establish statistical relationships between these three processes and measure how these change over time and across space. In order to address this analytical requirement we incorporate *Regression Analysis and Modelling* (i.e. power regression) within the temporal and spatial analyses aforementioned.

Finally, in order to successfully implement each of these analyses, appropriate data describing instances of hashtag activism, physical protest, and legislative action during the case study of #JusticeForGeorgeFloyd must be collected. This data must also satisfy the problem-specific dimensions of this research: it must carry detailed spatial, temporal, and contextual information relating to instances of hashtag activism, physical protest, and legislative action.

Thus, this chapter describes the data and methods utilized to leverage the #JusticeForGeorgeFloyd case study and approach each sub-question of this research. In the first section, the data required to conduct this research is introduced. For each data source the collection process is highlighted, the data structure is illustrated, and the pre-processing steps used to prepare it for analysis are described. In the second section of this chapter the computational procedures and modelling techniques chosen for this research are introduced. Accordingly, specific methods for time-series modelling, regression analysis, machine learning, and geo-spatial analysis are described and their implementation in relation to the research problem is clarified. Figure ?? provides a general overview of the experimental design proposed in for this research.

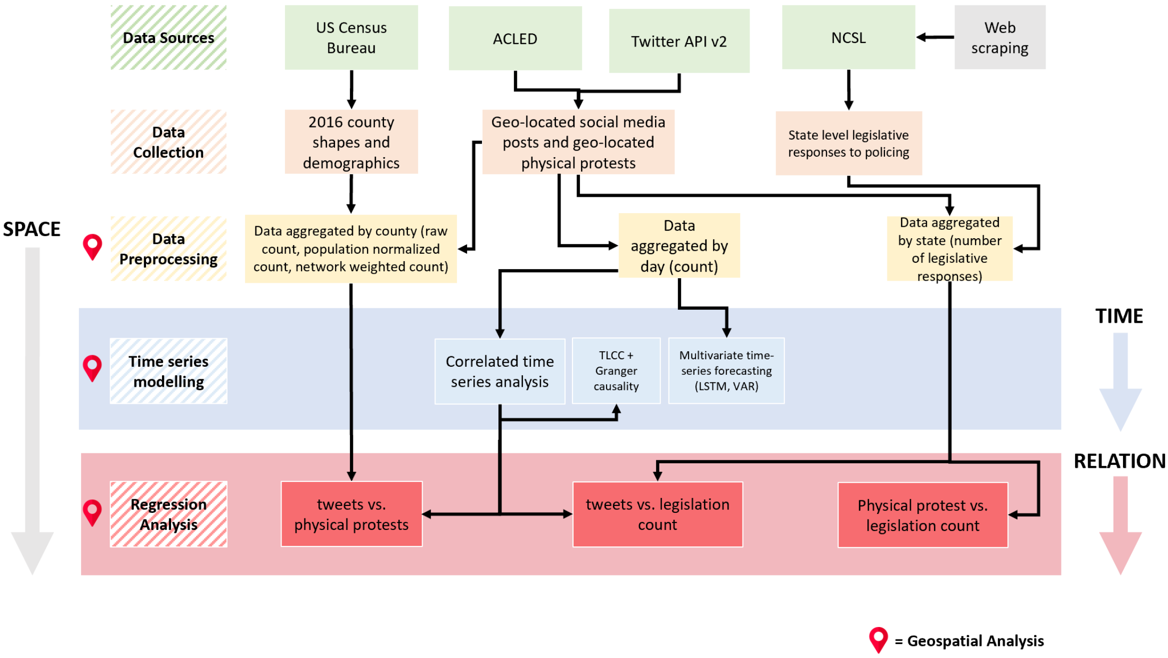

Figure 3.1: Experimental design chosen to conduct this research. This diagram is meant to be read from top to bottom, where each arrow represents a flow between an output and input. The top section (green) presents the data sources chosen for this research. The orange section presents the specific outputs resulting from the data collection phase. The yellow section describes the specific outputs resulting from the data pre-processing phase. The blue section presents the chosen methods for analysis related to time-series analysis. The red section presents the chosen methods for analysis related to regression analysis. Finally, the thick arrows on the right and left of the diagram highlight which research phases address which key problem-specific dimension.

3.1. DATA

3.1.1. DATA DESCRIPTION AND COLLECTION PROCESS

GEO-LOCATED SOCIAL MEDIA POSTS DATA

As argued in section 2.4 of chapter 2, Twitter was the social media platform of choice for hashtag activists during the #JusticeForGeorgeFloyd social movement. Thus, the primary data-set used for this research consists of a collection of geo-located social media posts (i.e. tweets) generated over the duration of the #JusticeForGeorgeFloyd social movement across the United States. This data was collected through the Twitter API v2 via the "Academic Research product track", an exclusive license that provides academic researchers with full access to the Twitter historical archive since March 26, 2006 (Twitter, 2021). In order to retrieve this data the Python programming language was used to connect to the `https://api.twitter.com/2/tweets/search/all` endpoint. In order to build a successful query the following constraints were taken into consideration:

1. Every tweet returned by the API must be related to the #JusticeForGeorgeFloyd social movement.
2. Every tweet returned by the API must originate from a location within the United States.

3. Every tweet returned by the API must have been created between May 20th 2020 and November 10th 2020.
4. Every tweet returned by the API must carry geo-information.

Additionally, in order to satisfy the data requirements of this research, fields pertaining to temporal, location, network, and textual features are specified.

The resulting query is presented as such:

```
query_params = {'query': f'(("george floyd" OR "elijah mcclain" OR "breonna taylor" OR (police brutality) OR blacklivesmatter OR (police violence) OR blm OR (racism police) OR (blacklives) OR (police accountability) OR "police reform" OR (defund police) OR (abolish police) OR "abolish the police" OR "police abolition" OR (police racist) OR (brutal police) OR (black lives) OR (police kill black) OR (black killed police) OR ("no justice" "no peace") OR "murdered by police" OR (police profiling) OR (police discrimination)) OR (#policereform OR #policeaccountability OR #defundthepolice OR #abolishthepolice OR #abolishpolice OR #defundpolice OR #criminaljusticereform OR #notoracism OR #blacklivesmatter OR #blm OR #georgefloyd OR #justiceforgeorgefloyd OR #icantbreathe OR #nojusticenopeace OR #justiceforfloyd OR #justiceforgeorge OR #sayhername OR #stopracismnow OR #blacklivesstillmatter OR #taketheknee OR #breonnataylor OR #justiceforbreonnataylor OR #blackwomenmatter)) has:geo point_radius:[{geo_query}]',
'tweet.fields': 'author_id, conversation_id, created_at, geo, entities, id, in_reply_to_user_id, referenced_tweets, public_metrics, text',
'place.fields': 'contained_within, country, country_code, full_name, geo, id, name, place_type',
'max_results': '500',
'expansions': 'geo.place_id, author_id, entities.mentions.username, in_reply_to_user_id, referenced_tweets.id, referenced_tweets.id.author_id',
'user.fields': 'created_at,description, id, location, name, protected, public_metrics, url, username, verified',
'start_time': '2020-05-20T00:00:00.00Z',
'end_time': '2020-11-10T00:00:00.00Z'}
```

Finally, because the Twitter API v2 allows for a maximum search radius of 25 miles, a grid consisting of equally sized 1770 km^2 hexagons covering the United States' landmass was created using the h3 Python library (Uber, 2021). Through an iterative process, and according to the above query, tweets within a 25 mile radius of the centroid of each hexagon were collected. As a result, the API returned a collection of 1,066,440 tweet objects each represented as a .json file structure as described in fig. 8.1 in the appendix. Each such tweet object contains temporal features, location features, network features and textual features. Table 3.1 describes the attributes returned for each tweet file.

	Description
id	The unique ID that serves to identify each tweet object.
created_at	The timestamp (year-month-day hour:min:second.timezone) at which the tweet was created.
public_metrics	The number of retweets, replies, likes, and quotes that the tweet received.
entities	Any special entities extracted from the tweet text (urls, user mentions, hashtags, place mentions).
geo	The geographical location where the tweet was created. If the user had enabled location tracking at the time the tweet was created, the latitude and longitude coordinates of the tweet are found within the "coordinates" sub-attribute. Otherwise if the user tweeted from a known place (e.g. restaurant, shop, neighborhood, city), this information is found within the "place_id" sub-attribute.
author_id	The unique ID that serves to identify the creator of the tweet.
text	The text that the user wrote when creating the tweet.
conversation_id	If a tweet is part of a conversation (e.g. a reply in a thread), this attribute is the unique ID that serves to identify the conversation of which a tweet may be part of.

Table 3.1: Attributes returned by the Twitter API v2 with the specified search query, for each tweet object

GEO-LOCATED PHYSICAL PROTEST DATA

Data on physical protest activity was retrieved from The Armed Conflict Location & Event Data Project (ACLED) (Raleigh *et al.*, 2010). ACLED is an event based data project that "codes the actions of rebels, governments, and militias within unstable states, specifying the exact location and date of battle events, transfers of military control, headquarter establishment, civilian violence, and rioting". In the case of #JusticeForGeorgeFloyd, ACLED is particularly relevant as it includes the most complete and reliable information regarding every Black Lives Matter physical protest that took place in the United States during 2020. ACLED data was acquired through the ACLED website as a .csv tabular file, of which relevant attributes pertaining to the timestamp, type, organizing actors, and location of each protest are extracted (table 3.2).

	Description
event_id_cnty	An individual identifier by number and country acronym.
event_date	The date the event occurred in the format: yyyy-mm-dd.
event_type	The type of conflict event.
actor1	The named actor involved in the event.
assoc_actor_1	The named actor allied with or identifying ACTOR1.
actor2	The named actor involved in the event.
assoc_actor_2	The named actor allied with or identifying ACTOR2.
interaction	A numeric code indicating the interaction between types of ACTOR1 and ACTOR2.
location	The location in which the event took place.
latitude	The latitude of the location.
longitude	The longitude of the location.
fatalities	The number of reported fatalities which occured during the event.

Table 3.2: Attributes extracted from ACLED data (Raleigh *et al.*, 2010)

STATE LEVEL LEGISLATIVE ACTION DATA

To date, the most reliable source of data on official U.S. legislation related to policing is the "Legislative Responses for Policing - State Bill Tracking Database" from the National Conference of State Legislatures (NCSL) (NCSL, 2020). This database has been actively tracking the status of every piece of law enforcement legislation for the United States' 50 states and the District of Columbia (NCSL, 2020). Data on legislative responses to policing is available for the year of 2020 and 2021. Although this data is publicly available, the NCSL does not facilitate public download in a user friendly format. Thus, in order to retrieve this data the Python library BeautifulSoup4 (PyPi, 2021) was used to scrape status information about each piece of legislation introduced since May 2020. In total, a data-set consisting of 2958 state policies, of which 743 were introduced in 2020 and 2215 were introduced in 2021, was created. Table 3.3 describes the attributes extracted for each piece of legislation.

	Description
policy_id	A unique identifier for legislation.
state	The state where the legislation is to be implemented.
year	The year in which the legislation was introduced.
date	The legal status of the legislation (i.e. pending, failed, enacted, vetoed, to mayor, adopted).
status	The named actor allied with or identifying ACTOR1.
date_of_last_action	The most recent date where action concerning the legislation took place.
author	The named author of the legislation.
topics	The topics related to the legislation.
summary	A brief description of the legislation.
history	A history of the different legal stages the legislation went through.

Table 3.3: Attributes extracted from NCSL data

2016 US-County and US-State Shapes and Demographics

In order to study any kind of geo-spatial relationships between social media activity, physical protest activity and legislative action, data must be aggregated across common spatial scales. For this research, the U.S. county and state spatial scales are chosen and the associated geo-spatial data is downloaded from the U.S. Census public API (Bureau, 2021). Additionally, basic demographic indicators at the U.S. county and State level are retrieved from the U.S. Census public API.

3.1.2. Data Pre-processing

Geo-located Social Media Posts Data

The .json format raw tweet objects returned by the Twitter API v2 are a complex data structure and carry more information than needed for this research. Thus, the first step in pre-processing the tweet data is to extract features of interest from each tweet object and store them in a tabular format that is appropriate for data analysis. This is performed by programmatically iterating through each .json file and extracting the following information:

1. Basic tweet information (tweet_id, conversation_id, created_at, text, hashtags).
2. Location information (coordinates, place_full_name, place_country_code, place_country, place_type, place_bbox, place_from_text).
3. Public metrics and engagement information (retweet_count, reply_count, like_count, quote_count).
4. Information regarding the user who created the tweet (user_id, username, user_full_name,user_following, user_followers, user_location, user_created_at, user_bio).
5. Information regarding any additional user that a tweet may be replying to or engaging with (in_reply_to_user_id, in_reply_to_username, in_reply_to_user_full_name, in_reply_to_user_following, in_reply_to_user_followers, in_reply_to_user_location, in_reply_to_user_created_at, in_reply_to_user_bio).
6. Information any additional tweet that a tweet may be referencing (referenced_tweet_id, referenced_tweet_type, referenced_tweet_text, referenced_tweet_created_at, referenced_tweet_author_id, referenced_tweet_conversation_id).

Once this iterative process has been completed, features from all tweets are combined together in one data-frame with the use of the Pandas Python package (Pydata, 2021). The result is a table of 1,066,440 tweets, where 267,033 tweets have location information at below-city level (i.e. exact location coordinates, coordinates of a point of interest such as a restaurant, an administrative bounding box such as for a neighborhood or a park), and 799,407 have location information at the city level. In order to extract a set of latitude-longitude coordinates for each tweet, the geographic centroid is calculated for each tweet that has a bounding box as location information. The geographical scope of the resulting data-set is illustrated in fig. 3.2.

In order to allow comparison between physical protest counts, tweet counts and legislation, tweet data must be further manipulated and aggregated at area and time units common to all three data-sets. Thus, a

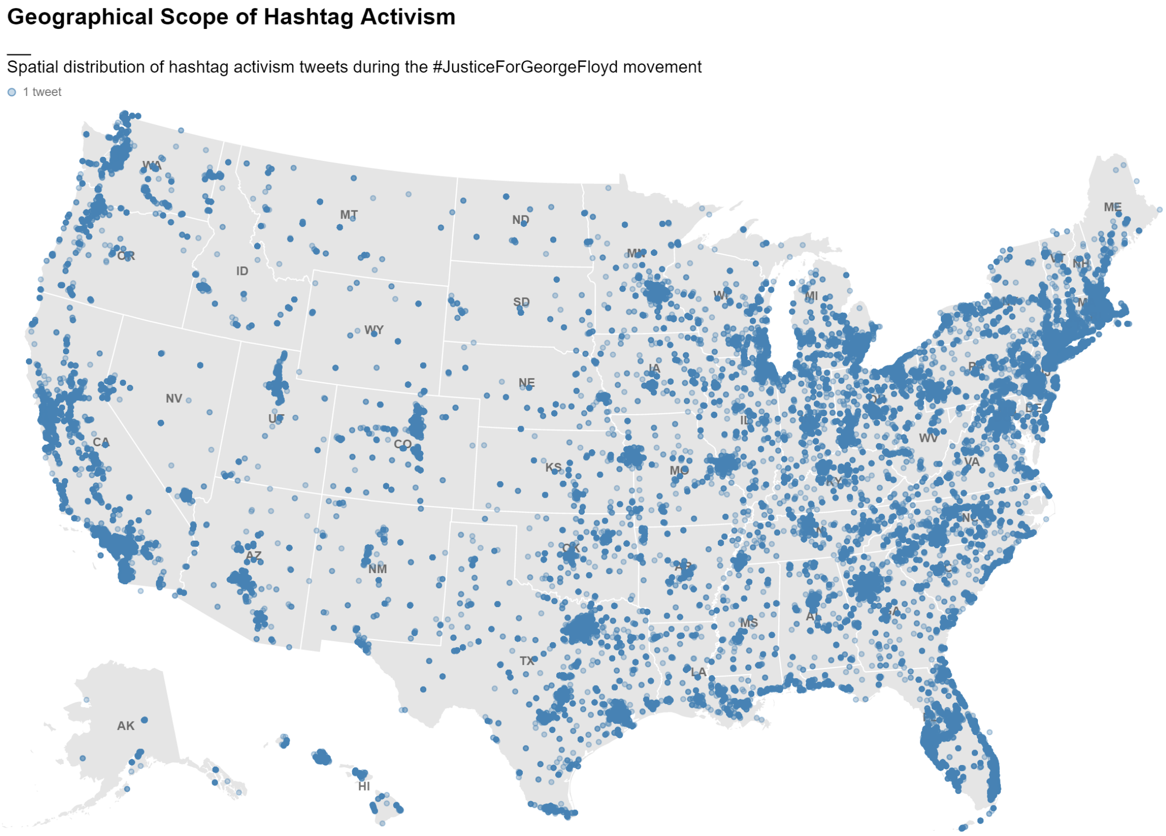

Figure 3.2: Geographical scope of hashtag activism tweets during #JusticeForGeorgeFloyd from tweet data-set. Each dot represents one tweet.

spatial join is first performed between tweet data and both 2016 US-county shapes and 2016 US-state shapes. Second, tweet data is aggregated spatially and temporally (by day). The result is five new data-sets:

1. Total number of daily tweets, irrespective of geographical location (168 records).
2. Total number of tweets by county, irrespective of time (3111 records).
3. Total number of tweets by State, irrespective of time (51 records).
4. Total number of daily tweets by county (522648 records).
5. Total number of daily tweets by state (3121 records).

GEO-LOCATED PHYSICAL PROTEST DATA

The ACLED data-set chosen for this research contains information about six different types of crisis events - Protests, Strategic developments, Violence against civilians, Riots, Battles, and Explosions/Remote violence - from January 1st 2020 to February 12th 2021 in the United States. Additionally, the ACLED data-set has recorded 403 actors directly or indirectly involved in one type of crisis event during this time-frame in the United States alone. Thus, the first step in pre-processing this data for analysis is to filter it to contain only crisis events that are relevant to the #JusticeForGeorgeFloyd case study. Thus, the following requirements are identified:

1. Each crisis event must either be a physical protest or a riot. For this the data-set is filtered according to the "EVENT_TYPE" field to equal "Protests" or "Riots".
2. Each crisis event must have taken place between May 20th 2020 and November 10th 2020. For this the data-set is filtered according to the "EVENT_DATE" field.
3. The actors (i.e. social group, political interest group, or social movement organization) directly or indirectly involved in each crisis event must either be in support of the #BlackLivesMatter movement or have its political interests closely aligned with those of #BlackLivesMatter. For this the data-set is filtered according to the "ASSOC_ACTOR_1" and "ASSOC_ACTOR_2" fields to match any of the follow-

ing actors: "BLM: Black Lives Matter", "NAN: National Action Network", "African American Group", "NAACP", "BSU: Black Student Union", "NABPP", "SURJ", "APSP: African People's Socialist Party", "Black Fist Coalition", "Black, Young, And Educated", "Black Panthers", "CCBP: Chesapeake Coalition of Black Pastors", "NABPP: New Afrikan Black Panther Party", "Black Leaders Movement LNK", "Black Unity", "RBPP: Revolutionary Black Panther Party", "BU: Black Unity", "Black Horse Militia", "APSP: African People's Socialist Party", "RAM: Revolutionary Abolitionist Movement", "NFAC: Not Fucking Around Coalition".

As a result, the ACLED data-set was reduced to 10278 protest events of which the geographical scope is illustrated in fig. 3.3.

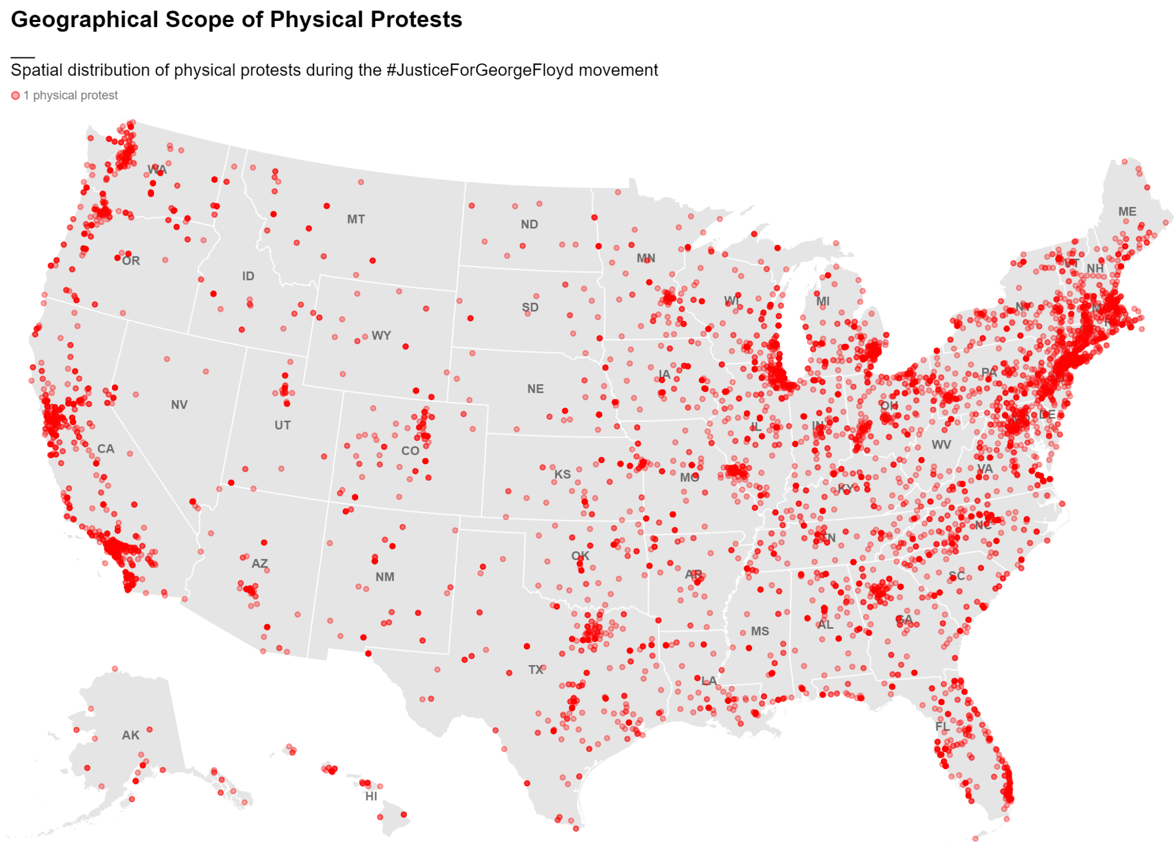

Figure 3.3: Geographical scope of physical protest events during #JusticeForGeorgeFloyd from ACLED data-set. Each dot represents one physical protest.

In order to allow comparison between physical protest counts, tweet counts and legislation, ACLED data must be further manipulated and aggregated at area and time units common to all three data-sets. Thus, a spatial join is first performed between ACLED data and both 2016 US-county shapes and 2016 US-state shapes. Second, ACLED data is aggregated spatially and temporally (by day). The result is five new data-sets:

1. Total number of daily physical protests, irrespective of geographical location (168 records).
2. Total number of physical protests by county, irrespective of time (3111 records).
3. Total number of physical protests by State, irrespective of time (51 records).
4. Total number of daily physical protests by county (522648 records).
5. Total number of daily physical protests by state (3121 records).

STATE LEVEL LEGISLATIVE ACTION DATA

The NCLS data on legislative responses to policing contains information about state legislature relating to fifteen different policing topics: 'Data and Transparency', 'Technology', 'Other Issues', 'Employment and Labor', 'Oversight', 'Decertification', 'Executive Orders', 'Policing Alternatives and Collaboration', 'Associated Bills: NJ

S 2801 - Identical', 'Officer Safety and Wellbeing', 'Use of Force', 'Investigations and Discipline', 'Certification', 'Training', and 'Standards'. Additionally, the legal status of each piece of legislation is available as one of the following categories: "Adopted", "Enacted", "Failed", "No Status", "Pending", "To Congress", "To Governor", "To Mayor", and "Vetoed". The geographical scope of this data-set is illustrated in fig. 3.2.

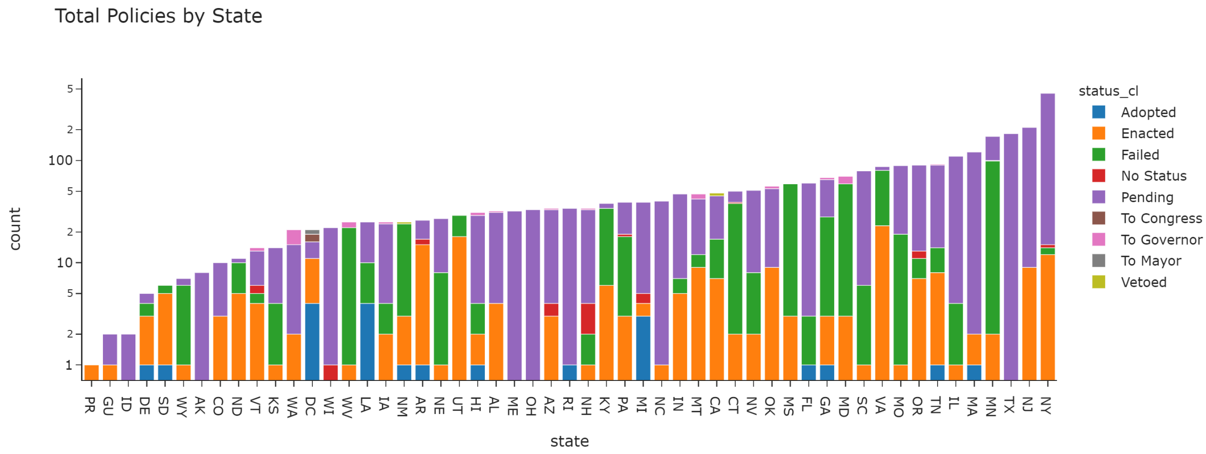

Figure 3.4: A histogram of state-level counts of legislative responses to policing, grouped by status categories (June 2020-June 2021)

In order to make this data comparable to physical protest data and tweet data, it is aggregated spatially and temporally (by day). The result is three new data-sets:

1. Total number of daily legislative responses, irrespective of geographical location (219 records).
2. Total number of legislative responses by State, irrespective of time (51 records).
3. Total number of daily legislative responses by state (1068 records).

COMBINING AND NORMALIZING DATA-SETS

Following the data pre-processing steps outlined above, data sources are combined together through tabular merging according to common keys. For example, the total count of physical protests by state data is merged with the total count of tweets by state data and the total count of legislative responses to policing data according to the "state" attribute. County level total counts data can be merged together in a similar fashion. Data describing the temporal dimension only (e.g. total daily count of tweets and total daily count of protests) can be merged according to "date" attribute. Finally, data describing both temporal and spatial dimensions (e.g. total daily count of tweets by state and total daily count of physical protests by state) can be merged according to both the "state" and "date" attributes. Finally, each data-set describing absolute counts irrespective of the temporal dimension (e.g. total count of tweets by county) is normalized by population count. This is done by dividing the absolute count variable of each specific area unit by the population of that area unit. Additional pre-processing steps specific to each analysis or modelling technique used in this research is described in the following sections of this chapter.

3.2. TIME: TIME-SERIES ANALYSIS AND FORECASTING

In nature and in society there are several processes, such as instances of physical protest or social media posts, that can be characterized by sequential events or occurrences. Such processes, when translated into data-set consisting of sequential observations, are defined as *time-series* (Wei, 2006). In this context, **time-series analysis** is the use of statistical tools and algorithms to extract temporal patterns from time-series, and **Time-series Forecasting** is the use of these temporal patterns to predict future values from past values (Wei, 2006). Time-series analysis and forecasting finds applications in a wide range of scientific fields, such as Economics, Sociology and Meteorology. In this research, time-series analysis is used to extract the temporal patterns and relationships that exist between hashtag activism, physical protest, and legislative action processes, and time-series forecasting is used to predict instances of physical protest and legislative action based on past instances of hashtag activism.

3.2.1. CORRELATED TIME-SERIES ANALYSIS

When two time-series are contextually related, they may or may not exhibit similar temporal behaviour. The study of this temporal relationship between two or more time-series is called **correlated time-series analysis**. Essentially, specific statistical tests can be executed to quantify local and global relationships between time-series. Correlated time-series analysis methods can generate useful insight regarding the flow of information between different time-series (Wei, 2006).

ROLLING WINDOW CROSS CORRELATION

3

A common method for performing correlated time-series analysis is a procedure called **rolling window correlation** (RWC), where the correlation between two time-series is repeatedly computed over a specific time window throughout the length of the time-series. This procedure involves (a) choosing two time-series that cover the same temporal extent, (b) choosing a temporal window of a specific length, (c) computing the Pearson correlation p value between the two time-series within that window of time and (d) shifting (or rolling) the window by some offset and repeating step (c). This process is executed for the entire length of the two time-series (Shakil *et al.*, 2016). The Pearson correlation p is calculated as such (Benesty *et al.*, 2009):

$$p = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{3.1}$$

While RWC has proven useful for a variety of applications ranging from neuroscience (Shakil *et al.*, 2016), finance (Andersson *et al.*, 2008), urban planning (Zhang and Fan, 2019) to social studies (Vomfell *et al.*, 2018), it has some limitations. Mainly, results of RWC are highly dependent on the size of the window that is used and, because this parameter is somewhat arbitrary, they may not be entirely objective . This limitation can be controlled for by performing RWC with different window sizes and taking note of how results change each time. This phenomenon is illustrated in fig. 3.5, where RWC analysis is computed on tweet count and protest count data during the #JusticeForGeorgeFloyd social movement for the example of New York State (NY) using a seven day (top) and three day (bottom) window. In this research, RWC is utilized to study the evolving temporal relationship between tweet count and protest count throughout the #JusticeForGeorgeFloyd social movement.

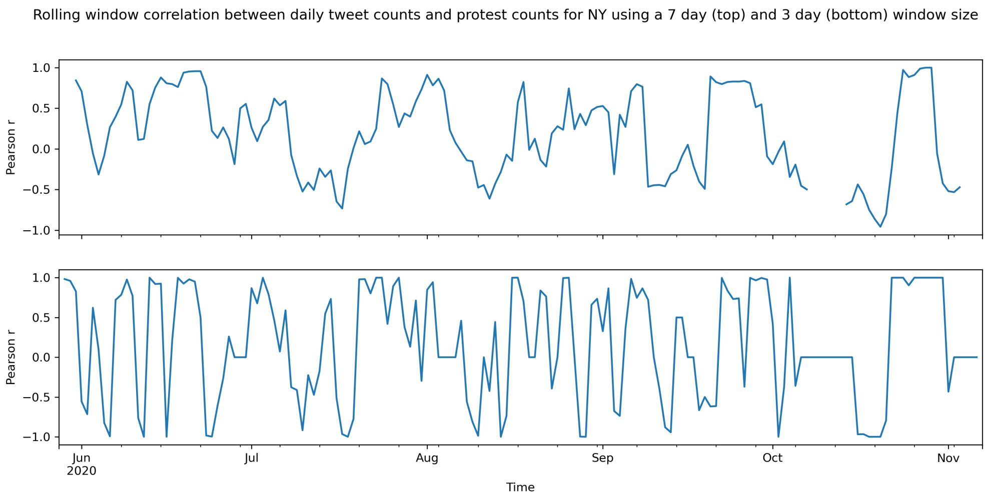

Figure 3.5: Results of Rolling Window Correlation analysis for tweet count and protest count time-series during the #JusticeForGeorge-Floyd social movement for NY using a seven day (top) and three day (bottom) window

TIME LAGGED CROSS CORRELATION ANALYSIS

While RWC analysis can generate information regarding the relationship between two time-series, it cannot generate insight regarding the flow of information between two time-series. That is, it cannot describe the directionality between two time-series, such as the existence of a leader-follower relationship where one variable expresses occurrence before the other variable. One method that is effective for this type of inference is **time lagged cross correlation analysis** (TLCC) (Shen, 2015). Essentially, TLCC studies the directionality between two time-series, x_t an y_t, by identifying lags of x that may be the best predictors of y_t. This method involves (a) choosing two time-series that cover the same temporal extent and (b) fitting a regression model $n * 2$ times using the y_t data and the x_{t-n} and x_{t+n} data, where $|n|$ is the maximum number of time lags. In this research TLCC is used to test the hypothesis that one of hashtag activism or physical protest happens before the other.

As an example, fig. 3.6 illustrates the results of TLCC using tweet count and protest count data during the #JusticeForGeorgeFloyd social movement for NY. In fig. 3.6, it is shown that, for NY, the optimal Pearson correlation p value between tweet count and protest count time-series is obtained when a time lag of one day is introduced in the tweet count data. In other words, during the case study of #JusticeForGeorgeFloyd in NY, the best predictor of physical protest activity at time t (in days) is the number of social movement related tweets posted at time $t1$ (i.e. the day before).

While TLCC is useful for identifying the directionality between two time-series, it is important to note that this method is not a measure of causality. While it can identify time lags at which best correlations are obtained, it does not provide statistically significant information for causal inference.

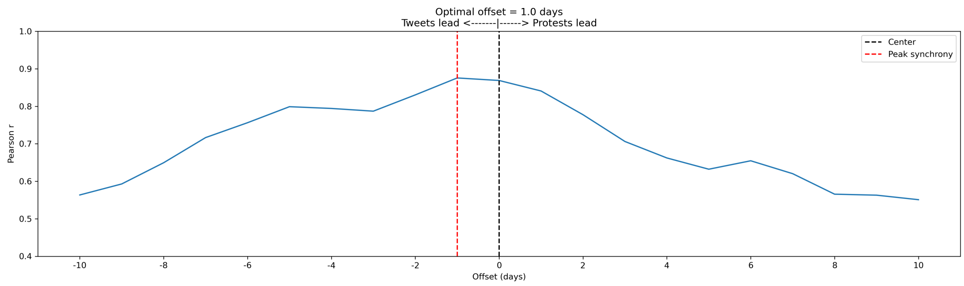

Figure 3.6: Example of TLCC analysis results for New York State (NY), where the optimal Pearson correlation p value is found with tweet count leading by a time lag of one day.

3.2.2. GRANGER CAUSALITY ANALYSIS

A more comprehensive method for analyzing the relationship between two time-series is **Granger Causality Analysis**. Granger causality is a statistical hypothesis test developed to study the exchange of information between two or more time-series (Granger, 1969; Zhao *et al.*, 2020). In the field of social movements research, Granger causality analysis has proven useful in a number of instances, such as for the study of causal relationships between social network nodes (Ver Steeg and Galstyan, 2012) and the study of causal relationships between protests and repression (Carey, 2006). More relevantly, and as previously discussed in 2, Bastos *et al.* (2015) and Brantly (2019) both implemented Granger causality to argue for the existence of causal relationships between social media activity and physical protest activity.

Essentially, Granger causality determines whether prior values from one time-series can be used to predict current and future values of another time-series. Thus, while Granger causality does not always imply *True* Causality, it is said that a time-series $x2$ *Granger-causes* another time-series $x1$ if *lagged* values of $x2$ provide statistically significant information about future values of $x1$ (Granger, 1969). As such, the Null hypothesis that $x2$ does not Granger-cause $x1$ is tested by first identifying the correct lagged values of $x1$ to use in a uni-variate auto-regression of $x1$:

$$x1_t = a_0 + a_1 x1_{t-1} + a_2 x1_{t-2} + \cdots + a_m x1_{t-m} + \text{ error }_t \tag{3.2}$$

Following this step, the auto-regression is modified to include lagged values of $x2$:

$$x1_t = a_0 + a_1 x1_{t-1} + a_2 x1_{t-2} + \cdots + a_m x1_{t-m} + b_1 x2_{t-1} + \cdots + b_q x2_{t-q} + \text{ error }_t \tag{3.3}$$

Next, the lagged values of $x2$ that (a) collectively add explanatory power to the auto-regression according to an F-test and (b) are found to be individually significant from the value of their t-statistic are retained in the auto-regression. In the above equation, q is the smallest lag and $t-q$ is the largest lag for which a lagged value of $x2$ is significant. Finally, we accept the Null hypothesis that $x2$ does not Granger-cause $x1$ if no lagged values of $x2$ are retained in the auto-regression (Granger, 1969).

In this research, Granger causality analysis is implemented using the *grangercausalitytests* function from the *statsmodels* Python library (Seabold and Perktold, 2010). This function performs four F-tests for granger non-causality of two time-series, where the Null hypothesis is that the second time-series $x2$ does not Granger cause the first time-series $x1$. In this implementation "Granger causality means that past values of $x2$ have a statistically significant effect on the current value of $x1$, taking past values of $x1$ into account as regressors. We reject the null hypothesis that $x2$ does not Granger cause $x1$ if the p-values are below a desired size of the test" (McKinney *et al.*, 2011). With national, state-level, and county-level data, Granger causality is implemented to study whether data on tweet counts can provide information that is statistically significant for the forecast of future physical protest counts.

3.2.3. MULTI-VARIATE TIME-SERIES FORECASTING

Time-series forecasting is the use of temporal patterns extracted during time-series analysis to predict future values from past values (Wei, 2006). When two or more time-series are found to be correlated, **multi-variate time-series forecasting** methods can be used to predict future instances of a time-series of interest based on its past values and the past values of other time-series it is correlated with (Wei, 2006). In this research, multi-variate time-series forecasting is used to predict future instances of physical protest and legislative action based on past instances of both physical protest and hashtag activism. As such, two predictive models are developed: a simple Long Short Term Memory (LSTM) neural network and a Long Short Term Memory neural network combined with a Vector Autoregressive (VAR) model (LSTM VAR). In the following sections, the technical characteristics of each model are described, the modelling choices are motivated, and the overall machine-learning modelling workflow is described.

VECTOR AUTOREGRESSIVE (VAR) MODEL

The Vector Autoregressive (VAR) Model is a statistical model designed to learn the relationship that exists between different temporal signals, and use this information for the purpose of time-series forecasting (Zivot and Wang, 2006). Across numerous disciplines, VAR is considered to be one of the most effective and intuitive methods for the analysis and forecasting of multi-variate time-series (Zivot and Wang, 2006). In particular, VAR models have demonstrated success for modelling the dynamic nature of time-series and have found many applications in the realm of policy analysis. These characteristics make VAR very attractive and suitable for modelling the multi-variate relationship between hashtag activism and physical protest, and for forecasting physical protest counts (Zivot and Wang, 2006). If $\mathbf{Y}_t = (y_{1t}, y_{2t}, \ldots, y_{nt})'$ "denote an $n \times 1$ vector of time-series variables, the basic p-lag vector autoregressive (VAR(p)) model has the form:

$$\mathbf{Y}_t = \mathbf{c} + \mathbf{\Pi}_1 \mathbf{Y}_{t-1} + \mathbf{\Pi}_2 \mathbf{Y}_{t-2} + \cdots + \mathbf{\Pi}_p \mathbf{Y}_{t-p} + \varepsilon_t, t = 1, \ldots, T \tag{3.4}$$

"where Π_i are n coefficient matrices and ε_t is an $n \times 1$ unobservable zero mean white noise vector process (serially uncorrelated or independent) with time invariant covariance matrix" (Zivot and Wang, 2006, p. 386). Since in this research, we have two time-series (tweet count and physical protest count), the associated bivariate (VAR(2)) model has the form:

$$\begin{aligned} \begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} &= \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{pmatrix} \pi_{11}^1 & \pi_{12}^1 \\ \pi_{21}^1 & \pi_{22}^1 \end{pmatrix} \begin{pmatrix} y_{1t-1} \\ y_{2t-1} \end{pmatrix} \\ &+ \begin{pmatrix} \pi_{11}^2 & \pi_{12}^2 \\ \pi_{21}^2 & \pi_{22}^2 \end{pmatrix} \begin{pmatrix} y_{1t-2} \\ y_{2t-2} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \end{aligned} \tag{3.5}$$

or

$$\begin{aligned} y_{1t} &= c_1 + \pi_{11}^1 y_{1t-1} + \pi_{12}^1 y_{2t-1} + \pi_{11}^2 y_{1t-2} + \pi_{12}^2 y_{2t-2} + \varepsilon_{1t} \\ y_{2t} &= c_2 + \pi_{21}^1 y_{1t-1} + \pi_{22}^1 y_{2t-1} + \pi_{21}^2 y_{1t-1} + \pi_{22}^2 y_{2t-1} + \varepsilon_{2t} \end{aligned} \tag{3.6}$$

where y_{1t} is tweet count and y_{2t} is physical protest count, and where $\text{cov}(\varepsilon_{1t}, \varepsilon_{2s}) = \sigma_{12}$ for $t = s; 0$ otherwise (Zivot and Wang, 2006, p. 386).

LONG SHORT-TERM MEMORY (LSTM) ARTIFICIAL NEURAL NETWORKS

Artificial Neural Networks (ANN) are a revolutionary machine learning technique (also known as *deep learning*), consisting of a set of inter-connected nodes (i.e. mathematical formulas), which mimics the functioning of biological networks (e.g. the human brain) to recognize patterns within large amounts of data (McCulloch and Pitts, 1943; Yegnanarayana, 2009). Neural networks are unique because they constantly update their knowledge about their environment (i.e. the data) as they face new experiences, through a process called *backpropagation*, where the connection *weights* within the network are adjusted based on the *error* found during the learning process (McCulloch and Pitts, 1943; Yegnanarayana, 2009). A high level representation of a simple ANN structure is illustrated by fig. 3.7.

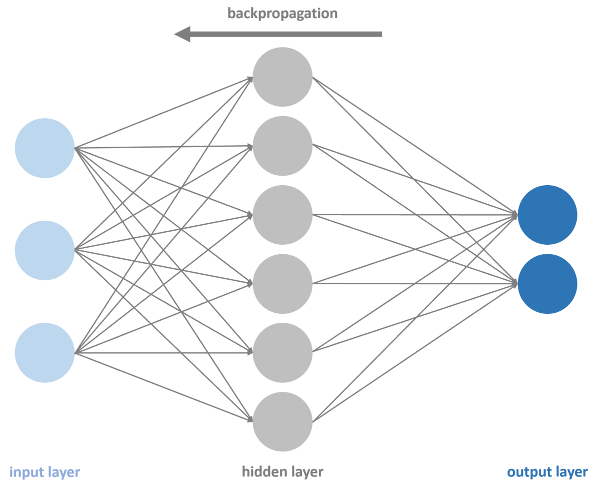

Figure 3.7: Example of a simple ANN structure

In recent years, through the increased availability of implementation methods for deep learning, **Long Short Term Memory (LSTM) neural networks** have become a popular and very effective method for the modelling of sequences, including multi-variate time-series forecasting (Sherstinsky, 2020). LSTM neural networks are a type of Recurrent Neural Networks (RNN) that were designed to overcome the *vanishing gradient* or *back-propagated error decay* problem associated with RNNs, where the backpropagated gradients become exponentially small and cause the learning process to stop, that arises during the learning of long sequences (Hochreiter and Schmidhuber, 1997). Essentially, LSTM neural networks use *memory blocks* to store long sequences of data and based on them predict the next values of a sequence (Hochreiter and Schmidhuber, 1997; Sherstinsky, 2020; Greff *et al.*, 2016). Fig. 3.8 illustrates the memory block of a simple RNN (SRN) (left) and compares it to the structure of a LSTM memory block (right).

LSTM neural networks have proven to be very effective for the forecasting of multi-variate time-series across several disciplines, such as traffic flows (Ma *et al.*, 2015), the concentration of air pollutants (X. Li *et al.*, 2017), stock prices (Bao *et al.*, 2017), and electric loads (Zheng *et al.*, 2017). In this research, LSTM neural networks are used in isolation and in combination with VAR for the forecasting of physical protest activity.

MACHINE-LEARNING MODELLING WORKFLOW

In this research, we develop two models for the forecast of next-day physical protest counts based on the previous seven days of physical protest counts and tweet counts from national level data. First, in order to prepare the data for modelling, national level data on daily counts of physical protests and tweets is up-sampled from daily intervals to hourly intervals using linear interpolation. This procedure allows to generate more data points (from 168 days to 4032 hours) for the models to work with. Next, each of the physical protest and tweet count time-series are split up into training (60%), validation (10%), and testing (30%) data. Finally, training, validation, and testing data-sets are each re-shaped into 7-day sequences of data-points. Fig. 3.9 illustrates the portion of data used for training (blue) and the portion of data used for testing (orange) for each of physical protest and tweet count data-sets.

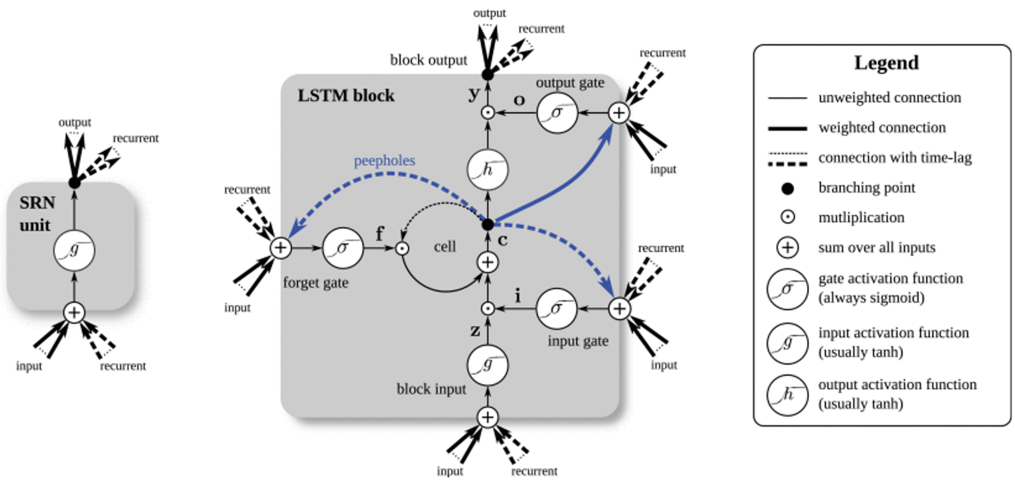

Figure 3.8: "Detailed schematic of the SRN unit (left) and an LSTM block (right) as used in the hidden layers of a recurrent neural network" (Greff *et al.*, 2016)

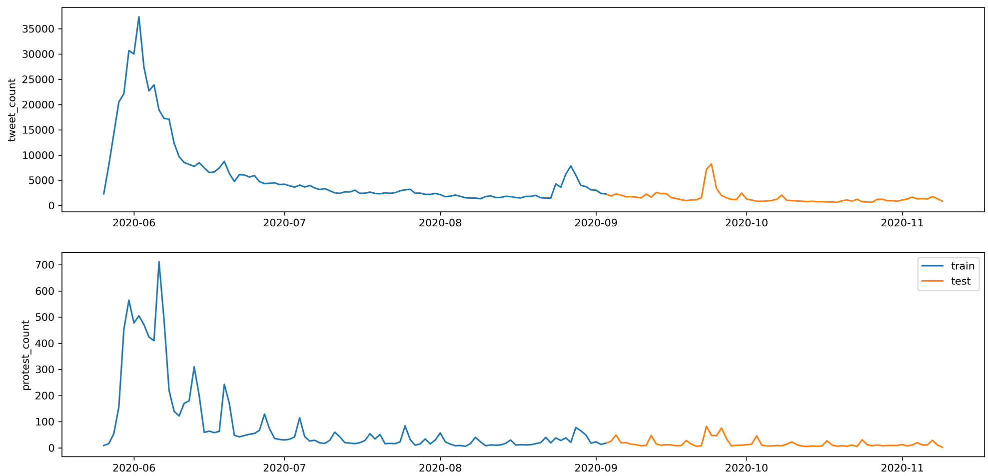

Figure 3.9: Training (blue) and testing (orange) data used for time-series forecasting with VAR

Using the training data, a VAR model is then fitted to the physical protest and tweet count time-series using the *var_model* implementation from the *statsmodels* Python library (Seabold and Perktold, 2010). The order selected for the VAR model is selected based on the lowest obtained Akaike Information Criterion (AIC) value, $\text{AIC}(n, L) = \log \hat{e}_L + 2L/n$ for $\hat{L}\left(1 \leq \hat{L} \leq L_{\max}^{(n)}\right)$, the state of the art metric for autoregressive model order selection (Ding *et al.*, 2017). This process is illustrated by fig. 3.10, where the AIC score at each model order (i.e. lag) is charted and the selected model order ($L = 146$) is highlighted in red. The resulting fitted values from the selected VAR model are then saved.

Next, two LSTM neural networks, LSTM and LSTM VAR, are trained and fine-tuned using the *LSTM* implementation from the *tensorflow* Python library and the *kerashypetune* Python library, respectively. Each model is a simple LSTM neural network which predicts the next day's value of physical protest counts based on the past seven days' values of physical protest counts and tweet counts. Each model consists of (a) an input layer of two nodes (physical protest count, tweet count), (b) three hidden layers of LSTM memory blocks using the *tanh* activation function, and (c) an output layer of two nodes (predicted physical protest count, predicted tweet count). The two models differ only in the data that they use for training: the first model, LSTM, uses the raw values of physical protest counts and tweet counts, whereas the second model, LSTM VAR, uses both

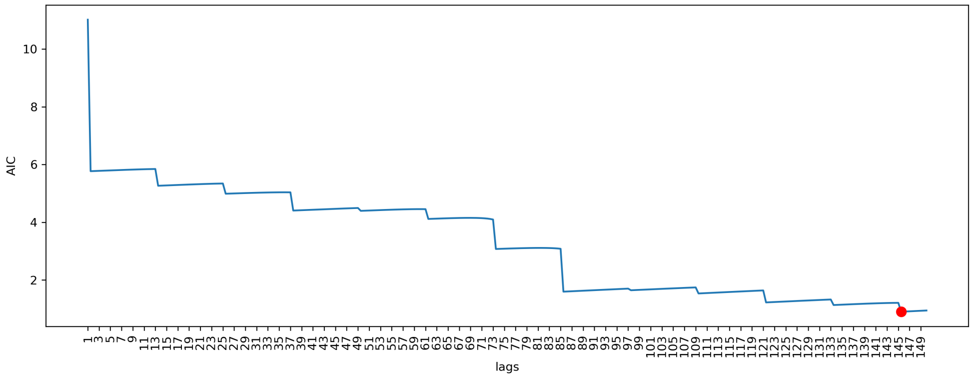

Figure 3.10: Optimal VAR model order chosen for time-series forecasting

the raw values of physical protest counts and tweet counts and the fitted values obtained from the previously trained VAR model. The objective behind the LSTM VAR model is to use what the VAR model has learned about the multivariate relationships between physical protests and tweet counts to improve the training and predictive power of the LSTM model. Thus, this model is trained and optimized by (a) feeding it seven-day sequences of both the multi-variate VAR fitted and raw time-series inputs, and adjusting link weights, (b) using the validation data to predict next-day physical protest counts and compute error values, (c) backpropagating error values and adjusting model link weights, and (d) repeating this process until the error values can no longer be reduced. In contrast, the LSTM model is trained and optimized through the same exact process except that it is only fed the raw time-series inputs. Finally, the optimized LSTM and LSTM VAR models are saved and utilized to predict national-level physical protest counts, using the national-level testing data, and state-level physical protest counts, using the full state-level data-set. Fig. 3.11 illustrates the high level machine-learning modelling workflow used to develop, train, fin-tune, and implement each of these models.

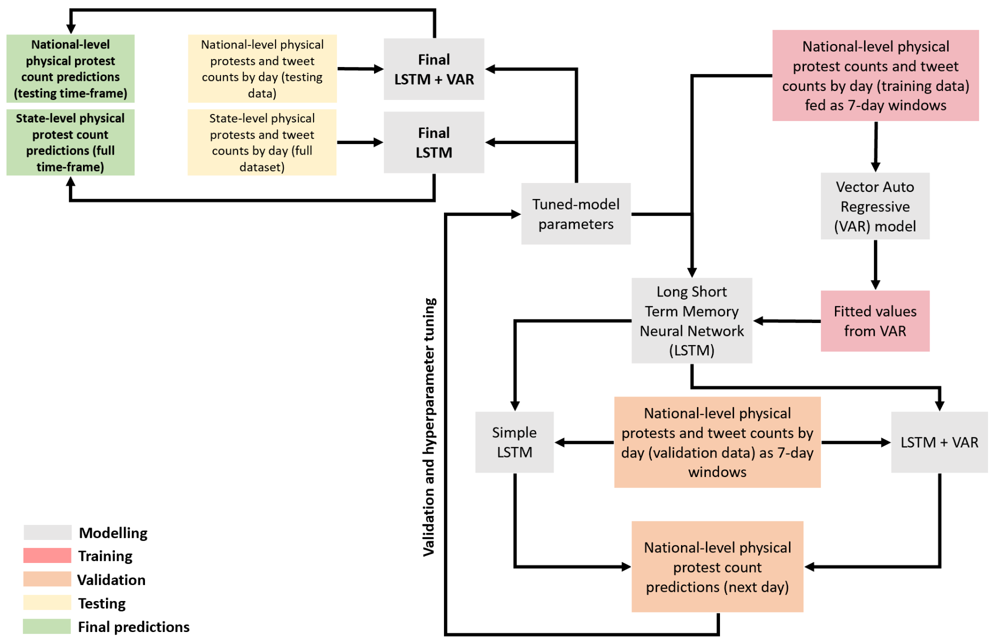

Figure 3.11: Machine learning modelling workflow for time-series forecasting

3.3. RELATIONSHIPS: REGRESSION ANALYSIS

Regression analysis describes a range of statistical methods for quantifying the relationships between a dependent variable and one or more independent variables (Draper and Smith, 1998). Regression is the most popularly used statistical procedure for the analysis of empirical problems in a wide range of fields such as biology and environmental sciences, social science, and economics (Fahrmeir *et al.*, 2007). In this research, regression is used to quantify the statistical relationship between hashtag activism and physical protest, the statistical relationship between hashtag activism and legislative responses to policing, and the statistical relationship between physical protest and legislative responses to policing. More specifically, due to the left skewed (log-normal) distribution of both tweet counts and physical protest counts (please see chapter 4) the power regression model is chosen for modelling such relationships. This regression model is used throughout all applications of regression modelling of this research. The power regression model is a non-linear regression model useful for modelling problems where the dependent variable is equal to the independent variable raised to a power, as such (Draper and Smith, 1998):

$$y = \alpha x^{\beta} \tag{3.7}$$

In this equation, y is the dependent variable, x is the independent variable, and α and β are the regression coefficients that describe the relationship between x and y. In this research, the performance of power regression models is evaluated by using the Pearson correlation score p (see eq. (3.1)) and the coefficient of determination R^2:

$$R^2 = 1 - \frac{RSS}{TSS} \tag{3.8}$$

where RSS is the Residual Sum of Squares, calculated as $RSS = \sum_{i=1}^{n} \left(y_i - f(x_i)\right)^2$, and TSS is the Total Sum of Squares, calculated as TSS $= \sum_{i=1}^{n} \left(y_i - \bar{y}\right)^2$, where y_i is a given sample value to be predicted, $f(x_i)$ is the predicted value of y_i, where $\bar{y}$ is the mean sample value, and n is the total number of observations in the sample. The Pearson correlation score p and the coefficient of determination R^2 are two of the most commonly used evaluation metrics in regression analysis (Fahrmeir *et al.*, 2007).

3.4. SPACE: GEOSPATIAL ANALYSIS

Geospatial analysis is a set of statistical procedures used for gathering, manipulating and analyzing data through their topological, geometric, or geographic attributes (De Smith *et al.*, 2007). Such data can come from a variety of data sources, such as satellite imagery, GPS coordinates and spatial networks, and geographic attributes can be defined explicitly (i.e. longitude-latitude coordinates) or implicitly (i.e. street addresses). In both cases, geospatial analysis is required to manipulate and combine spatial data together, and to establish statistical relationships between different spatial data. Geospatial analysis is used as a fundamental tool in a wide range of fields such as urban planning, environmental sciences, sociology and transportation(De Smith *et al.*, 2007). In this research, each data source contains explicit and implicit spatial information, such as geographical coordinates, the shapes of administrative regions, and the name of places and administrative regions. Thus, as described in section 3.1.1, several traditional geospatial analysis procedures are used to gather, process, manipulate and combine these different data sources analyze spatial relationships that exist between them.

4
Results

This chapter presents the results of the analysis conducted to investigate sub-question 1 (i.e. *What is the relationship between location-based hashtag activism and physical protest activity?*) and sub-question 2 (i.e. *What is the relationship between location-based hashtag activism and legislative action?*) of this research. In applying each step of the proposed experimental framework presented in chapter 3 to the case study of #JusticeForGeorgeFloyd, insight into the relationship between hashtag activism and two forms of political processes - physical protest and legislative action - are presented. First, the spatio-temporal relationship between hashtag activism and physical protest is explored. For this, time-series modelling, regression analysis and geo-spatial analysis steps are applied to explore how hashtag activism and physical protest are related in time and space. Then, a deep learning model is developed and tested for the forecasting of physical protest activity based on past instances of hashtag activism. This section concludes with the development of a novel index, aimed at providing an accurate and granular picture of where hashtag activism is most likely to translate into physical protest activity and what that may mean for local policy makers. Second, the spatio-temporal relationship between hashtag activism and legislative action is explored. For this, time-series modelling, regression analysis and geo-spatial analysis steps are applied to explore how hashtag activism and legislative action are related in time and space.

4.1. The spatio-temporal relationship between location-based hashtag activism and physical protest activity

As discussed in chapter 2 of this research, the study of the relationship between hashtag activism and physical protest has been inconclusive, with evidence both in support and against the existence of such a relationship. This, scholars conclude, is partly due to a lack of empirical research and the use of limited case studies. Additionally, most of the literature within this area of research has studied this relationship through a broad lens and does not consider the geo-spatial relationships that may exist between digital and physical protests. In an attempt to address this knowledge gap, this section presents the results of time-series modelling, regression analysis, and geo-spatial analysis of hashtag activism and physical protests using the case study of #JusticeForGeorgeFloyd. Thus, this section answers one of the main two research questions: ***What is the relationship between location-based hashtag activism and physical protest activity?***

Correlated Time Series Analysis

As described in chapter 3, point data describing instances of hashtag activism and physical protests were aggregated in time and space, so that the temporal relationship between hashtag activism and physical protest could be studied at different spatial scales. Figure 4.1 shows, in the right half of the chart, the temporal patterns of both hashtag activism (measured in number of tweets) and physical protest at the national level. On the left side of the chart the temporal relationship between daily number of physical protests and daily number of tweets is illustrated through a scatter-plot. From fig. 4.1 it can be observed that, at the national scale, hashtag activism bears a very strong positive relationship with physical protest activity. On days where people protested more on social media, people also protested more in the street, and vice versa. In the case of #JusticeForGeorgeFloyd, people may have used Twitter to express outrage about events of police brutality and to organize physical protests. During protests, more acts of police brutality were documented and shared on Twitter, which could have led to yet more protests and to more people showing up to protests, and so on. This trend can be modelled by a power regression function ($y_n = -0.004x_n^{1.116}$), represented by the red dashed line, where y_n is the number of physical protests that take place on a given day and x_n is the number of social movement tweets that are posted on a given day. This model is considerably accurate, with a Pearson correlation score p of 0.852 and an R^2 score of 0.726.

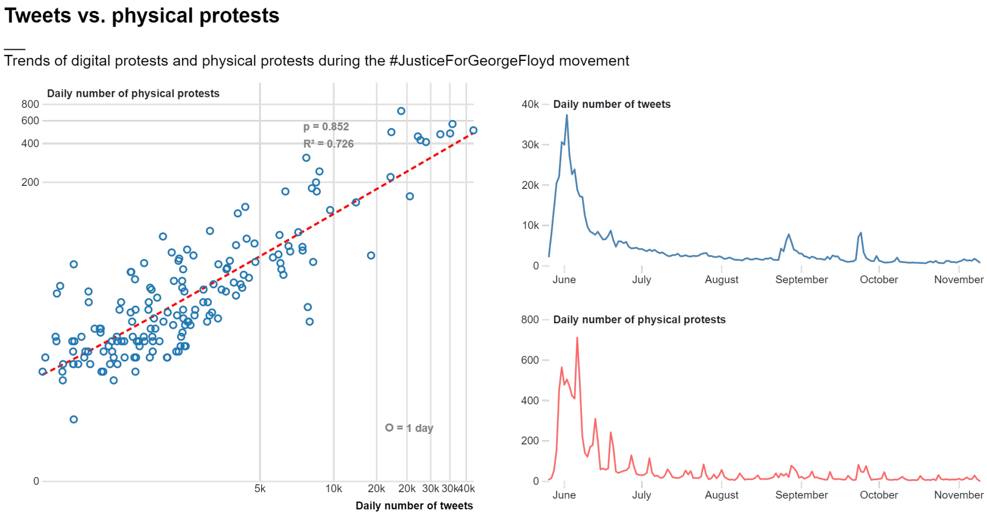

Figure 4.1: Temporal (right) and statistical (left) relationship between hashtag activism and physical protests during #JusticeForGeorge-Floyd, where each dot represents a day. The statistical relationship is modelled by the red dashed line as ($y_n = -0.004x_n^{1.116}$).

While this model provides a good description of the global relationship or synchrony between hashtag activism and physical protests, it offers little insight into local synchrony (how a relationshio changes over time) between the two variables. One method for exploring local synchrony is to compute the Pearson correlation between the two variables for a small subset of time (i.e. a rolling window) and repeat this for the entire duration of the time-series. Thus, fig. 4.2 illustrates the rolling window Pearson correlation between hashtag activism and physical protests with the use of a seven day window. The top subplot illustrates the two time series normalized (as $x' = \frac{x - \min(x)}{\max(x) - \min(x)}$) on a scale from 0 to 1 and the bottom subplot illustrates the seven day rolling window Pearson correlation between the two time series. The figure shows that the Pearson correlation is generally a high value, but oscillates to very low values for brief periods of time. This is due to the existence of some differences between the seasonality pattern of each time series, mainly the physical protest time series occasionally experiencing spikes in protest activity on different days than the hashtag activism time series. This difference could be explained by the existence of a time lag between the two time series, where higher Pearson correlation values would be obtained through the use of a time lag between the two time series. This is investigated later in this section.

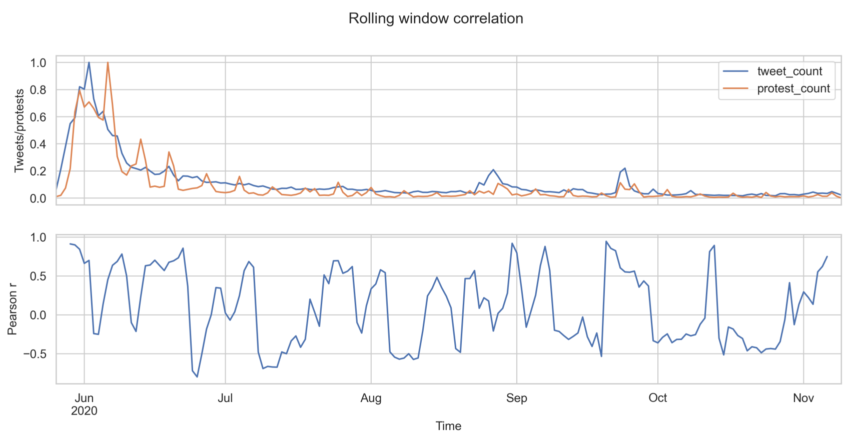

Figure 4.2: Rescaled hashtag acvitism and physical protest time series (top) and seven day rolling window correlation (bottom).

While fig. 4.1 and fig. 4.2 provide useful insight into the temporal relationship between hashtag activism and physical protest, they cannot answer whether this relationship exists in space as well. That is, if an *overall* increase in hashtag activism results in an *overall* increase in physical protest counts, does an increase in hashtag activism in *a particular location* result in an increase in physical protest counts in *that particular location*? With the use of data on social movement tweet counts and physical protest counts that is aggregated both in time and in space (i.e. at the state level), fig. 4.3 illustrates the temporal and statistical relationship between daily number of physical protests and daily number of tweets as scatter-plots for each of the U.S.'s 50 states. From this figure it is evident that the pattern previously observed at the national level also exists at the state level. Thus, in spite of the existence of a few outliers, the relationship between hashtag activism and physical protest counts in each state is described by a strong positive correlation. Across 50 states, the median Pearson correlation p value is 0.74 and there are 44 states associated with Pearson correlation p values greater than 0.6 (table 8.1. Additionally, similarly to the trend observed at the national scale, the relationship between state-level hashtag activism and physical protest counts can be modelled by a power regression function ($y_n = mx_n^b$), represented by the red dashed line in each subplot, where y_n is the number of physical protests that take place on a given day in a given state and x_n is the number of social movement tweets that are posted on a given day in a given state.

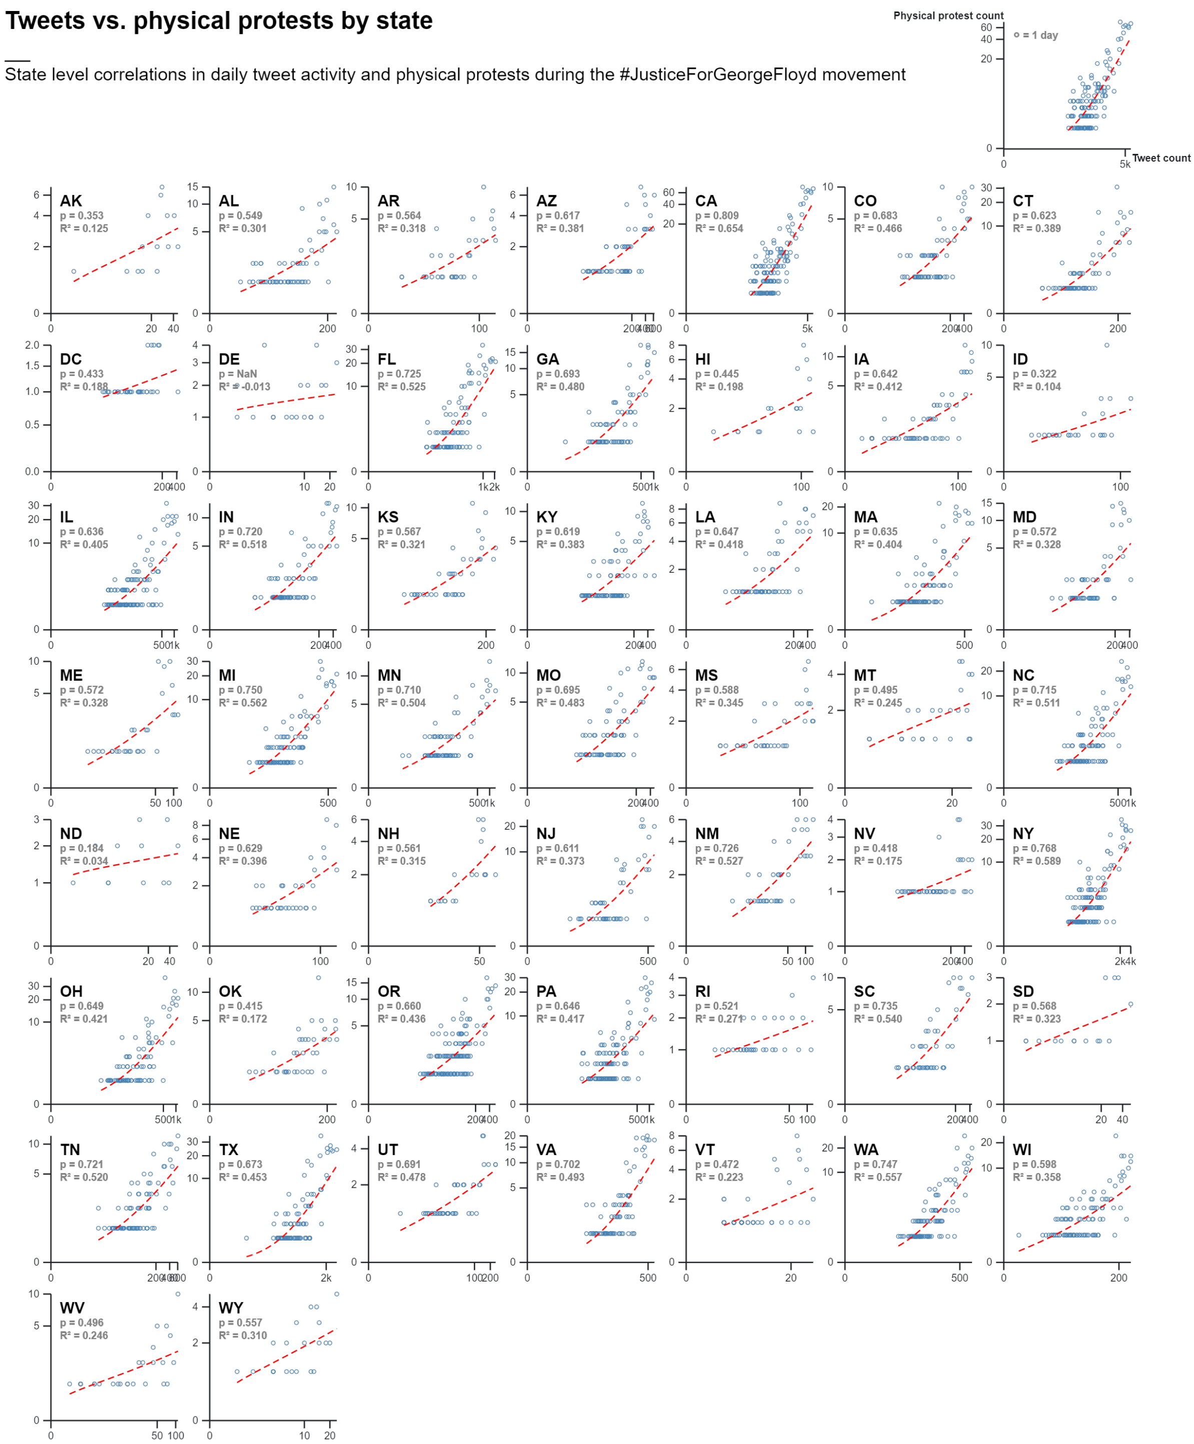

Figure 4.3: State-level statistical relationships between hashtag activism and physical protests using the #JusticeForGeorgeFloyd social movement as a case study.

In general, the accuracy of these models is described by a median R^2 value of 0.55 across 50 states, and 37 states that are associated with R^2 values greater than 0.5. Interestingly, fig. 4.4 shows that across the 25% most densely populated states, the power regression model yields considerably higher scores, with the median Pearson correlation p value climbing to 0.8 and the median R^2 values climbing to 0.64. Across the 25% least densely populated states, the opposite trend is observed. Summary statistics for each of the U.S. 50 states can be viewed in the 8.

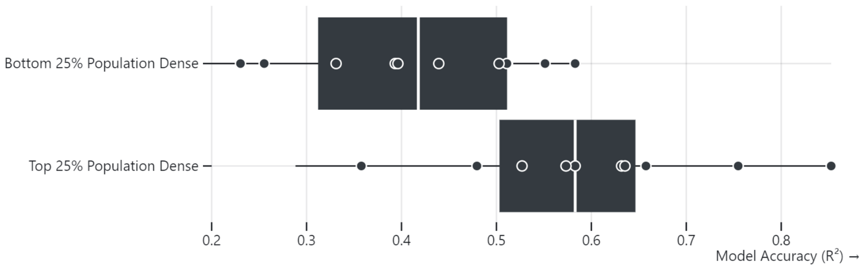

Figure 4.4: Accuracy of Power regression models (R^2) across the 25% most population dense U.S. states and the 25% least population dense U.S. states

In this context, the accuracy of power regression models at the state level seems to be very sensitive to population density. This is further illustrated in fig. 4.5, where population density is visualized as a power law function of model accuracy (R^2) for each state. In the right subplot, states with a land area smaller than the bottom 0.05 quantile (i.e. DE) are removed, as considerably less data-points are recorded within those states. Thus, from the right subplot of fig. 4.5, we can observe that model accuracy increases exponentially with population density (i.e. each axis is on a logarithmic scale).

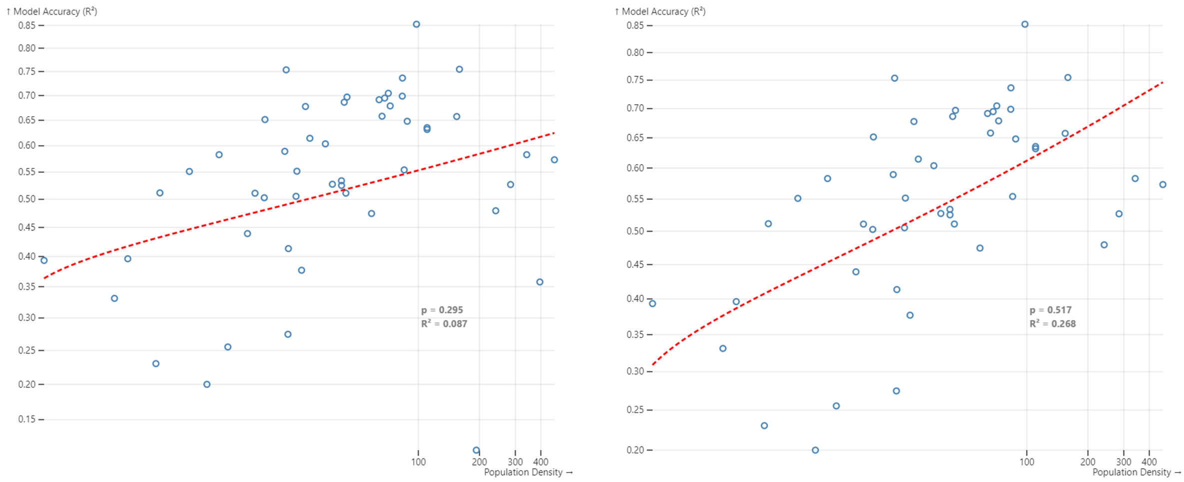

Figure 4.5: Accuracy of power regression models (R^2) as a function of population density. This relationship is shown for all 50 states (left) and for the 49 states of which land area is greater than the bottom 0.05 quantile (right).

With state-level data, the seven day rolling window Pearson correlation between hashtag activism and physical protests is computed for each of the U.S.' 50 states. Results from this analysis are illustrated by fig. 8.2 in the appendix. Similar to national level results, it can be seen from this figure that the Pearson correlation is generally a high value, but oscillates to very low values for brief periods of time. While this is the case in the majority of states, some states instead show a more steady temporal pattern in Pearson correlation p values. It can be observed that these states are not only the same states that obtained the lowest power regression model scores in fig. 4.3 (i.e. HI, ID, ME, ND, AK, AR, DE, MS, SD, WY, VT), they are also among the least densely populated states.

MOBILIZATION SYNERGY INDEX

Following the same methodology, the spatial scale is further disaggregated and the temporal relationship between hashtag activism and physical protest counts is investigated at the county level. Once again, we model this relationship as a power regression function ($y_n = mx_n^b$) for each of U.S. county where tweet and physical protest data was available (i = 2178 counties). Results from this analysis are illustrated by the bottom left and right subplots of fig. 4.7, where the obtained Pearson correlation p values and R^2 values are illustrated as kernel density distributions due to the large number of counties analysed. fig. 4.6 shows that, on average, the county level Pearson correlation p value is positive and relatively high. fig. 4.6 also shows that while a power regression model achieves a high R^2 score in a majority of counties, there is more variance in the accuracy of the power regression model.

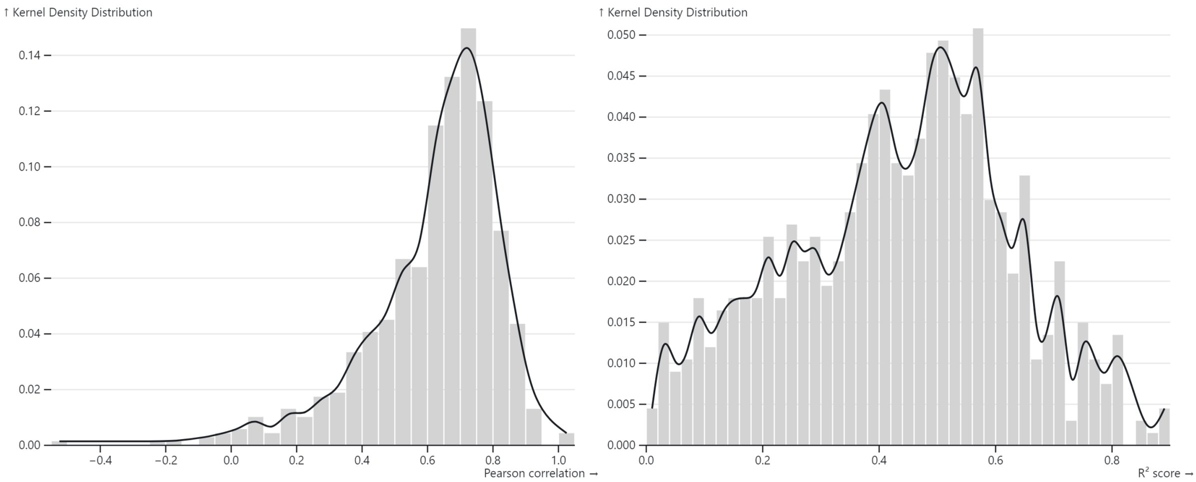

Figure 4.6: Distribution of Pearson correlation p values (left) and R^2 values (right) obtained from power regression modelling between tweet count and physical protest counts at the county level using the #JusticeForGeorgeFloyd case study.

One reason that could explain this variance could be the lack of data-points available in some less densely populated counties. Another reason is that in some counties one of hashtag activism or physical protest was more commonly used to engage in activism during the #JusticeForGeorgeFloyd social movement. In other words, in some counties people use physical protest to greater extents, while in other counties people may use hashtag activism more instead. In this context, modelling the relationship between hashtag activism and physical protest and visualizing this trade-off at the county level can provide unique insight into which locations experience synergy across digital and physical activism mediums, and which do not. Thus, on the one hand it can help identify which counties have activist groups that mobilize both digitally and physically, and thus where physical protests can be expected if people are already protesting on social media. On the other hand, for the counties that exhibit a weaker relationship between both forms of protest, it can generate insight as to *why* this relationship is weak. Such tendency can be understood with the help of residual values resulting from each county's power regression model. As such, the number of positive residual values represents the number of data-points (i.e. days) in which physical protests were underestimated by the regression model, and thus where the proportion of physical protests was larger than that of social movement tweets. Instead, the number of negative residual values represents the number of data-points in which physical protests were overestimated by the regression model, and thus where the proportion of physical protests was smaller than that of social movement tweets.

This type of classification can be achieved with the computation of a statistical index which we call the ***Mobilization Synergy Index (MSI)***. Such index is achieved by (a) calculating the number of positive and negative residuals r_{pos} and r_{neg} respectively from each individual county's power regression model where each residual r is $r = y - y_0$, y is the actual number of physical protests on a given day, and y_0 is the number of physical protests predicted by the regression model, and (b) calculating the proportions of positive residual values (days in which physical protest counts are disproportionally intense) and negative residual values (days in which tweets counts are disproportionally intense) as $r_p = r_{pos}/(r_{pos} + r_{neg})$ and $r_t = r_{neg}/(r_{pos} + r_{neg})$ respectively, and (c) multiplying each of these values by the product of the power regression model's slope m and R^2 score (so as to capture both the strength of the power relationship between tweet count and physical protest activity and the accuracy of the power regression model). As such, two measurements are recorded:

$$MSI_{protests} = r_p \times (m \times R^2) \tag{4.1}$$

$$MSI_{tweets} = r_t \times (m \times R^2) \tag{4.2}$$

As such, if in a given county $MSI_{protests} > MSI_{tweets}$ then physical protest was disproportionately intense over the course of #JusticeForGeorgeFloyd, if instead $MSI_{protests} > MSI_{tweets}$ then tweet activity was disproportionately intense over the course of #JusticeForGeorgeFloyd, and finally if $MSI_{protests} \approx \mathrm{MSI}_{tweets}$ then physical protests and tweet counts followed a strong power law relationship where more tweets on a given day translates in more physical protests on that same day.

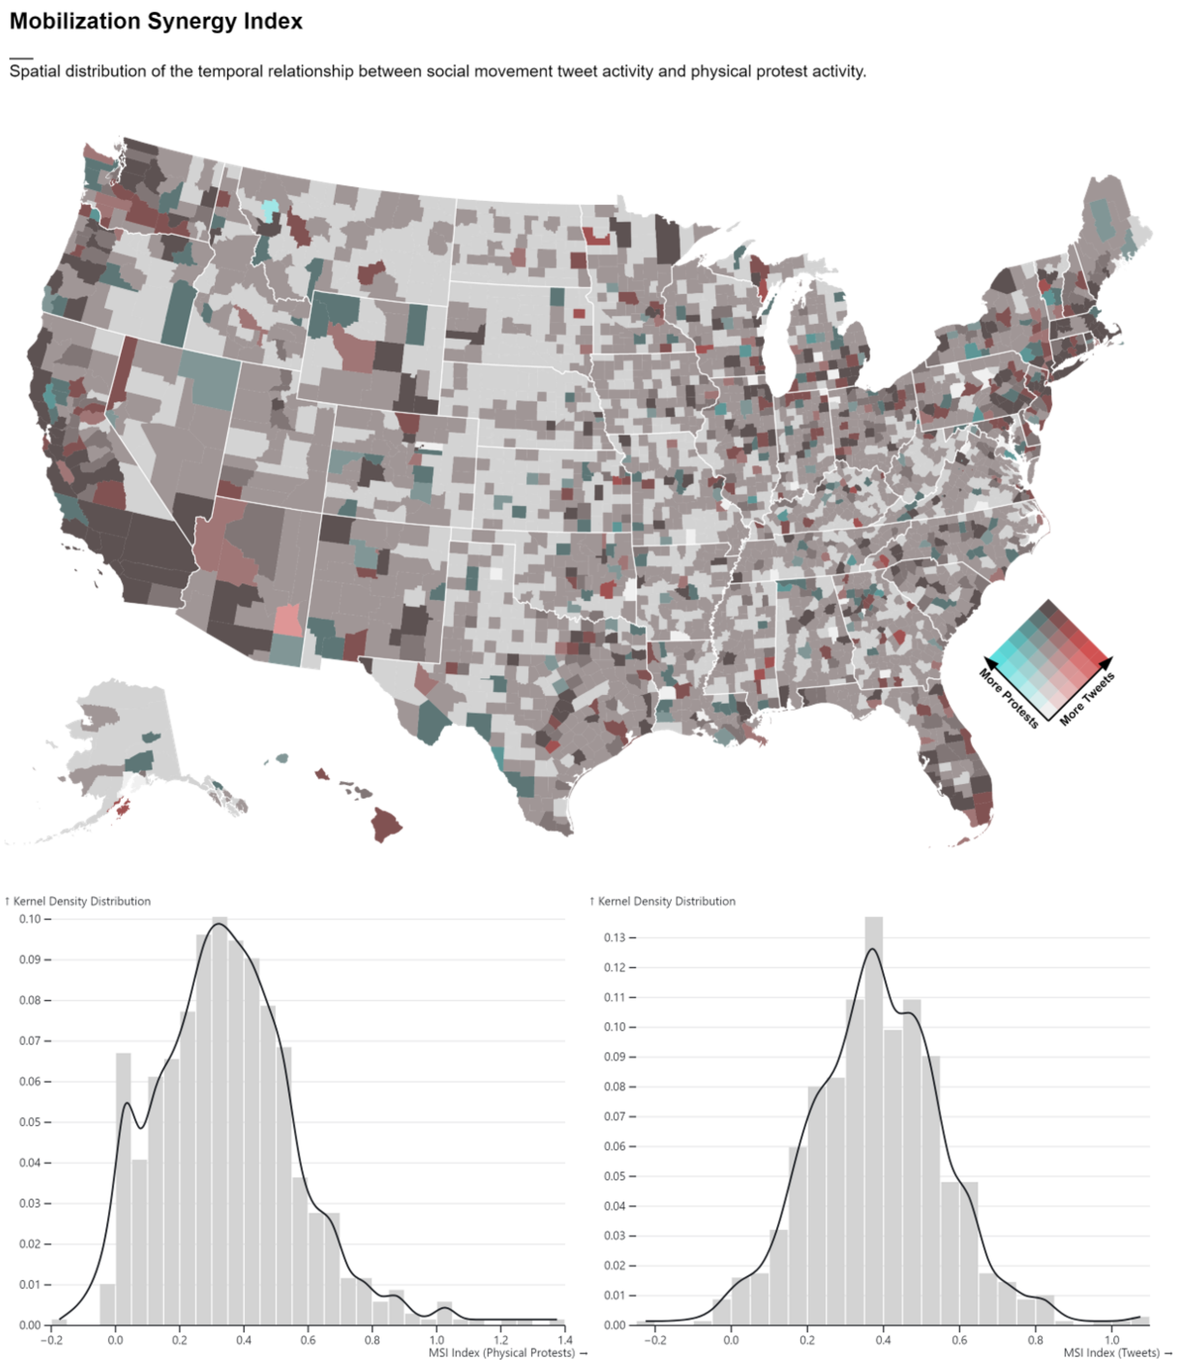

Figure 4.7: Spatial (top) and statistical (bottom) distribution of the Mobility Synergy Index (MSI): darker shades represent counties in which people are likely to protest both digitally and physically on a given day. Shades of blue represent counties in which people are more likely to protest physically on a given day. Shades of red represent counties in which people are more likely to protest digitally on a given day.

The *Mobilization Synergy Index (MSI)* and its spatial distribution is visualized through a bi-variate choropleth map in the top subplot of fig. 4.7. In this map, darker shades represent counties in which $MSI_{protests}$ and MSI_{tweets} converge, and thus where people are likely to protest both digitally and physically on a given day. In contrast, shades of blue represent counties where $MSI_{protests}$ outweighs MSI_{tweets}, and thus where people are more likely to protest physically on a given day. Finally, shades of red represent counties where MSI_{tweets} outweighs $MSI_{protests}$, and thus where people are more likely to protest digitally on a given day. From fig. 4.7, it can be observed that the areas where the temporal relationship between tweets and physical protests is stronger are concentrated along the east and west coastlines, while the areas where this relationship is weaker are concentrated in the country's interior landmass. Areas where people choose to engage in hashtag activism over physical protest, and vice versa, doe not exhibit any discernible spatial pattern.

TIME-LAGGED CROSS-CORRELATION ANALYSIS

Political mobilization processes such as hashtag activism and physical protest are complex in nature as they occur in uncontrolled environments where a multitude of influences are at play. For this reason, it is virtually impossible to establish causal relations between them. Still, some insight can be generated regarding the way in which these two processes may influence each-other in space and time.

As discussed in chapter 3, Time Lagged Cross Correlation (TLCC) or Lagged Regressions is a way to study the directionality between two time series, x_t an y_t, and identify lags of x that may be the best predictors of y_t. In the case of hashtag activism and physical protests, this method involves fitting a power regression model $n * 2$ times using the y_t data (daily physical protest counts) and the x_{t-n} and x_{t+n} (daily tweet counts) data, where $|n|$ is the maximum number of time lags. fig. 4.8 illustrates the results of performing TLCC on national level data using a time lag of $n = 10$. The top subplot shows that the optimal Pearson correlation p value and R^2 score is obtained when hashtag activism (i.e. daily number of tweets posted) leads physical protests by one day. In other words, at the national scale and during the case study of #JusticeForGeorgeFloyd, the best predictor of physical protest activity at time t (in days) is the number of social movement related tweets posted at time $t-1$ (i.e. the day before). This is relationship is further illustrated in the bottom subplot, where the outputs of power regression models using x_t (left) and x_{t-1} (right) can be compared.

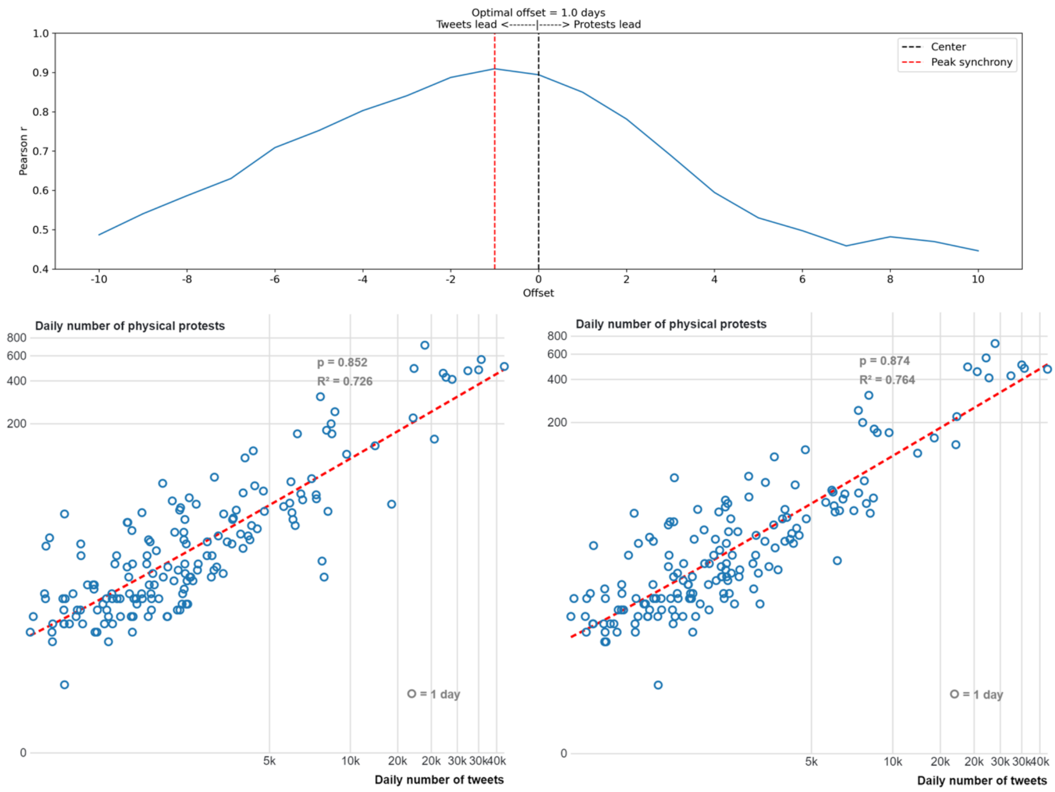

Figure 4.8: Time Lagged Cross Correlation results for hashtag activism and physical protests at $n = 10$ time lags (top). Power regression models for hashtag activism and physical protests using same day tweet counts x_t (bottom left) and previous day tweet counts x_{t-1} (bottom right).

With state-level data, the Time Lagged Cross Correlation (TLCC) between hashtag activism and physical protests is computed for each of the U.S.' 50 states. Results from this analysis are summarized by fig. 4.9 and illustrated in detail by fig. 8.3 in appendix 8. Similar to national level results, in 41 out of 50 states the optimal Pearson correlation p value is obtained when tweets lead by one or more days. For 20 of those states, the optimal Pearson correlation p is obtained with a time lag of one day.

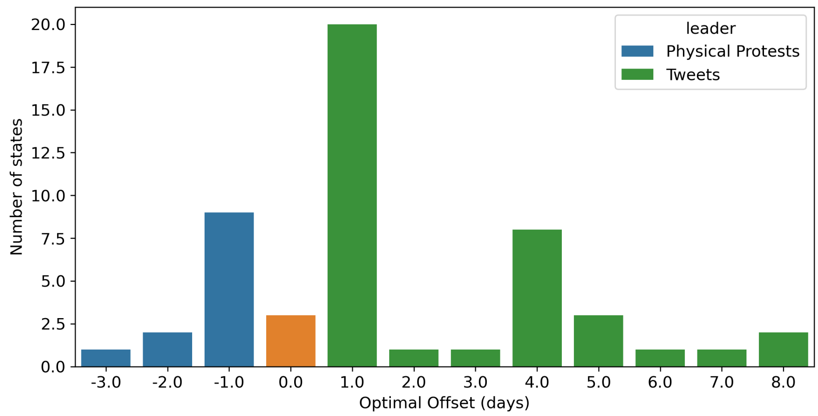

Figure 4.9: Optimal time lags for modelling the relationship between hashtag activism and physical protest by state.

Granger causality analysis

As discussed in chapter 3, Granger causality is a statistical hypothesis test developed to study the flow of information between two or more time series (Granger, 1969). Essentially, Granger causality determines whether prior values from one time series can be used to predict current and future values of another time series. Thus, while Granger causality is not a test of *True* Causality, it is said that a time series x *Granger-causes* y if *lagged* values of x provide statistically significant information about future values of y (Granger, 1969).

With national and state-level data, Granger causality is implemented to study whether data on tweet counts can provide information that is statistically significant for the forecast of future physical protest counts. Results from this analysis are summarized in fig. 4.10 and presented in table 8.2 of appendix 8. Overall, at 95% confidence level, at the national level and for 49 states we can reject the null-hypothesis that tweet count does not cause protest count. Thus, it is found that tweet count *Granger causes* protest count at the national level and in 49 out of 50 States. These results are in harmony with results obtained from the TLCC method.

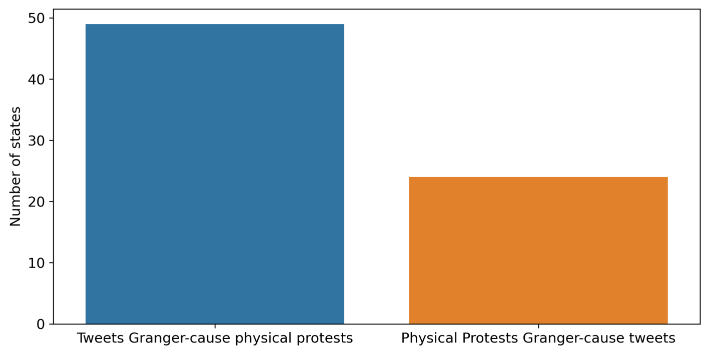

Figure 4.10: State-level results of Granger causality analysis for the hypothesis that tweet count Granger causes protest count (blue) and the hypothesis that protest count Granger causes tweet count (orange).

With national and state-level data, Granger causality is also implemented to study whether data on protest counts can provide information that is statistically significant for the forecast of future tweet counts. Results from this analysis are summarized in fig. 4.10 and presented in table 8.4 of appendix 8. Overall, at 95% confidence level, at the national level and for 27 states we cannot reject the null-hypothesis that protest count does not cause tweet count. Thus, it is found that protest count does not *Granger cause* tweet count at the national level and in 27 out of 50 States.

MULTIVARIATE TIME SERIES FORECASTING

As discussed in chapter 3 of this research, time-series forecasting is the use of temporal patterns extracted during time-series analysis to predict future values from past values (Wei, 2006). In the context of hashtag activism and physical protest, the coincident temporal patterns of these two processes can be leveraged by time-series forecasting to make future predictions for either process. In this research, we developed a simple Long Short Term Memory (LSTM) artificial neural network as well as a LSTM artificial neural network combined with a Vector Autoregressive (VAR) model (LSTM VAR) to predict future counts of physical protests based on past values of both physical protests and tweet counts. Each model was trained and fine-tuned on aggregated national level data, and was used to predict physical protest counts at the national and state-level. Fig. 4.11 illustrates the predictive performance of the LSTM (red) and the LSTM VAR (blue) models. The left subplot of fig. 4.11 illustrates the predictions made by each model for physical protest counts over a period of two months (September 6^{th} to November 8^{th}) alongside *actual* or *observed* physical protest counts over the same time-frame shown in grey. The right subplot illustrates the prediction error y^e, as $y^e = y - y^0$ where y is the actual value and y^0 is the predicted value, made by each model at each time-step. At the national level, we can observe that both models where able to predict the general temporal pattern of physical protest activity. Interestingly, however, each model seems to have overestimated larger spikes of physical protest, while at the same time underestimated smaller spikes of physical protest. Additionally, the LSTM VAR model overestimated the largest spike of physical protest (on September 25^{th}) by a considerably larger margin than the simple LSTM model.

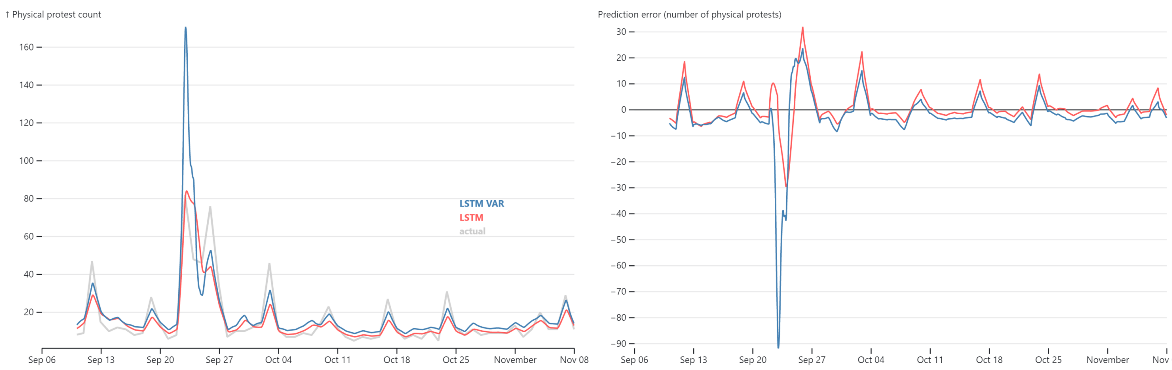

Figure 4.11: National-level physical protest count predictions (left) and prediction error (right) with LSTM neural network (red) and combined VAR + LSTM neural network (blue).

The predictive performance of each model at the national level is further illustrated in fig. 4.12, where we plot predicted physical protest counts on the x-axis against actual physical protest counts on the y-axis for the LSTM (red) and LSTM VAR (red) models. In the case of both models, we fit a power-regression model to the values of predicted and actual physical protest counts and compute its accuracy or R^2 score. Fig. 4.12 shows that while on average both the LSTM and LSTM VAR models were able to make accurate predictions of physical protest counts, the LSTM model scored a much higher accuracy score. This is likely due to the single instance (September 25^{th}) where the LSTM VAR model largely overestimated physical protest counts.

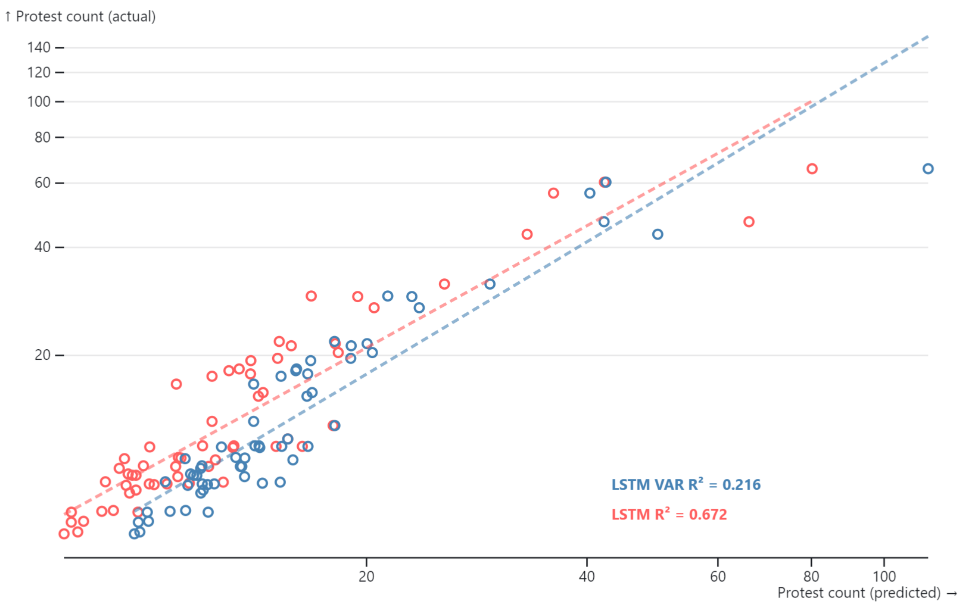

Figure 4.12: National-level predicted vs. actual physical protest counts with LSTM neural network (red) and combined VAR + LSTM neural network (blue).

In order to test the generalization capability of both the LSTM and LSTM VAR models, the predictive performance of each model was tested for state-level physical protest and tweet count data. Accordingly, the models were used to predict state-level physical protest counts throughout the entire time-frame of the #JusticeForGeorgeFloyd social movement (five months). For both models, state-level physical protest count predictions, aggregated model prediction errors, and disaggregated model prediction errors are illustrated in fig. 8.6, table 8.5, and fig. 8.7 of the appendix, respectively. State-level model performance results are illustrated in fig. 4.14, where we plot state-level predicted physical protest counts on the x-axis against actual physical protest counts on the y-axis for the LSTM (red) and LSTM VAR (red) models. From fig. 4.14, we can observe that both models were successful at accurately predicting daily physical protest counts for each state. In all cases, both models were able to generalize the temporal patterns learned at the national level to predict future instances of physical protest based on past instances of physical protest and tweet activity at the state level.

The state-level performance of each model is further illustrated in fig. 4.13, where the R^2 scores obtained from power regressions of actual values vs. predicted values are compared for each model. Fig. 4.13 shows that, in most cases, the LSTM VAR model was consistently more accurate in predicting daily, state-level, physical protest counts, than the simple LSTM model. The LSTM VAR model scored higher than the LSTM model for 37 states with an average R^2 accuracy score of 0.87, whereas the LSTM model only scored higher than the LSTM VAR model in 13 states with an average R^2 accuracy score of 0.79.

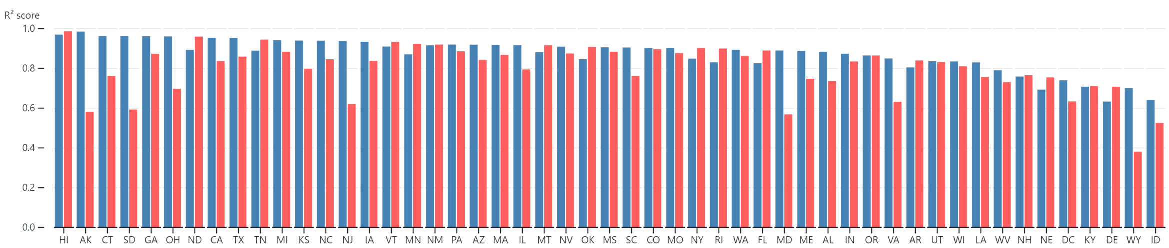

Figure 4.13: State-level physical protest prediction accuracy scores (R^2) with LSTM neural network (red) and combined VAR + LSTM neural network (blue).

Figure 4.14: State-level predicted vs. actual physical protest counts with LSTM neural network (red) and combined VAR + LSTM neural network (blue).

4.2. THE SPATIO-TEMPORAL RELATIONSHIP BETWEEN LOCATION-BASED HASHTAG ACTIVISM AND LEGISLATIVE ACTIVITY

As discussed in chapter 2 of this research, empirical evidence about the relationship between hashtag activism and legislative action is lacking in the field of social movement research. In an attempt to address this knowledge gap, this section presents the results of time-series modelling, regression analysis, and geo-spatial analysis of hashtag activism and legislative action using the case study of #JusticeForGeorgeFloyd. Thus, this section answers the second sub-question of this research: ***What is the relationship between location-based hashtag activism and legislative action?***

REGRESSION ANALYSIS ON AGGREGATED DATA

One way to quantify the impact of hashtag activism on legislative responses is to model the linear relationship between tweet counts and counts of legislative responses aggregated at the state level. Figure 4.15 (left) illustrates this relationship as a scatter plot where each dot represents a U.S. state, the x-axis represents the total number tweets (normalized by population) generated within that state, and the y-axis represents the total number of legislative responses to policing within that same state since May 25th, 2020. Similarly, the right subplot of fig. 4.15 illustrates the relationship between physical protest and legislative responses to policing as a scatter plot where each dot represents a U.S. state, the x-axis represents the total number physical protests (normalized by population) that took place within that state, and the y-axis represents the total number of legislative responses to policing within that same state since May 25th, 2020. The relationship between both tweet counts and physical protest counts with legislative responses to policing is modelled as a power regression function, represented by the red dashed lines in fig. 4.15.

Figure 4.15 shows that increased tweet counts per-capita in one state bear a positive correlation with legislative responses in the same state, while increased protest activity per-capita does not relate to more legislative responses. In fact, without consideration of the temporal dimension physical protests bear no relationship with subsequent state-level legislative change.

Figure 4.15: Power regression models for hashtag activism and number of legislative responses to policing (left), and physical protests and number of legislative responses to policing (right)

TIME-LAGGED CROSS-CORRELATION ANALYSIS

While state-level aggregated tweet counts and physical protest counts can provide some insight into the overall effects of hashtag activism and physical protest on legislative responses, they may also provide misleading information as they are time independent. As we have seen in 4.1, hashtag activism and physical protest processes are highly time-dependent, and thus deeper insight into their relationship with legislative responses may be generated by considering this dimension for analysis. As described in chapter 3, point data describing instances of hashtag activism and physical protests, and counts of legislative responses to policing were aggregated in time and space, so that the temporal relationship between (a) hashtag activism and legislative responses and (b) physical protest and legislative responses could be studied at different spatial scales. fig. 4.16 illustrates the temporal patterns of each of legislative responses to policing (top), hashtag activism (middle), and physical protest activity (bottom), at the national level, and reveals that legislative responses are also time dependent. Indeed, fig. 4.16 indicates that the highest counts of legislative responses to policing took place many months after peaks of hashtag activism and physical protest activity.

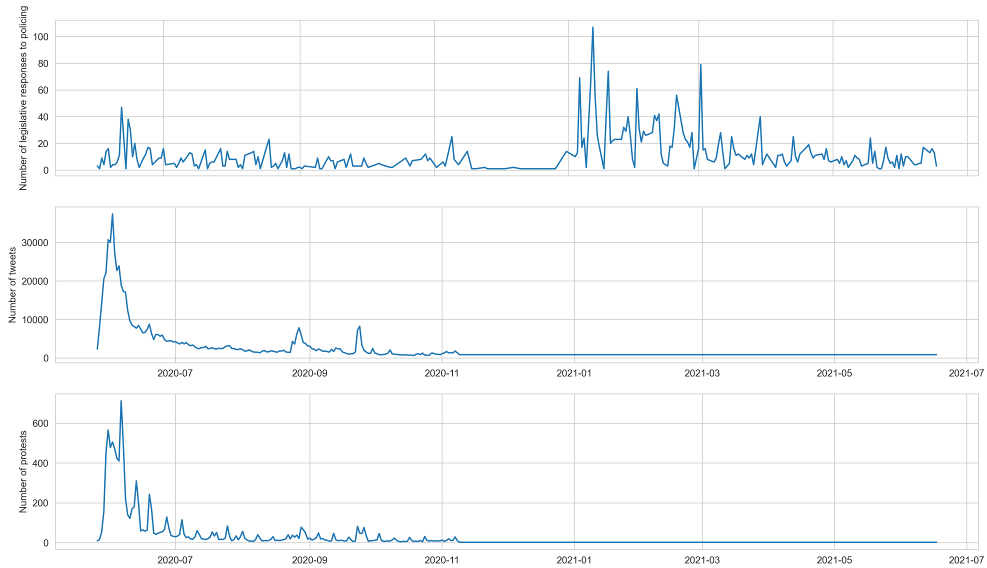

Figure 4.16: Temporal patterns of legislative responses to policing (top), hashtag activism (middle) and physical protest (bottom) at the national level

Similar to the analysis presented in 4.1, the temporal relationship between both of hashtag activism ($x1$) and physical protest ($x2$) and legislative responses to policing (y), is studied through TLCC analysis, where the Pearson correlation p value between $x1$ and y and $x2$ and y is computed repeatedly such that lags of $x1$ and $x2$ that yield the highest correlations y_t are identified. fig. 4.18 illustrates the results of performing TLCC on national level data using a maximum time lag of $n = 400$. The top subplot shows that the optimal Pearson correlation p value between hashtag activism and legislative responses is obtained when hashtag activism (i.e. daily number of tweets posted) leads legislative responses by 263 days. The bottom subplot shows that the optimal Pearson correlation p value between physical protests and legislative responses is obtained when physical protests lead legislative responses by 264 days.

This temporal relationship between both hashtag activism and physical protest and legislative responses to policing is further investigated by performing Lagged Regressions on national level data. Similar to the analysis presented in 4.1, we use Lagged Regressions to study the directionality between each of $x1_t$ and $x2_t$, and y_t, and identify lags of $x1$ and $x2$ that may be the best predictors of y_t. As such, a power regression model is fitted $n*2$ times using the y_t data (daily counts of legislative responses) and the $x1_{t+-n}$ (daily tweet counts) and $x2_{t+-n}$ (daily physical protest counts) data, where $|n|$ is the maximum number of time lags. Tables 4.1 and 8.1 illustrate the results of performing Lagged Regressions on national level data using a maximum time lag of $n = 400$. For the power regression between tweet counts and counts of legislative responses to policing the optimal time lag is 235 days, while for the power regression between physical protest counts and counts of legislative responses to policing the optimal time lag is 236 days.

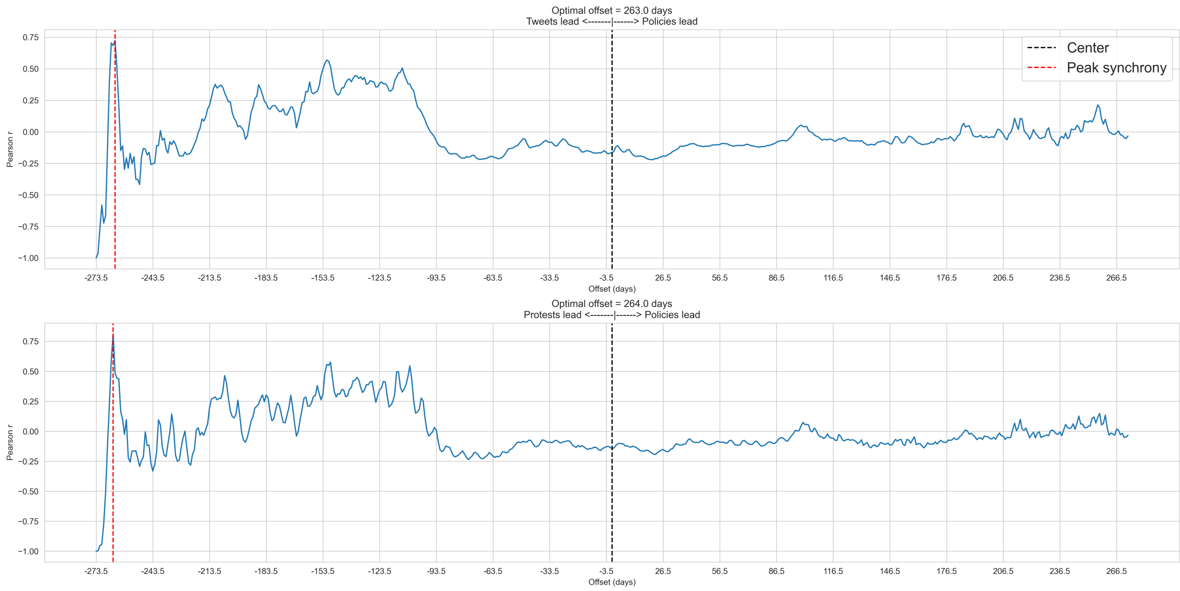

Figure 4.17: Time Lagged Cross Correlation results for hashtag activism and legislative responses (top) and physical protests and legislative responses (bottom) at $n = 400$ time lags.

lag (days)	Parson r	R^2
236	0.494	0.244
237	0.469	0.220
230	0.447	0.200
229	0.441	0.195
238	0.440	0.194
235	0.432	0.186
228	0.425	0.181
244	0.425	0.181
224	0.420	0.177
239	0.418	0.175

Table 4.1: Top 10 optimal time lags of the tweet count variable for power regressions between tweet counts and counts of legislative responses to policing, and associated Pearson correlations p and R^2 values.

lag (days)	Parson r	R^2
235	0.500	0.250
220	0.478	0.228
234	0.471	0.221
221	0.458	0.210
236	0.454	0.206
262	0.449	0.201
228	0.447	0.200
263	0.447	0.200
227	0.444	0.197
242	0.436	0.190

Table 4.2: Top 10 optimal time lags of the physical protest count variable for power regressions between physical protest counts and counts of legislative responses to policing, and associated Pearson correlations p and R^2 values.

In other words, at the national scale and during the case study of #JusticeForGeorgeFloyd, the best predictor of the number of legislative responses at time t (in days) is the number of social movement related tweets posted at time $t-236$ (in days) or the number of social movement related physical protests that took place at time $t-235$ (in days). This relationship is further illustrated in fig. 4.18, where the outputs of both power regression models ($x1_{t-235}$ and y_t (left), $x2_{t-236}$ and y_t (right)) can be compared. In each subplot each dot is a day, the x-axis represents the daily number of tweets at time $t-235$ (left) and the daily number of physical protests at time $t-236$ (right), and the y-axis represents the daily number of legislative responses to policing at time t.

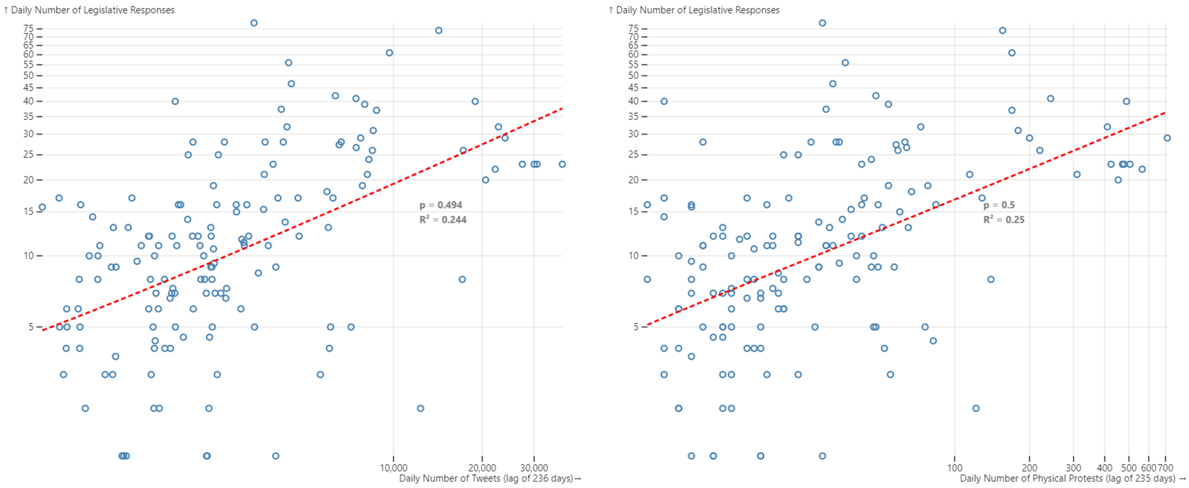

Figure 4.18: Power regression functions for tweet count and number of legislative responses (left) and physical protest count and number of legislative responses (right) at optimal time lags of 235 days and 236 days respectively.

With state-level data, the Time Lagged Cross Correlation (TLCC) between both hashtag activism and legislative responses and physical protest and legislative responses is computed for each of the U.S.' 50 states. Results from this analysis can be viewed in detail in fig. 8.4 and fig. 8.5 of the appendix. On average across states, the optimal Pearson correlation p value is obtained when tweets lead by 179 days. In other words, there is an average time lag of 179 days between the time in which hashtag activism peaks and the time in which the majority of legislative responses occur. The average time lag between the time in which physical protests peak and the time in which the majority of legislative responses occur is also 179 days. Figure 4.19 illustrates the optimal time lags for modelling the relationship between state-level hashtag activism and legislative responses (top) and the relationship between state-level physical protest and legislative responses (bottom). From fig. 4.19, it can be observed that, while a majority there is considerable variance in time lags across states. In other words, some states experienced the majority of their legislative responses very close in time to when they experienced the highest peaks in hashtag activism and physical protest, while others experienced the majority of their legislative responses very far in time to when they experienced the highest peaks in hashtag activism and physical protest.

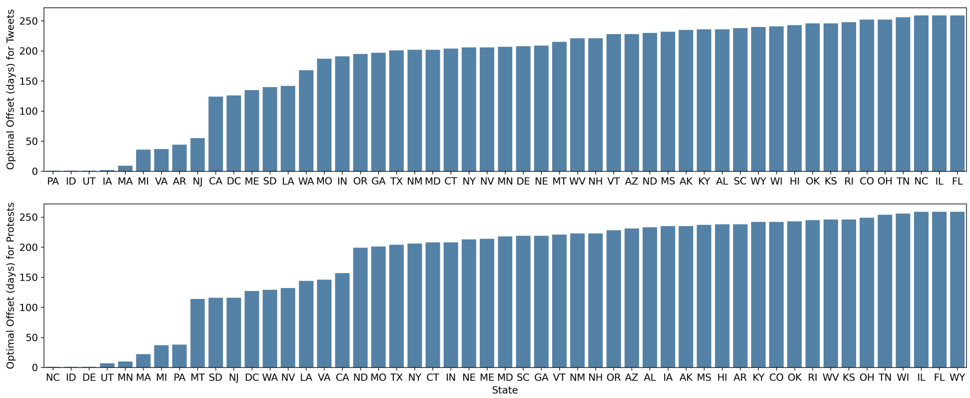

Figure 4.19: Optimal time lags between legislative responses to policing and tweet counts (top) and protest counts (bottom), by state.

5

DISCUSSION

THE digitization of our world in the past two decades has drastically changed the functioning of our societies. Perhaps most drastic is how digital technologies have transformed the way in which we communicate with others and participate in public discourse. In this context, social media has (a) transferred public discourse online and (b) made it possible to connect with people well beyond our "traditional" or physical social circles. One of the biggest impacts of this phenomenon is the use of social media by marginalized or underrepresented social groups as a tool to voice their opinions, world-views, and experiences in a world where white, patriarchal, and hetero-normative ideals are the status-quo (Ortiz *et al.*, 2019; Nemer, 2016; J. J. George and Leidner, 2018; Devito *et al.*, 2019; Trevisan, 2019; Johnson and Callahan, 2013). Through this process, the evolution and strengthening of collective identities has become amplified and the ability to engage in **hashtag activism**, organize politically, and take collective action has become easier (Ray *et al.*, 2017; Lăzăroiu *et al.*, 2018; Gerbaudo and Treré, 2015). As a result, social media has been at the heart of important social movements, such as the Arab Spring, #MeToo, and more recently #JusticeForGeorgeFloyd, critical for the progress of our societies towards more just, accountable, and inclusive standards (Khondker, 2011; Manikonda *et al.*, 2018).

But ***what are the true political impacts of location-based hashtag activism during and after these social movements****, and can they be measured?*

This question is precisely what this research set out to investigate, as an effort to narrow the gap between our theoretical and qualitative understanding of the political consequences of hashtag activism, and the empirical evidence that supports this understanding. More specifically, the objectives of this research were to leverage contemporary data-sources and state of the art quantitative methods to explore and quantify the impacts of hashtag activism on the processes of physical protest and legislative action. Through a combination of *time-series analysis*, *regression analysis*, *geo-spatial analysis*, and *machine learning* for the analysis of data-sets originating from *social media*, *non-profit organizations*, and *government*, the temporal, dynamic, and spatial dimensions of hashtag activism, physical protest, and legislative action were unraveled, using the 2020 #JusticeForGeorgeFloyd social movement as a case study.

This fifth chapter discusses the main results presented in chapter 4 and the possible implications of this research within the realms of academia, policy, and activism. The chapter first discusses each result in (a) the scientific context by outlining if and how it constitutes a contribution to the broader research landscape; (b) the socio-political context by outlining its societal meaning and possible implications for policy making. Second, it discusses the scientific limitations of the research and describes areas of future research to address these limitations.

5.1. DISCUSSION OF KEY RESULTS

5.1.1. LOCATION-BASED HASHTAG ACTIVISM AND PHYSICAL PROTEST ACTIVITY

Physical protest, "the public use of protest by groups or organizations that seek to influence a political decision or process, which they perceive as having negative consequences for themselves, another group or society as a whole" (Rucht, 2007), is the most tangible form of political mobilization in modern societies. As such, in the twentieth century physical protest evolved from its previously illegitimate connotation of "irrational outburst of the dangerous classes" to a legitimate and effective political strategy to influence political action at various levels of government (Rucht, 2007). In the twenty-first century, digital technologies have enabled protesters to also engage in protest through the internet space and, mostly social media, and thus hashtag activism was born. As a result, contemporary social movements are characterized by sustained episodes of both physical and digital protest, but the interplay between these two forms of protest is largely undetermined. Thus, in order to quantify the political impacts of location-based hashtag activism during social movements this research set out to investigate a first relevant question:

What is the relationship between location-based hashtag activism and physical protest activity?

Through implementation of the proposed experimental design introduced in section **??** of chapter 3, this research resulted in several key findings that provide useful insights into this question. First, in the case of #JusticeForGeorgeFloyd, it was found that at the national scale, hashtag activism bears a very strong positive temporal relationship with physical protest activity. Through the statistical modelling of this relationship, it is possible to accurately predict the intensity of national level protest activity on a given day based on the intensity of national level hashtag activism the day before. This observation is in accordance with the qualitative evidence generated by Tufekci and C. Wilson (2012) on the role of social media in increasing the likelihood of participation in physical protest during the Tahrir Square protests, and more importantly with the empirical evidence generated by Brantly (2019) and Bastos *et al.* (2015) on the potential existence of a one day time-lag between hashtag activism and physical protest activity.

Second, it was found that this temporal relationship between hashtag activism and protest activity exists at the state level as well. By modelling this relationship at the state level, through the use of traditional regression models and the use of more complex deep learning frameworks, it is possible to predict the intensity of state level protest activity on a given day based on the intensity of state level hashtag activism the day before. In general, it was found that the accuracy of such predictions increases exponentially with population density. This could be explained by the fact that population dense states may have a larger number of urban centers, where social (media) networks may be more connected and thus where people are more aware of both local and national news. At the same time, urban centers may make it easier for people to gather in public spaces to engage in protest. The leader-follower relationship between hashtag activism and physical protest activity was further supported by the results of *Granger causality analysis*, where it was found at both the national and state level that hashtag activism Granger causes physical protest activity in the majority of cases. As a result, this research proposes that hashtag activism may act as a causal factor for physical protest activity, and not the other way around, where this causal relationship is also strongly related to physical space. This would be in accordance with other research within the broader topic of hashtag activism and social movements which has revealed unidirectional relationships between hashtag activism and other characteristics of social movements (e.g. the framing of political narratives; the polarization of social groups) (Jackson *et al.*, 2020; Kent, 2013; Carney, 2016; M. Li *et al.*, 2021; Xiong *et al.*, 2019). In this context, a causal spatial relationship between hashtag activism and physical protest would imply that today social media is not only a vital tool for activist groups to mobilize people physically, but also to mobilize people physically in targeted geographical locations.

Third, it was found that the spatio-temporal relationship between hashtag activism and protest activity continues to hold at the county level. At such small scale, the accuracy of predictions remained generally high, but considerably more variance was observed and thus, a wider error margin is noted. As discussed in section 4.1 of chapter 4, this is explained by the fact that in some "outlier" counties the temporal relationship between hashtag activism and physical protest is characterized by a disproportionate amount of physical protest compared to the intensity of hashtag activism, while in other counties this relationship is characterized by a disproportionate amount of hashtag activism compared to the intensity of physical protests. In order to make sense of this tendency this research proposed the *Mobilization Synergy Index*, which makes it possible to statistically quantify this trade-off, and to visualize its spatial distribution in an intuitive way. Indeed, both policy makers and activists could gain valuable insight from this index. On the one hand, local policy makers could utilize it to understand how their communities mobilize in times of crisis, and become responsive to local demands accordingly. For example, if a specific county is categorized by the index as a county where the relationship between hashtag activism and physical protest is strong, local policy makers could avoid potential public disruption by fulfilling community needs rapidly once local levels of hashtag activism peak. In contrast, if a specific county is categorized by the index as a county where people use hashtag activism disproportionally, this could be a sign for local policy makers to pay closer attention to social media platforms for understanding community needs. On the other hand, activists could utilize the index to plan their mobilization efforts more efficiently and effectively. For example, by identifying counties in which people engage in hashtag activism intensely but do not mobilize physically, local activist groups could modify their strategies for physical mobilization so that more people show up at protests. In contrast, in counties where people engage in physical protest intensely but not in hashtag activism, local activist groups may increase their use of offline mobilization strategies to maximize protest turnout.

5.1.2. LOCATION-BASED HASHTAG ACTIVISM AND LEGISLATIVE ACTION

Legislative action, the development, reform, or creation of legal structures, is the most tangible form of political action in democratic countries. For this reason, a majority of social movements have as a primary objective to influence political processes through triggering concrete action at the legislative level (Amenta *et al.*, 2010). While there exists limited evidence of instances in which legislative action has occured during or after social movements (Amenta *et al.*, 2010; Enos *et al.*, 2019; Weldon *et al.*, 2011; Lipsky, 1968), the relationship between legislative action and political protest, especially hashtag activism, has largely been unexplored through a quantitative lens. For this reason, this research set out to investigate a second relevant question:

What is the relationship between location-based hashtag activism and legislative action?

Through the implementation of the proposed experimental design introduced in chapter 3, several key findings were obtained that provide useful insight into this question. First, data on instances of hashtag activism aggregated by state and across the full time-frame of the #JusticeForGeorgeFloyd social movement revealed that states in which people engage more in hashtag activism on average also experienced more legislative responses related to policing in total. Unexpectedly, the opposite relationship was observed with regards to the relationship between average physical protest activity and the number of subsequent legislative responses to policing. One interpretation for these opposite findings could be that, during the #JusticeForGeorgeFloyd social movement, hashtag activism in a specific state was more effective at influencing legislative action in that same state than physical protest. In this case, state representatives would appear to be more sensitive to their voters' concerns when they are expressed through social media. An argument that could support this theory is that digital protests involve more people than physical protests (e.g. one tweet can be retweeted millions of times) and, as a result, policymakers may feel more pressure to act in response to them. What is more, social media platforms like Twitter make it easier for activists to publicly target and pressure specific politicians through "cancel culture" (D. Clark, 2020). As a result, political issues can quickly turn into personal issues. However, results above-discussed on the spatio-temporal relationship between hashtag activism and physical protest suggest that this interpretation is likely not accurate. Indeed, since state level physical protest activity and hashtag activism are strongly correlated in time, state level results aggregated across the full time-frame of the #JusticeForGeorgeFloyd social movement likely reflect the presence of outliers in the data (i.e. days with unusually high levels of physical protest or hashtag activism) that when aggregated may have skewed some data-points towards high or low extremes.

In this context, results from temporally disaggregated national level and state level data confirm this hypothesis. At both spatial scales it was found that the number of legislative responses on a given day bears a very similar positive relationship with both hashtag activism activity and physical protest activity taken at a specific time lag prior to the day in which the legislative responses occurred. At the national level, the number of legislative responses at time t (in days) can be predicted from the number of social movement related tweets posted at time t_{-236} (in days) or the number of social movement related physical protests that took place at time t_{-235} (in days). The accuracy of these predictions, however, was associated with considerable variance. Thus, while hashtag activism activity and physical protest activity can inform a general idea of future legislative developments, they may not be used to predict specific counts of legislative responses in the future. This research argues, however, that such an objective is not relevant in the context of measuring the political impacts of hashtag activism, and that the insights generated from this analysis may prove very useful for activists. As such, the simple observation of a positive relationship between the two of hashtag activism and physical protest activity and legislative action represents important positive feedback for the efforts of activists during #JusticeForGeorgeFloyd. What's more, these findings imply that hashtag activism and physical protest are both effective strategies for influencing legislative action, and that the more intensely they are used, the more likely it is that activist demands will be reflected into law and policy.

Finally, at the state level it was found that while both hashtag activism and physical protest activity result in future legislative responses, this temporal relationship is highly variable by state. As observed in fig. 4.19 of chapter 4, some states experienced the majority of their legislative responses very close in time to when they experienced their highest peaks of hashtag activism and physical protest activity, while others experienced the majority of their legislative responses very far in time to when they experienced their highest peaks of hashtag activism and physical protest activity. This tendency does not necessarily suggest that state-level representatives were more responsive in some states and less responsive in other states, as a state may experience faster legislative responses but less legislative responses in total. Nonetheless, this tendency indicates that more complex, state-specific activism strategies may have taken place during the #JusticeForGeorgeFloyd movement which could have influenced the speed at which each state responded politically to the social movement. However, exploring these issues was outside the scope of this study and is left for future research.

5.2. RESEARCH IMPACT

This research represents the first national-level quantitative study aimed at measuring the temporal relationships between hashtag activism, physical protest activity, and legislative action, at various spatial resolutions. It adds several potential scientific contributions to the field of social movements research and, more broadly, the computational social sciences.

1. The research offers a computational framework that leverages open source social media data and physical protest data which are available for most countries in the world, thus allowing its approach and results of this study to be scaled up to other social movements in other parts of the world.
2. The research contributes additional empirical evidence that hashtag activism is strongly related to physical protest activity through time, and that a unidirectional relationship between hashtag activism and physical protest may exist. This evidence is grounded in robust statistical analysis and the use of an expansive and timely geo-located data-set.
3. The research demonstrates that hashtag activism and physical protest activity are not only related through time, but also across the spatial dimension. As such, it contributes novel empirical evidence that the temporal relationship between hashtag activism and physical protest activity exists at different spatial resolutions, including the state and county levels.
4. The research provides evidence that deep learning frameworks can be used effectively for the prediction of physical protest activity during social movements based on the raw count of relevant social media posts. Through this research a deep learning model, namely a Long Short Term Memory (LSTM) neural network combined with a traditional Vector Autoregression (VAR) model, was trained and validated for the prediction of physical protest activity for a chosen case-study. For obvious ethical implications, however, this model will not be shared publicly.
5. The research offers the *Mobilization Synergy Index*, a novel statistical index that sheds new insight into the spatial distribution of the temporal relationship between hashtag activism and physical protest. Such an index makes use of simple statistical concepts to explain why some geographical areas exhibit a weaker temporal relationship between hashtag activism and physical protest activity. The index may serve as a tool for local policy makers and activists to understand how their communities mobilize in times of crisis, and thus improve the effectiveness of local political and activism strategies.
6. The research demonstrates that hashtag activism may have real political consequences that go beyond political mobilization for physical protest. As such, it offers novel empirical evidence that the intensity of hashtag activism today can provide insight into the significance of future political action in terms of new legislative responses. Additionally, the research provides statistical evidence that this trend may also be related over space, where state level hashtag activism informs state level legislative responses.

5.2.1. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

Due to the scope of this research, this study comes with a number of limitations that may inform future scientific efforts:

1. **External validity**: due to the analysis of a single case study, the results of this research could only be valid for the specific case of #JusticeForGeorgeFloyd and may not be observed for other social movements in other parts of the world. Future research may address this limitation by implementing the proposed experimental framework to the case of other relevant social movements in other geographical areas.
2. **Spatial resolution**: While this research makes use of a large data-set with considerable geographical coverage, the spatial resolution of this data is still limited. For this reason, the analysis could not be implemented for sub-county level patterns in hashtag activism and physical protest activity (e.g. the neighborhood or precinct level). This limitation is difficult to address, as the availability of sub-city level geo-located social media data is scarce due to a number of ethical considerations (e.g. privacy). Nonetheless, future research could address this limitation by using sources of social media data that come with improved spatial resolution.
3. **Restrictive Legislative data**: Due to the lack of a concerted effort to collect local level (i.e. county-level or city-level) legislative responses to policing, this research could only analyze state-level data relating to legislative action. Future research leveraging on local-level data on legislative responses could greatly improve this limitation and provide more accurate insights into the impact of hashtag activism on legislative responses. With such analysis, a similar index to the *Mobilization Synergy Index* proposed in this research could be developed.

6
Conclusion

Through this research, the political impacts of location-based hashtag activism during social movements were analyzed with the help of two main guiding questions. Let us revisit each one of them, and describe how this research has answered them. Finally, let us discuss some important ethical implications associated with the results of this research.

6.1. Research Questions

1. *What is the relationship between location-based hashtag activism and physical protest activity?*

   The relationship between location-based hashtag activism and physical protest activity can be described as a spatio-temporal, positive power-law correlation, and potentially as a unidirectional relationship, where the highest spikes in hashtag activism would occur one day before the highest spikes in physical protest activity. In other words, when people with similar values, experiences, world-views, and interests, come together as a collective voice on social media, it is likely that they will also come together physically the next day to express their collective voice in public spaces. As we have seen with county-level results, however, there exist geographical areas in which this relationship does not hold, and where one of hashtag activism or physical protest activity is disproportionally intense over time.

2. *What is the relationship between location-based hashtag activism and legislative action?*

   The relationship between location-based hashtag activism and legislative action can also be described as a spatio-temporal, positive power-law correlation, and potentially as a unidirectional relationship, where hashtag activism may influence future legislative action. The temporal nature of this relationship, however, varies greatly by geographical area. The reasons for these differences are outside the scope of this study but could be explored in future research.

6.2. Ethical Considerations

Personal data, code, models, and advanced computation are exciting things for us scientists. In so little time, their increased accessibility has empowered us to study the world in revolutionary ways. In particular, these tools have created a bridge across seemingly distant scientific fields, such as computer science and sociology, and have allowed us to deepen our understanding of complex human behaviours and social structures. Through this research, the analysis and modelling of geo-located social media data has taught us that the way in which we use social media in times of crisis is a strong indicator of where, when, and with what intensity we choose to take the streets and protest against injustice.

While this is a fascinating research outcome it may also raise some ethical concerns because we live in a world where powerful institutions such as law enforcement agencies often harness their power and knowledge of social mechanisms to oppress the most vulnerable and disenfranchised groups. In the United States, research has demonstrated that police use of force is among the leading causes of death for Black and Brown men, and that racially-rooted police violence is on the rise. Similar bias in police practices has been observed in many other countries around the world. Additionally, as paradoxical as it sounds, state-sanctioned violence of this kind is widespread during physical protests, where the very people that demand justice and accountability for people like George Floyd are oppressed further. In this light, 2020 and 2021 were no exceptions, as it was documented during the #JusticeForGeorgeFloyd social movement in the United States, the case study of this research, and as it is being documented during the #ColombiaResiste social movement which is happening as this paragraph is being written.

From this point of view, the social impact of this research may not be as exciting as its scientific impact. While its findings might empower activist groups with tools to better allocate their resources and communicate with local communities more efficiently, they could also be diverted to other less desirable uses. For instance, the models and insights of this research could equip law enforcement agencies with tools to respond to physical protests more rapidly and to use more oppressive tactics against protesters. As scientists, it is our responsibility to conduct and communicate novel scientific research, but we have little scope to ensure that such research is used solely for the betterment of society. Such a task is the role of policy makers and (to a lesser extent) private companies, who do have regulatory and veto power at their disposal.

Fortunately, there are effective ways in which policy makers and social media companies could prevent unethical applications of this research. With regards to law and policy making, regulatory policies can be put in place for law enforcement agencies to respect the personal privacy of individuals. For example, such policies could limit the use of location data from social media by law enforcement agencies for the purpose of tracking demonstrations and ongoing physical protests. With regards to social media companies such as Twitter, usage policies can be put in place to limit access of specific data attributes to law enforcement agencies. For example, just like Twitter has recently increased access to its data for academic purposes, it may limit access to its data for law enforcement purposes (i.e. by not providing location information to law enforcement developers). By creating boundaries to access, the use of data for social good can be ensured and activist groups can safeguard their opportunity to lawfully voice their concerns for more progressive policies through physical protest.

#JusticeForGeorgeFloyd

7

Bibliography

Amenta, E., Caren, N., Chiarello, E., and Su, Y. (2010). "The political consequences of social movements". In: *Annual Review of Sociology* 36, pp. 287–307.

Andersson, M., Krylova, E., and Vähämaa, S. (2008). "Why does the correlation between stock and bond returns vary over time?" In: *Applied Financial Economics* 18.2, pp. 139–151.

Bao, W., Yue, J., and Rao, Y. (2017). "A deep learning framework for financial time series using stacked autoencoders and long-short term memory". In: *PloS one* 12.7, e0180944.

Bastos, M. T., Mercea, D., and Charpentier, A. (2015). "Tents, tweets, and events: The interplay between ongoing protests and social media". In: *Journal of Communication* 65.2, pp. 320–350.

Benesty, J., Chen, J., Huang, Y., and Cohen, I. (2009). "Pearson correlation coefficient". In: *Noise reduction in speech processing*. Springer, pp. 1–4.

Benford, R. D. and Snow, D. A. (2000). "Framing processes and social movements: An overview and assessment". In: *Annual review of sociology* 26.1, pp. 611–639.

Bennett, W. L. and Segerberg, A. (2012). "The logic of connective action: Digital media and the personalization of contentious politics". In: *Information, communication & society* 15.5, pp. 739–768.

Boling, K. S. (2020). "# ShePersisted, Mitch: a memetic critical discourse analysis on an attempted Instagram feminist revolution". In: *Feminist Media Studies* 20.7, pp. 966–982.

Bonilla, Y. and Rosa, J. (2015). "# Ferguson: Digital protest, hashtag ethnography, and the racial politics of social media in the United States". In: *American ethnologist* 42.1, pp. 4–17.

Brantly, A. F. (2019). "From cyberspace to independence square: understanding the impact of social media on physical protest mobilization during Ukraine's Euromaidan revolution". In: *Journal of information technology & politics* 16.4, pp. 360–378.

Bureau, U. C. (2021). *Cartographic Boundary Files - Shapefile*. data retrieved from the U.S. Census API, `https://www.census.gov/geographies/mapping-files/time-series/geo/carto-boundary-file.html`.

Burghart, D. B. (2020). *Fatal Encounters*. data retrieved from Fatal Encounters, `https://fatalencounters.org/`.

Burns, A. and Eltham, B. (2009). "Twitter free Iran: An evaluation of Twitter's role in public diplomacy and information operations in Iran's 2009 election crisis". In:

Carey, S. C. (2006). "The dynamic relationship between protest and repression". In: *Political Research Quarterly* 59.1, pp. 1–11.

Carney, N. (2016). "All lives matter, but so does race: Black lives matter and the evolving role of social media". In: *Humanity & Society* 40.2, pp. 180–199.

Curtis, L., Edwards, C., Fraser, K. L., Gudelsky, S., Holmquist, J., Thornton, K., and Sweetser, K. D. (2010). "Adoption of social media for public relations by nonprofit organizations". In: *Public relations review* 36.1, pp. 90–92.

D. Clark, M. (2020). "DRAG THEM: A brief etymology of so-called "cancel culture"". In: *Communication and the Public* 5.3-4, pp. 88–92.

De Smith, M. J., Goodchild, M. F., and Longley, P. (2007). *Geospatial analysis: a comprehensive guide to principles, techniques and software tools*. Troubador publishing ltd.

Della Porta, D. and Mattoni, A. (2015). “Social movements”. In: *The international encyclopedia of political communication*, pp. 1–8.

Devito, M. A., Walker, A. M., Birnholtz, J., Ringland, K., Macapagal, K., Kraus, A., Munson, S., Liang, C., and Saksono, H. (2019). “Social technologies for digital wellbeing among marginalized communities”. In: *Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing*, pp. 449–454.

Ding, J., Tarokh, V., and Yang, Y. (2017). “Bridging AIC and BIC: a new criterion for autoregression”. In: *IEEE Transactions on Information Theory* 64.6, pp. 4024–4043.

Draper, N. R. and Smith, H. (1998). *Applied regression analysis*. Vol. 326. John Wiley & Sons.

Earl, J. and Kimport, K. (2011). *Digitally enabled social change: Activism in the internet age*. Mit Press.

Edwards, F., Lee, H., and Esposito, M. (2019). “Risk of being killed by police use of force in the United States by age, race–ethnicity, and sex”. In: *Proceedings of the National Academy of Sciences* 116.34, pp. 16793–16798.

Eisinger, P. K. (1973). “The conditions of protest behavior in American cities”. In: *The American political science review* 67.1, pp. 11–28.

Enos, R. D., Kaufman, A. R., and Sands, M. L. (2019). “Can violent protest change local policy support? Evidence from the aftermath of the 1992 Los Angeles riot”. In: *American Political Science Review* 113.4, pp. 1012–1028.

Fahrmeir, L., Kneib, T., Lang, S., and Marx, B. (2007). *Regression*. Springer.

Freelon, D., McIlwain, C., and Clark, M. (2018). “Quantifying the power and consequences of social media protest”. In: *New Media & Society* 20.3, pp. 990–1011.

Fuchs, C. and Sandoval, M. (2014). “Introduction: Critique, social media and the information society in the age of capitalist crisis”. In: Routledge.

George, J. J. and Leidner, D. E. (2018). “Why marginalized groups struggle to leverage digital activism”. In:

Gerbaudo, P. and Treré, E. (2015). *In search of the ‘we’of social media activism: introduction to the special issue on social media and protest identities*.

7

Granger, C. W. (1969). “Investigating causal relations by econometric models and cross-spectral methods”. In: *Econometrica: journal of the Econometric Society*, pp. 424–438.

Greff, K., Srivastava, R. K., Koutnık, J., Steunebrink, B. R., and Schmidhuber, J. (2016). “LSTM: A search space odyssey”. In: *IEEE transactions on neural networks and learning systems* 28.10, pp. 2222–2232.

Grover, P., Kar, A. K., Dwivedi, Y. K., and Janssen, M. (2019). “Polarization and acculturation in US Election 2016 outcomes–Can twitter analytics predict changes in voting preferences”. In: *Technological Forecasting and Social Change* 145, pp. 438–460.

Heaney, M. T. (2020). “Protest at the center of American politics”. In: *Journal of International Affairs* 73.2, pp. 195–208.

Hochreiter, S. and Schmidhuber, J. (1997). “Long short-term memory”. In: *Neural computation* 9.8, pp. 1735–1780.

Howard, P. N., Agarwal, S. D., and Hussain, M. M. (2011). “When do states disconnect their digital networks? Regime responses to the political uses of social media”. In: *The Communication Review* 14.3, pp. 216–232.

Ince, J., Rojas, F., and Davis, C. A. (2017). “The social media response to Black Lives Matter: How Twitter users interact with Black Lives Matter through hashtag use”. In: *Ethnic and racial studies* 40.11, pp. 1814–1830.

Jackson, S. J., Bailey, M., and Welles, B. F. (2020). *# HashtagActivism: Networks of race and gender justice*. MIT Press.

Jacobs, L. A., Kim, M. E., Whitfield, D. L., Gartner, R. E., Panichelli, M., Kattari, S. K., Downey, M. M., McQueen, S. S., and Mountz, S. E. (2021). “Defund the police: Moving towards an anti-carceral social work”. In: *Journal of Progressive Human Services* 32.1, pp. 37–62.

Johnson, J. L. and Callahan, C. (2013). “Minority cultures and social media: Magnifying Garifuna”. In: *Journal of Intercultural Communication Research* 42.4, pp. 319–339.

Juris, J. S. (2012). “Reflections on# Occupy Everywhere: Social media, public space, and emerging logics of aggregation”. In: *American ethnologist* 39.2, pp. 259–279.

Karduni, A. and Sauda, E. (2020). “Anatomy of a protest: spatial information, social media, and urban space”. In: *Social Media+ Society* 6.1, p. 2056305119897320.

Kent, M. L. (2013). “Using social media dialogically: Public relations role in reviving democracy”. In: *Public relations review* 39.4, pp. 337–345.

Khondker, H. H. (2011). “Role of the new media in the Arab Spring”. In: *Globalizations* 8.5, pp. 675–679.

Lăzăroiu, G. *et al.* (2018). “Participation environments, collective identities, and online political behavior: The role of media technologies for social protest campaigns”. In: *Geopolitics, History, and International Relations* 10.2, pp. 58–63.

Lebron, C. J. (2017). *The making of black lives matter: A brief history of an idea*. Oxford University Press.

Li, M., Turki, N., Izaguirre, C. R., DeMahy, C., Thibodeaux, B. L., and Gage, T. (2021). “Twitter as a tool for social movement: An analysis of feminist activism on social media communities”. In: *Journal of community psychology* 49.3, pp. 854–868.

Li, X., Peng, L., Yao, X., Cui, S., Hu, Y., You, C., and Chi, T. (2017). “Long short-term memory neural network for air pollutant concentration predictions: Method development and evaluation”. In: *Environmental pollution* 231, pp. 997–1004.

Lipsky, M. (1968). “Protest as a political resource”. In: *The American Political Science Review* 62.4, pp. 1144–1158.

Ma, X., Tao, Z., Wang, Y., Yu, H., and Wang, Y. (2015). “Long short-term memory neural network for traffic speed prediction using remote microwave sensor data”. In: *Transportation Research Part C: Emerging Technologies* 54, pp. 187–197.

Manikonda, L., Beigi, G., Kambhampati, S., and Liu, H. (2018). “# metoo through the lens of social media”. In: *International conference on social computing, behavioral-cultural modeling and prediction and behavior representation in modeling and simulation*. Springer, pp. 104–110.

Matter, B. L. (n.d.). *About: Black Lives Matter*. https://blacklivesmatter.com/about/. Accessed: 2021-05-25.

McCulloch, W. S. and Pitts, W. (1943). “A logical calculus of the ideas immanent in nervous activity”. In: *The bulletin of mathematical biophysics* 5.4, pp. 115–133.

McKinney, W., Perktold, J., and Seabold, S. (2011). “Time series analysis in Python with statsmodels”. In: *Jarrodmillman Com*, pp. 96–102.

NCSL (2020). *Legislative Responses for Policing-State Bill Tracking Database*. data retrieved from the National Conference of State Legislatures, https://www.ncsl.org/research/civil-and-criminal-justice/legislative-responses-for-policing.aspx#Database.

Nemer, D. (2016). “Online favela: The use of social media by the marginalized in Brazil”. In: *Information technology for development* 22.3, pp. 364–379.

Ortiz, J. *et al.* (2019). “Giving voice to the voiceless: The use of digital technologies by marginalized groups”. In:

Ouassini, A. (2021). “We are all amina filali: social media, civil society, and rape legislation reform in Morocco”. In: *Women & Criminal Justice* 31.1, pp. 77–82.

Priniski, J. H., Mokhberian, N., Harandizadeh, B., Morstatter, F., Lerman, K., Lu, H., and Brantingham, P. J. (2021). “Mapping Moral Valence of Tweets Following the Killing of George Floyd”. In: *arXiv preprint arXiv:2104.09578*.

Pydata (2021). *pandas*. code retrieved from Pydata, https://pandas.pydata.org/.

PyPi (2021). *Beautifulsoup 4*. code retrieved from PyPi, https://pypi.org/project/beautifulsoup4/.

Raleigh, C., Linke, A., Hegre, H., and Karlsen, J. (2010). “Introducing ACLED: an armed conflict location and event dataset: special data feature”. In: *Journal of peace research* 47.5, pp. 651–660.

Ray, R., Brown, M., Fraistat, N., and Summers, E. (2017). “Ferguson and the death of Michael Brown on Twitter:# BlackLivesMatter,# TCOT, and the evolution of collective identities”. In: *Ethnic and Racial Studies* 40.11, pp. 1797–1813.

Rucht, D. (2007). “The spread of protest politics”. In: *The Oxford handbook of political behavior*.

Schussman, A. and Soule, S. A. (2005). “Process and protest: Accounting for individual protest participation”. In: *Social forces* 84.2, pp. 1083–1108.

Seabold, S. and Perktold, J. (2010). “statsmodels: Econometric and statistical modeling with python”. In: *9th Python in Science Conference*.

Shakil, S., Lee, C.-H., and Keilholz, S. D. (2016). “Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states”. In: *Neuroimage* 133, pp. 111–128.

Shen, C. (2015). “Analysis of detrended time-lagged cross-correlation between two nonstationary time series”. In: *Physics Letters A* 379.7, pp. 680–687.

Sherstinsky, A. (2020). “Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network”. In: *Physica D: Nonlinear Phenomena* 404, p. 132306.

Shirky, C. (2011). “The political power of social media: Technology, the public sphere, and political change”. In: *Foreign affairs*, pp. 28–41.

Tan, L., Ponnam, S., Gillham, P., Edwards, B., and Johnson, E. (2013). “Analyzing the impact of social media on social movements: a computational study on Twitter and the Occupy Wall Street movement”. In: *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 1259–1266.

The New York Times, T. (2020). *How George FLoyd Died and What Happened Next.*

— (2021). *George Floyd Protests: A Timeline.*

Tombleson, B. and Wolf, K. (2017). “Rethinking the circuit of culture: How participatory culture has transformed cross-cultural communication”. In: *Public Relations Review* 43.1, pp. 14–25.

Trevisan, F. (2019). “Using the internet to mobilize marginalized groups: people with disabilities and digital campaign strategies in the 2016 US presidential election”. In: *International Journal of Communication* 13, p. 20.

Tufekci, Z. and Wilson, C. (2012). “Social media and the decision to participate in political protest: Observations from Tahrir Square”. In: *Journal of communication* 62.2, pp. 363–379.

Twitter (2021). *Twitter API v2: Early Access.* data retrieved from the Twitter API v2, https://developer.twitter.com/en/docs/twitter-api/tweets/search/api-reference/get-tweets-search-all.

Uber (2021). *H3: A Hexagonal Hierarchical Geospatial Indexing System.* data retrieved from Uber, https://github.com/uber/h3.

Ver Steeg, G. and Galstyan, A. (2012). “Information transfer in social media”. In: *Proceedings of the 21st international conference on World Wide Web*, pp. 509–518.

Vomfell, L., Härdle, W. K., and Lessmann, S. (2018). “Improving crime count forecasts using Twitter and taxi data”. In: *Decision Support Systems* 113, pp. 73–85.

Wei, W. W. (2006). “Time series analysis”. In: *The Oxford Handbook of Quantitative Methods in Psychology: Vol. 2.*

Weldon, L. *et al.* (2011). *When protest makes policy: How social movements represent disadvantaged groups.* University of Michigan Press.

Wilson, S. L. (2017). “Detecting mass protest through social media”. In: *The Journal of Social Media in Society* 6.2, pp. 5–25.

Xiong, Y., Cho, M., and Boatwright, B. (2019). “Hashtag activism and message frames among social movement organizations: Semantic network analysis and thematic analysis of Twitter during the# MeToo movement”. In: *Public relations review* 45.1, pp. 10–23.

Yegnanarayana, B. (2009). *Artificial neural networks.* PHI Learning Pvt. Ltd.

Zhang, D. and Fan, G.-Z. (2019). “Regional spillover and rising connectedness in China’s urban housing prices”. In: *Regional Studies* 53.6, pp. 861–873.

Zhao, L., Zeng, W., Shi, Y., Nie, W., and Yang, J. (2020). “Dynamic visual cortical connectivity analysis based on functional magnetic resonance imaging”. In: *Brain and behavior* 10.7, e01698.

Zheng, J., Xu, C., Zhang, Z., and Li, X. (2017). “Electric load forecasting in smart grids using long-short-term-memory based recurrent neural network”. In: *2017 51st Annual Conference on Information Sciences and Systems (CISS).* IEEE, pp. 1–6.

Zhuravskaya, E., Petrova, M., and Enikolopov, R. (2020). “Political effects of the internet and social media”. In: *Annual Review of Economics* 12, pp. 415–438.

Zivot, E. and Wang, J. (2006). “Vector autoregressive models for multivariate time series”. In: *Modeling Financial Time Series with S-Plus®*, pp. 385–429.

8

Appendix

Example tweet object


```
  id: "1287220607411781633"
  created_at: "2020-07-26T02:58:02.000Z"
▼ public_metrics:
    retweet_count: 2
    reply_count: 1
    like_count: 5
    quote_count: 0
▼ entities:
  ▶ urls: [] 1 item
  ▶ mentions: [] 1 item
  ▶ hashtags: [] 2 items
  ▶ annotations: [] 3 items
▼ geo:
    place_id: "00f751614d8ce37b"
  ▼ coordinates:
      type: "Point"
    ▼ coordinates: [] 2 items
        0: -77.4371
        1: 37.5373
  author_id: "17838014"
  text: "HAPPENING NOW: Richmond, Va • Protestors marching in solidarity with Portland and against police violence • This video was taken by Richmon
   d photographer @carlyIsands ▫▪▫▪ • • #blacklivesmatter #solidarity… https://t.co/YXSDIwwabG"
  conversation_id: "1287220607411781633"
```


Figure 8.1: Example of data structure for one tweet object, as provided by the Twitter API

ROLLING WINDOW CORRELATION AT THE U.S. STATE LEVEL

state	twt_pvalue	protst_pvalue	twt_statistic	protst_statistic	tweet_normresult	protest_normresult	corr_p_valu	parson_r	r2
AK	2.175×10^{-14}	1.221×10^{-28}	62.918	128.546	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.748×10^{-12}	0.627	0.393
AL	3.463×10^{-27}	1.789×10^{-41}	121.855	187.649	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.278×10^{-28}	0.726	0.528
AR	4.303×10^{-28}	7.729×10^{-37}	126.026	166.301	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.619×10^{-18}	0.643	0.413
AZ	8.982×10^{-24}	7.747×10^{-32}	106.134	143.271	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.284×10^{-28}	0.743	0.552
CA	1.648×10^{-32}	2.011×10^{-32}	146.367	145.968	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	4.343×10^{-70}	0.923	0.853
CO	1.049×10^{-21}	8.085×10^{-25}	96.613	110.949	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.216×10^{-28}	0.768	0.589
CT	4.955×10^{-31}	1.692×10^{-46}	139.560	210.786	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.784×10^{-28}	0.726	0.527
DC	7.965×10^{-36}	2.640×10^{-32}	161.636	145.424	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.805×10^{-6}	0.391	0.153
DE	1.333×10^{-23}	3.236×10^{-23}	105.344	103.570	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	0.000	0.329	0.108
FL	1.339×10^{-31}	9.501×10^{-28}	142.176	124.442	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.324×10^{-40}	0.811	0.657
GA	8.047×10^{-26}	1.525×10^{-34}	115.564	155.732	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	4.341×10^{-46}	0.839	0.705
HI	1.402×10^{-23}	7.280×10^{-39}	105.243	175.631	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.475×10^{-26}	0.744	0.554
IA	9.836×10^{-25}	3.298×10^{-36}	110.557	163.399	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	6.524×10^{-50}	0.868	0.754
ID	5.406×10^{-35}	1.960×10^{-55}	157.806	251.939	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.635×10^{-8}	0.447	0.200
IL	1.671×10^{-28}	1.669×10^{-32}	127.918	146.341	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.410×10^{-38}	0.805	0.648
IN	1.068×10^{-26}	1.712×10^{-32}	119.603	146.291	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.485×10^{-41}	0.824	0.679
KS	4.254×10^{-30}	3.869×10^{-39}	135.260	176.896	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.605×10^{-19}	0.663	0.439
KY	6.943×10^{-27}	6.775×10^{-30}	120.464	134.329	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.756×10^{-27}	0.715	0.511
LA	9.173×10^{-31}	1.484×10^{-28}	138.328	128.155	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.620×10^{-28}	0.725	0.525
MA	2.991×10^{-33}	1.298×10^{-32}	149.779	146.844	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.551×10^{-33}	0.764	0.583
MD	1.583×10^{-26}	1.097×10^{-38}	118.815	174.810	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.346×10^{-24}	0.692	0.479
ME	2.047×10^{-29}	8.550×10^{-40}	132.117	179.915	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	9.207×10^{-24}	0.709	0.503
MI	1.892×10^{-33}	2.880×10^{-35}	150.695	159.065	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.258×10^{-44}	0.834	0.695
MN	7.860×10^{-33}	1.527×10^{-35}	147.847	160.334	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.881×10^{-39}	0.823	0.678
MO	8.334×10^{-29}	2.526×10^{-26}	129.309	117.881	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.367×10^{-34}	0.777	0.604
MS	1.935×10^{-25}	5.902×10^{-38}	113.810	171.446	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.259×10^{-25}	0.711	0.505
MT	3.272×10^{-22}	4.048×10^{-26}	98.943	116.938	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	3.773×10^{-18}	0.629	0.396
NC	2.505×10^{-26}	7.105×10^{-32}	117.898	143.444	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	9.411×10^{-43}	0.836	0.699
ND	5.278×10^{-26}	1.437×10^{-7}	116.408	31.511	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.904×10^{-8}	0.480	0.230
NE	6.901×10^{-21}	4.506×10^{-26}	92.845	116.723	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.352×10^{-19}	0.763	0.583
NH	8.647×10^{-12}	6.520×10^{-22}	50.948	97.564	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.279×10^{-13}	0.689	0.475
NJ	4.310×10^{-28}	8.867×10^{-36}	126.023	161.422	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	8.679×10^{-29}	0.757	0.573
NM	7.303×10^{-21}	8.418×10^{-21}	92.732	92.448	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.573×10^{-24}	0.742	0.551
NV	1.043×10^{-22}	3.464×10^{-37}	101.229	167.906	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.177×10^{-9}	0.505	0.255
NY	1.239×10^{-28}	2.011×10^{-29}	128.516	132.152	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.657×10^{-51}	0.869	0.755
OH	7.792×10^{-29}	1.031×10^{-39}	129.444	179.541	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	5.538×10^{-37}	0.795	0.632
OK	3.723×10^{-25}	8.620×10^{-42}	112.500	189.109	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	4.196×10^{-11}	0.524	0.275
OR	2.838×10^{-25}	3.653×10^{-28}	113.043	126.354	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.326×10^{-39}	0.807	0.651
PA	5.372×10^{-29}	6.829×10^{-34}	130.188	152.733	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.515×10^{-37}	0.797	0.635
RI	5.670×10^{-24}	2.664×10^{-20}	107.054	90.144	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.775×10^{-13}	0.598	0.358
SC	1.822×10^{-34}	1.099×10^{-30}	155.376	137.966	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	2.321×10^{-39}	0.811	0.658
SD	3.430×10^{-28}	4.382×10^{-37}	126.480	167.436	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	3.048×10^{-20}	0.715	0.512
TN	2.230×10^{-26}	2.066×10^{-28}	118.130	127.493	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.055×10^{-42}	0.832	0.691
TX	4.755×10^{-34}	1.322×10^{-34}	153.457	156.018	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	3.736×10^{-44}	0.828	0.686
USA	3.252×10^{-30}	9.731×10^{-32}	135.797	142.815	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	8.723×10^{-60}	0.894	0.799
UT	2.754×10^{-29}	2.537×10^{-29}	131.524	131.688	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.152×10^{-23}	0.715	0.511
VA	3.881×10^{-26}	1.605×10^{-29}	117.022	132.604	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	9.503×10^{-49}	0.858	0.736
VT	4.018×10^{-19}	8.432×10^{-25}	84.717	110.865	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	1.122×10^{-13}	0.614	0.377
WA	8.205×10^{-27}	7.306×10^{-35}	120.130	157.204	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.642×10^{-44}	0.835	0.697
WI	1.102×10^{-22}	1.105×10^{-40}	101.120	184.008	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	7.982×10^{-30}	0.731	0.534
WV	2.790×10^{-28}	5.838×10^{-48}	126.893	217.519	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	3.088×10^{-32}	0.784	0.615
WY	2.014×10^{-11}	1.214×10^{-28}	49.256	128.557	reject NH: tweet data not gaussian	reject NH: protest data not gaussian	9.384×10^{-12}	0.575	0.331

Table 8.1: State-level normality tests and rolling window correlation for tweet count and protest count

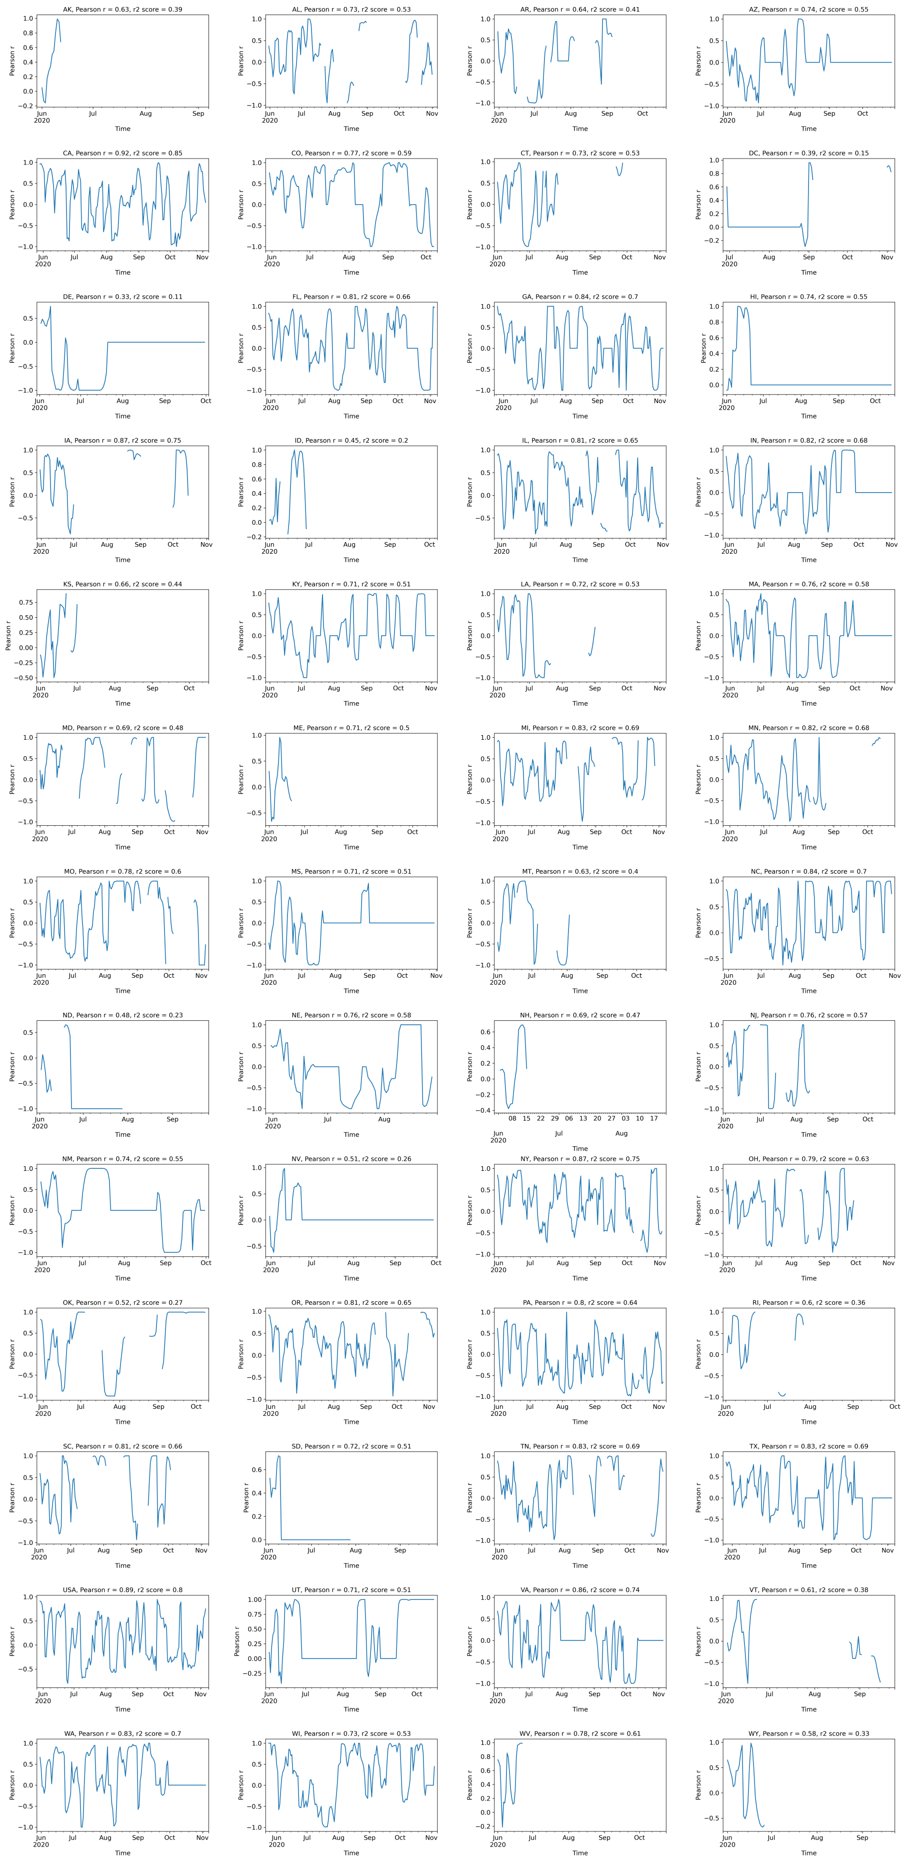

Figure 8.2: Seven day rolling window correlation for each of the U.S.'s 50 states using the #JusticeForGeorgeFloyd as a case study.

TIME-LAGGED CROSS-CORRELATION ANALYSIS AT THE U.S. STATE LEVEL

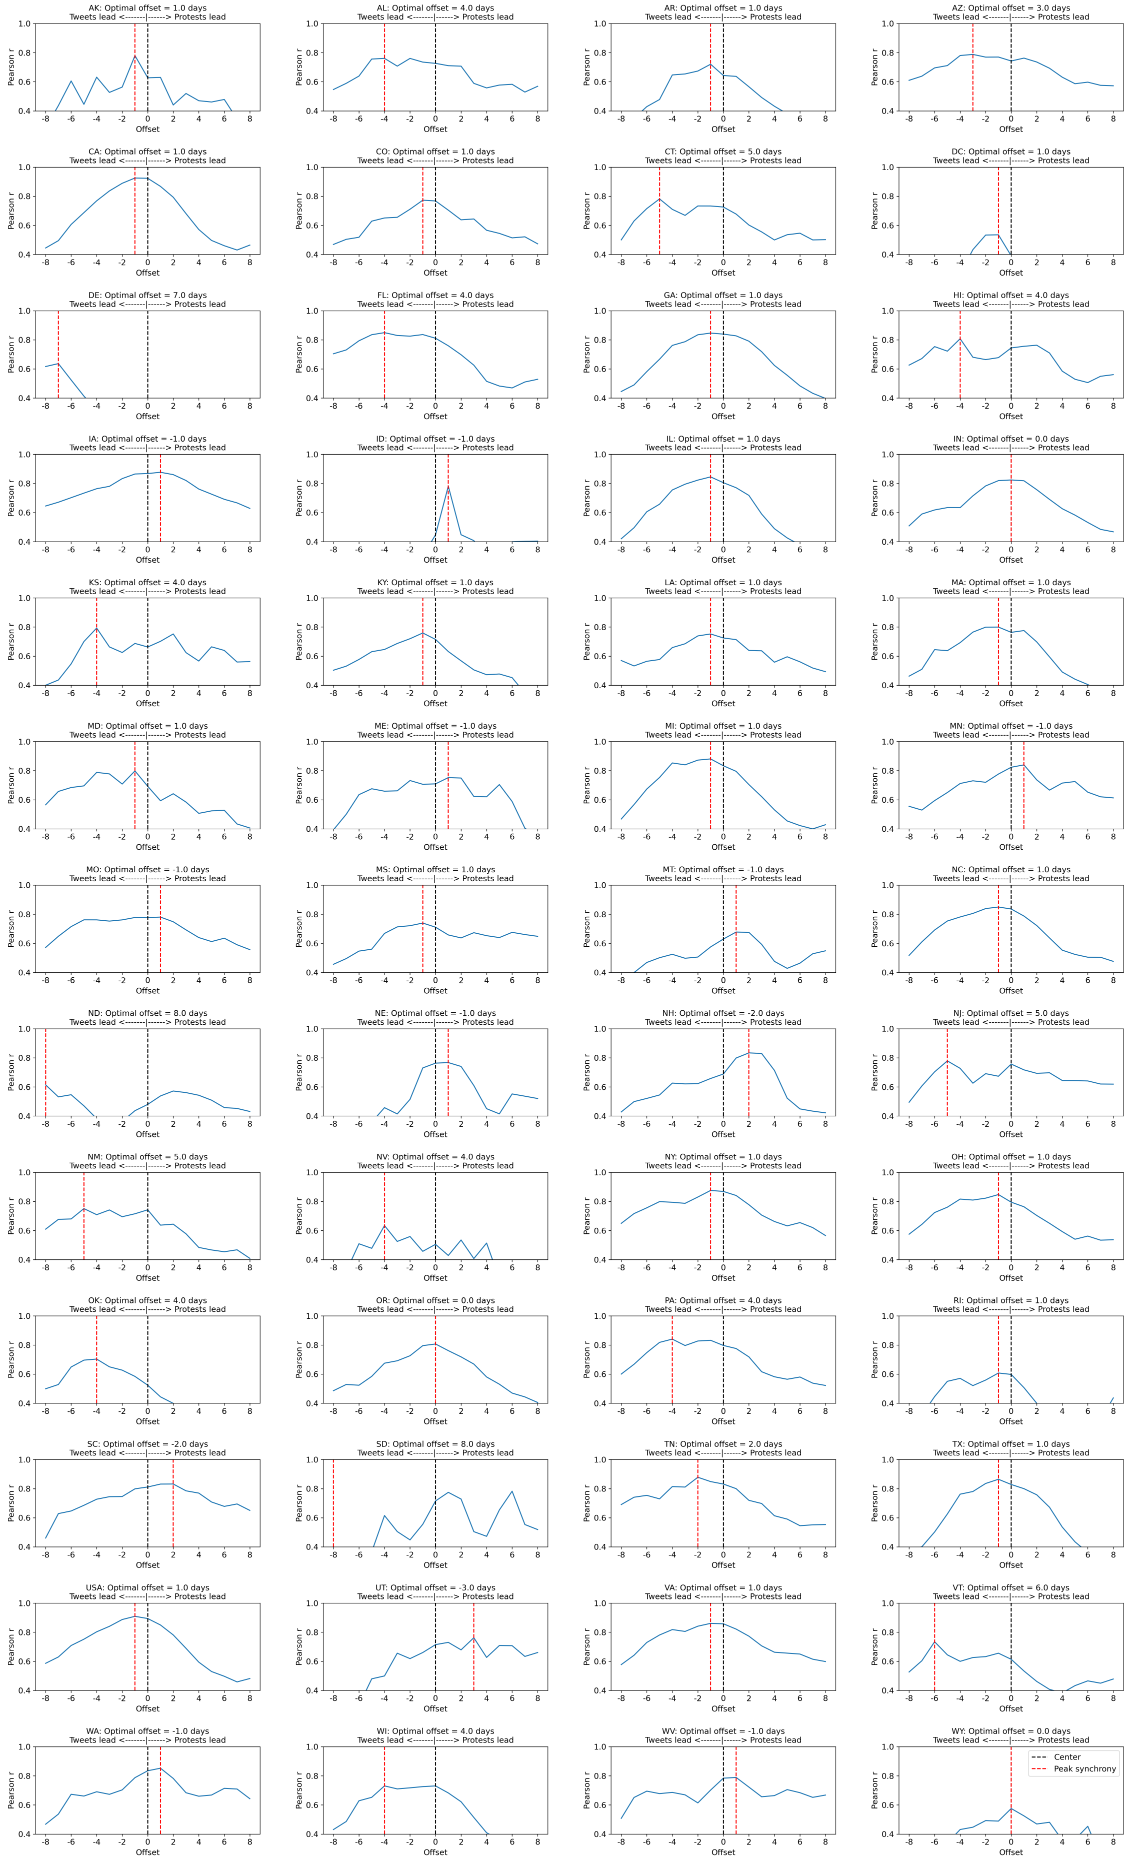

Figure 8.3: Time-lagged cross-correlation analysis for each of the U.S.'s 50 states using the #JusticeForGeorgeFloyd as a case study.

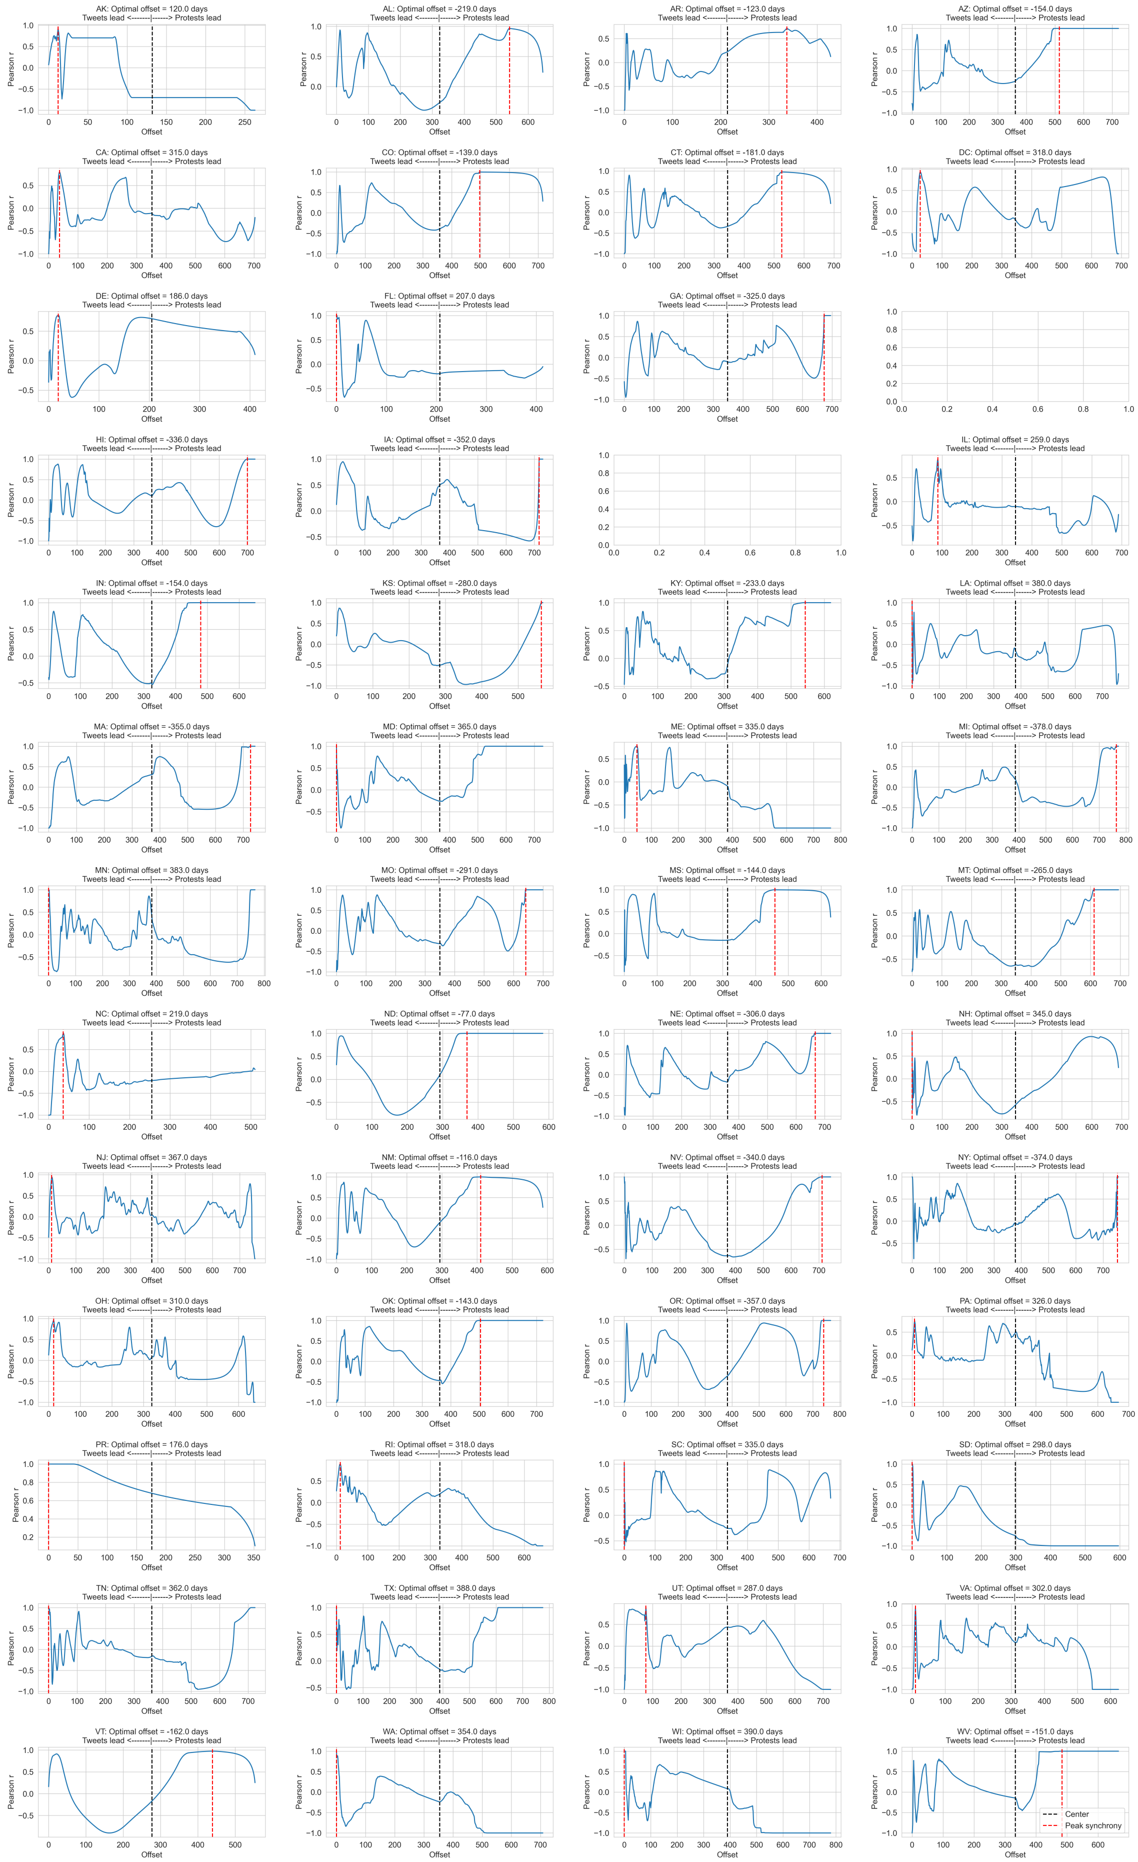

Figure 8.4: Time-lagged cross-correlation analysis (tweets and legislation) for each of the U.S.'s 50 states using the #JusticeForGeorgeFloyd as a case study.

8

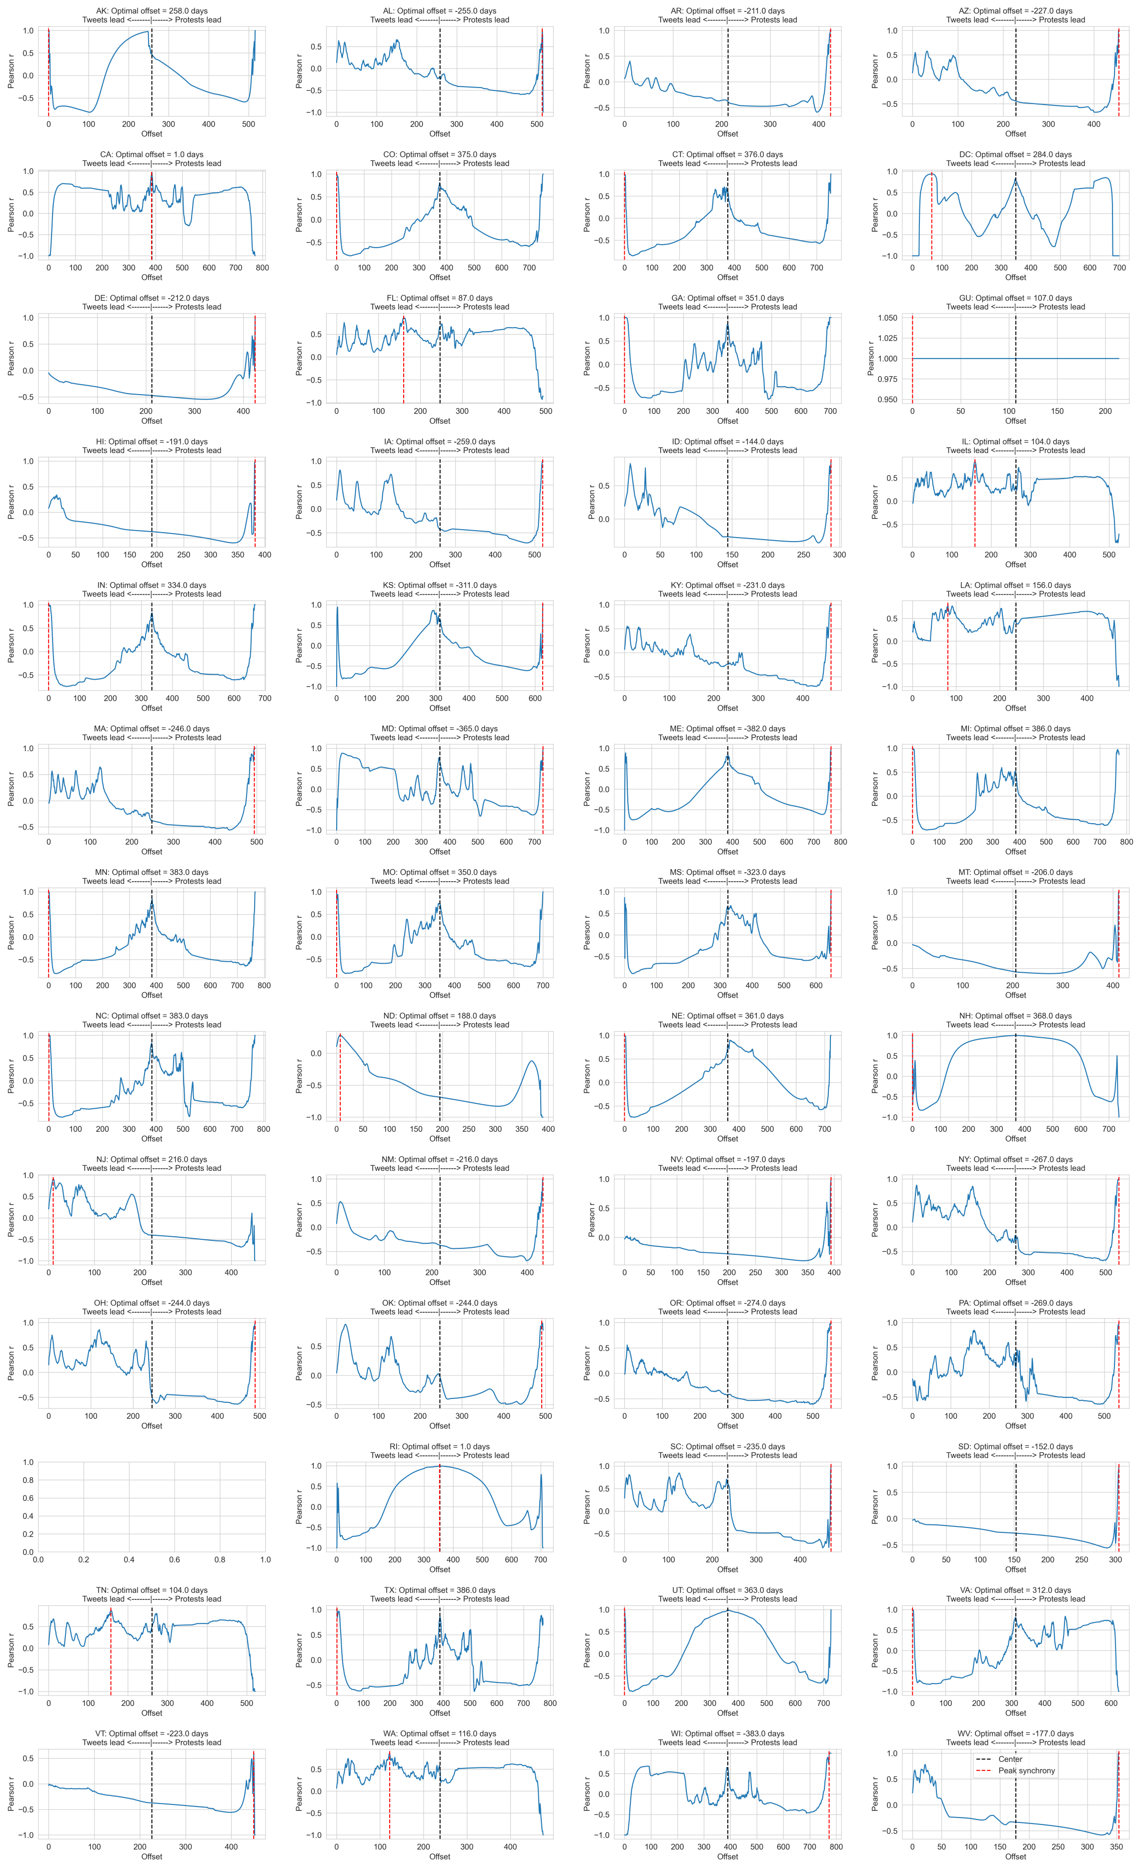

Figure 8.5: Time-lagged cross-correlation analysis (physical protests and legislation) for each of the U.S.'s 50 states using the #JusticeForGeorgeFloyd as a case study.

8

GRANGER CAUSALITY ANALYSIS AT THE U.S. STATE LEVEL

granger_relationship	p_value	state
tweet_count ->protest_count	0.0	AK
tweet_count ->protest_count	0.0	AL
tweet_count ->protest_count	0.0	AR
tweet_count ->protest_count	0.0	AZ
tweet_count ->protest_count	0.0	CA
tweet_count ->protest_count	0.0	CO
tweet_count ->protest_count	0.0	CT
tweet_count ->protest_count	0.0	DC
tweet_count ->protest_count	0.4189	DE
tweet_count ->protest_count	0.0	FL
tweet_count ->protest_count	0.0	GA
tweet_count ->protest_count	0.0	HI
tweet_count ->protest_count	0.0	IA
tweet_count ->protest_count	0.0247	ID
tweet_count ->protest_count	0.0	IL
tweet_count ->protest_count	0.0	IN
tweet_count ->protest_count	0.0	KS
tweet_count ->protest_count	0.0	KY
tweet_count ->protest_count	0.0	LA
tweet_count ->protest_count	0.0	MA
tweet_count ->protest_count	0.0	MD
tweet_count ->protest_count	0.0	ME
tweet_count ->protest_count	0.0	MI
tweet_count ->protest_count	0.0	MN
tweet_count ->protest_count	0.0	MO
tweet_count ->protest_count	0.0	MS
tweet_count ->protest_count	0.0071	MT
tweet_count ->protest_count	0.0	NC
tweet_count ->protest_count	0.2303	ND
tweet_count ->protest_count	0.0	NE
tweet_count ->protest_count	0.0061	NH
tweet_count ->protest_count	0.0	NJ
tweet_count ->protest_count	0.0	NM
tweet_count ->protest_count	0.0	NV
tweet_count ->protest_count	0.0	NY
tweet_count ->protest_count	0.0	OH
tweet_count ->protest_count	0.0	OK
tweet_count ->protest_count	0.0	OR
tweet_count ->protest_count	0.0	PA
tweet_count ->protest_count	0.0062	RI
tweet_count ->protest_count	0.0	SC
tweet_count ->protest_count	0.0	SD
tweet_count ->protest_count	0.0	TN
tweet_count ->protest_count	0.0	TX
tweet_count ->protest_count	0.0	USA
tweet_count ->protest_count	0.0	UT
tweet_count ->protest_count	0.0	VA
tweet_count ->protest_count	0.0	VT
tweet_count ->protest_count	0.0	WA
tweet_count ->protest_count	0.0	WI
tweet_count ->protest_count	0.0	WV
tweet_count ->protest_count	0.1253	WY

Table 8.2: National and State-level results of Granger causality analysis for the hypothesis that tweet count Granger causes protest count.

granger_relationship	p_value	state
protest_count ->tweet_count	0.0167	AK
protest_count ->tweet_count	0.1941	AL
protest_count ->tweet_count	0.138	AR
protest_count ->tweet_count	0.0008	AZ
protest_count ->tweet_count	0.1678	CA
protest_count ->tweet_count	0.1262	CO
protest_count ->tweet_count	0.2926	CT
protest_count ->tweet_count	0.9055	DC
protest_count ->tweet_count	0.0673	DE
protest_count ->tweet_count	0.5815	FL
protest_count ->tweet_count	0.0328	GA
protest_count ->tweet_count	0.0027	HI
protest_count ->tweet_count	0.0002	IA
protest_count ->tweet_count	0.0	ID
protest_count ->tweet_count	0.1795	IL
protest_count ->tweet_count	0.0213	IN
protest_count ->tweet_count	0.0001	KS
protest_count ->tweet_count	0.1958	KY
protest_count ->tweet_count	0.08	LA
protest_count ->tweet_count	0.0002	MA
protest_count ->tweet_count	0.0841	MD
protest_count ->tweet_count	0.0	ME
protest_count ->tweet_count	0.4056	MI
protest_count ->tweet_count	0.0	MN
protest_count ->tweet_count	0.0041	MO
protest_count ->tweet_count	0.0146	MS
protest_count ->tweet_count	0.0	MT
protest_count ->tweet_count	0.0531	NC
protest_count ->tweet_count	0.0001	ND
protest_count ->tweet_count	0.0116	NE
protest_count ->tweet_count	0.0	NH
protest_count ->tweet_count	0.8118	NJ
protest_count ->tweet_count	0.0	NM
protest_count ->tweet_count	0.6764	NV
protest_count ->tweet_count	0.3411	NY
protest_count ->tweet_count	0.6884	OH
protest_count ->tweet_count	0.0259	OK
protest_count ->tweet_count	0.7193	OR
protest_count ->tweet_count	0.2883	PA
protest_count ->tweet_count	0.6172	RI
protest_count ->tweet_count	0.0	SC
protest_count ->tweet_count	0.0	SD
protest_count ->tweet_count	0.3117	TN
protest_count ->tweet_count	0.2259	TX
protest_count ->tweet_count	0.6817	USA
protest_count ->tweet_count	0.0006	UT
protest_count ->tweet_count	0.3075	VA
protest_count ->tweet_count	0.375	VT
protest_count ->tweet_count	0.0	WA
protest_count ->tweet_count	0.1443	WI
protest_count ->tweet_count	0.0108	WV
protest_count ->tweet_count	0.3425	WY

Table 8.3: National and State-level results of Granger causality analysis for the hypothesis that protest count Granger causes tweet count.

CO-INTEGRATION TESTS AT THE U.S. STATE LEVEL

variable	**test_statistic**	**C(95%)**	**signif**	**state**
tweet_count	24.64	12.3212	True	AK
protest_count	6.13	4.1296	True	AK
tweet_count	40.58	12.3212	True	AL
protest_count	6.68	4.1296	True	AL
tweet_count	43.85	12.3212	True	AR
protest_count	9.03	4.1296	True	AR
tweet_count	34.59	12.3212	True	AZ
protest_count	8.23	4.1296	True	AZ
tweet_count	55.18	12.3212	True	CA
protest_count	11.65	4.1296	True	CA
tweet_count	39.47	12.3212	True	CO
protest_count	13.79	4.1296	True	CO
tweet_count	42.49	12.3212	True	CT
protest_count	17.84	4.1296	True	CT
tweet_count	3.14	12.3212	False	DC
protest_count	0.74	4.1296	False	DC
tweet_count	426.93	12.3212	True	DE
protest_count	4.2	4.1296	True	DE
tweet_count	34.42	12.3212	True	FL
protest_count	13.12	4.1296	True	FL
tweet_count	31.55	12.3212	True	GA
protest_count	3.72	4.1296	False	GA
tweet_count	16.71	12.3212	True	HI
protest_count	0.54	4.1296	False	HI
tweet_count	36.54	12.3212	True	IA
protest_count	7.84	4.1296	True	IA
tweet_count	33.97	12.3212	True	ID
protest_count	4.27	4.1296	True	ID
tweet_count	51.4	12.3212	True	IL
protest_count	18.44	4.1296	True	IL
tweet_count	41.22	12.3212	True	IN
protest_count	14.06	4.1296	True	IN
tweet_count	32.94	12.3212	True	KS
protest_count	13.2	4.1296	True	KS
tweet_count	22.12	12.3212	True	KY
protest_count	5.61	4.1296	True	KY
tweet_count	28.92	12.3212	True	LA
protest_count	7.15	4.1296	True	LA
tweet_count	75.45	12.3212	True	MA
protest_count	28.19	4.1296	True	MA
tweet_count	23.49	12.3212	True	MD
protest_count	5.98	4.1296	True	MD
tweet_count	30.2	12.3212	True	ME
protest_count	3.14	4.1296	False	ME
tweet_count	43.93	12.3212	True	MI
protest_count	13.6	4.1296	True	MI
tweet_count	42.9	12.3212	True	MN
protest_count	8.98	4.1296	True	MN
tweet_count	61.22	12.3212	True	MO
protest_count	18.05	4.1296	True	MO
tweet_count	26.9	12.3212	True	MS
protest_count	9.74	4.1296	True	MS
tweet_count	72.16	12.3212	True	MT
protest_count	6.34	4.1296	True	MT
tweet_count	59.98	12.3212	True	NC
protest_count	7.97	4.1296	True	NC
tweet_count	1.42	12.3212	False	ND
protest_count	0.06	4.1296	False	ND
tweet_count	46.16	12.3212	True	NE
protest_count	3.31	4.1296	False	NE
tweet_count	438.89	12.3212	True	NH
protest_count	7.15	4.1296	True	NH
tweet_count	29.02	12.3212	True	NJ
protest_count	11.81	4.1296	True	NJ
tweet_count	26.81	12.3212	True	NM
protest_count	4.81	4.1296	True	NM
tweet_count	46.75	12.3212	True	NV
protest_count	20.7	4.1296	True	NV
tweet_count	95.21	12.3212	True	NY
protest_count	16.91	4.1296	True	NY
tweet_count	68.68	12.3212	True	OH
protest_count	14.33	4.1296	True	OH
tweet_count	40.75	12.3212	True	OK
protest_count	14.24	4.1296	True	OK
tweet_count	37.08	12.3212	True	OR
protest_count	12.08	4.1296	True	OR
tweet_count	84.82	12.3212	True	PA
protest_count	22.35	4.1296	True	PA
tweet_count	34.0	12.3212	True	RI
protest_count	12.48	4.1296	True	RI
tweet_count	9.63	12.3212	False	SC
protest_count	2.89	4.1296	False	SC
tweet_count	1.46	12.3212	False	SD
protest_count	0.05	4.1296	False	SD
tweet_count	24.4	12.3212	True	TN
protest_count	8.88	4.1296	True	TN
tweet_count	38.48	12.3212	True	TX
protest_count	11.05	4.1296	True	TX
tweet_count	22.81	12.3212	True	UT
protest_count	8.58	4.1296	True	UT
tweet_count	25.04	12.3212	True	VA
protest_count	2.44	4.1296	False	VA
tweet_count	44.02	12.3212	True	VT
protest_count	12.56	4.1296	True	VT
tweet_count	49.37	12.3212	True	WA
protest_count	11.78	4.1296	True	WA
tweet_count	27.77	12.3212	True	WI
protest_count	11.8	4.1296	True	WI
tweet_count	32.16	12.3212	True	WV
protest_count	3.87	4.1296	False	WV
tweet_count	32.56	12.3212	True	WY
protest_count	2.52	4.1296	False	WY
tweet_count	52.21	12.3212	True	USA
protest_count	17.19	4.1296	True	USA

Table 8.4: State-level co-integration tests

LSTM prediction error at the national and state-level

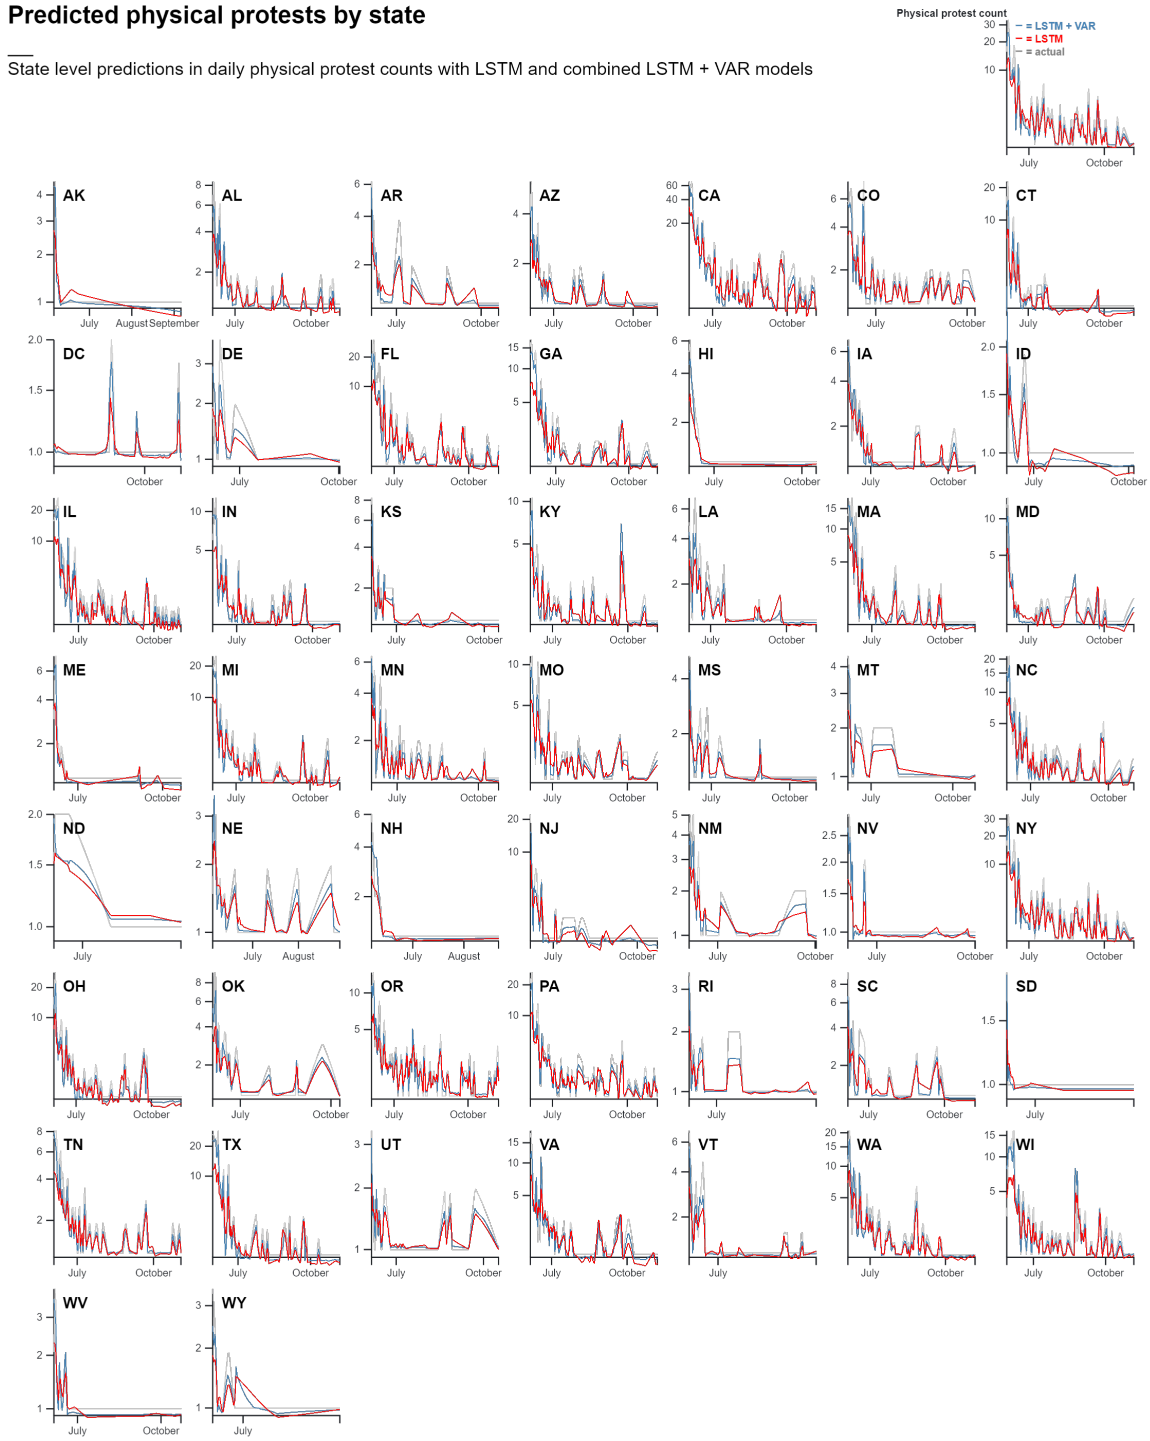

Figure 8.6: National-level physical protest count predictions with LSTM neural network (red) and combined VAR + LSTM neural network (blue).

tweet_count_lstm	tweet_count_lstm_var	protest_count_lstm	protest_count_lstm_var	state
1.63174465309194	0.9851871057985988	0.34166806706859554	0.254757023054979	AK
11.786798575957118	9.479509345884614	0.7181662198401891	0.5324006983283389	AL
6.650622331336926	4.590145606948933	0.4403163881744126	0.33943155017375976	AR
28.77921582992491	21.23793296574892	0.4485823686682536	0.30707144941654047	AZ
404.4084227098777	336.7806295058118	6.007457786461597	2.9490765513812516	CA
30.01621425355721	26.59957928280264	0.7993148435610226	0.5649572078856558	CO
15.334285431189157	11.02505529408219	1.896353718992696	1.3939363271591505	CT
23.736083785650997	23.81068069179058	0.11086437356012052	0.07793381179396766	DC
0.7325222082128885	0.7159716193011622	0.3421247586505661	0.2908898635219365	DE
88.12077214157925	69.51680985979907	2.7537371048288377	1.7902235603602197	FL
86.19234353603818	73.44897694097348	1.6713623859351712	0.8724158856359889	GA
11.792400674815646	9.984452420722613	0.40686146920549854	0.24948583249649936	HI
8.744681662462858	7.232269350463589	0.4478089436769372	0.29619341233604535	IA
2.4082968930696698	1.7056638770117498	0.16536612496970338	0.1531584145572045	ID
87.44650952462074	77.17844138526935	2.5512773094083148	1.4318422797671726	IL
29.64258375965722	24.856338336057984	1.0406612790715757	0.6146814918287979	IN
7.596514682829078	4.6646413313963055	0.538870446451288	0.4273335084921411	KS
35.27923582287327	33.346640730585136	0.858949893716185	0.6930373347654135	KY
15.082717459282277	12.49877140264469	0.6605243445159249	0.5197826014082174	LA
44.371240531710946	40.80492706779965	1.9651804282308123	1.2183749864623408	MA
17.220220305719156	12.134207658043596	1.054893421176555	0.6049653200942864	MD
4.712338896284972	4.779128526357501	0.4312755912392352	0.29121200653814017	ME
39.26284626658697	32.649136507918755	2.2736716827165804	1.2908681860732523	MI
37.60891257499972	30.988163629976786	0.5637418400213641	0.4416140288511977	MN
23.584773396686856	19.414096657499528	1.1517504417466091	0.7516604567751498	MO
7.6928909975281705	6.095523738663566	0.3597730114512824	0.2520372894607872	MS
2.461001533697911	1.9912504818000762	0.30638775656288	0.20538945196792338	MT
53.048239071872416	41.72391583153716	1.8599233018289973	1.1259518981366636	NC
1.0845398726706077	0.8065042804164961	0.23576849319214016	0.18078224158984635	ND
4.497626957375036	3.2732644772140986	0.2814055141637685	0.21466245161975894	NE
6.917417567540771	5.957940319807103	0.3279720760202953	0.31228440069476443	NH
22.033960911131445	14.502067721326917	1.29235694449712	0.9038380907979418	NJ
4.887234961705204	3.8938534158028233	0.4993964532394889	0.3597711472784601	NM
31.740694851241958	25.811818953071192	0.22938030479483743	0.20473067668484166	NV
193.32098471055713	155.19002733485738	2.7253931111086476	1.7080818533848026	NY
43.71758675433871	32.76663751839004	2.2714852356348474	1.6789330965437155	OH
11.236223729190684	8.593552591084453	0.8608258513145886	0.7041220905616566	OK
31.46619025087569	25.49808775889185	1.2707095962331014	0.9618626634376153	OR
64.45493012022322	48.564322163796064	2.1110752934090766	1.2551703442554134	PA
3.634083897780312	3.0198348841343443	0.23951033290628052	0.17125112133223994	RI
12.261648992214887	8.650423363451342	0.5935496190537264	0.44972230293626536	SC
1.094090331933008	0.6800541164605686	0.05360725929713853	0.07534496892633583	SD
31.0187440318061	25.165153238529058	0.8179922869302388	0.5046314795751303	TN
318.3269475667483	291.18826928563954	3.0199716619188743	1.696242563175341	TX
6.61191603109787	4.992015052961891	0.21566761713590943	0.1840298138968798	UT
26.172762526691542	19.803941462250105	1.478927738798053	1.0383042192858312	VA
1.8152779739658185	1.306537454018416	0.5080003443358236	0.32388305396713185	VT
55.959254478781034	45.185026789100725	1.6681479250630276	1.1719343201862271	WA
35.375162133862275	30.831447877571232	1.7882649235851182	1.303275537353162	WI
3.284356242592954	2.4738900360095832	0.23132859531003916	0.17959034062514084	WV
1.4774450121743548	1.2560884846069857	0.25293273451736775	0.18375998517876252	WY
900.9194036005148	673.1375783191256	11.79289524971523	12.795728123580574	USA

Table 8.5: National and state-level prediction error with both LSTM and LSTM + VAR models

8

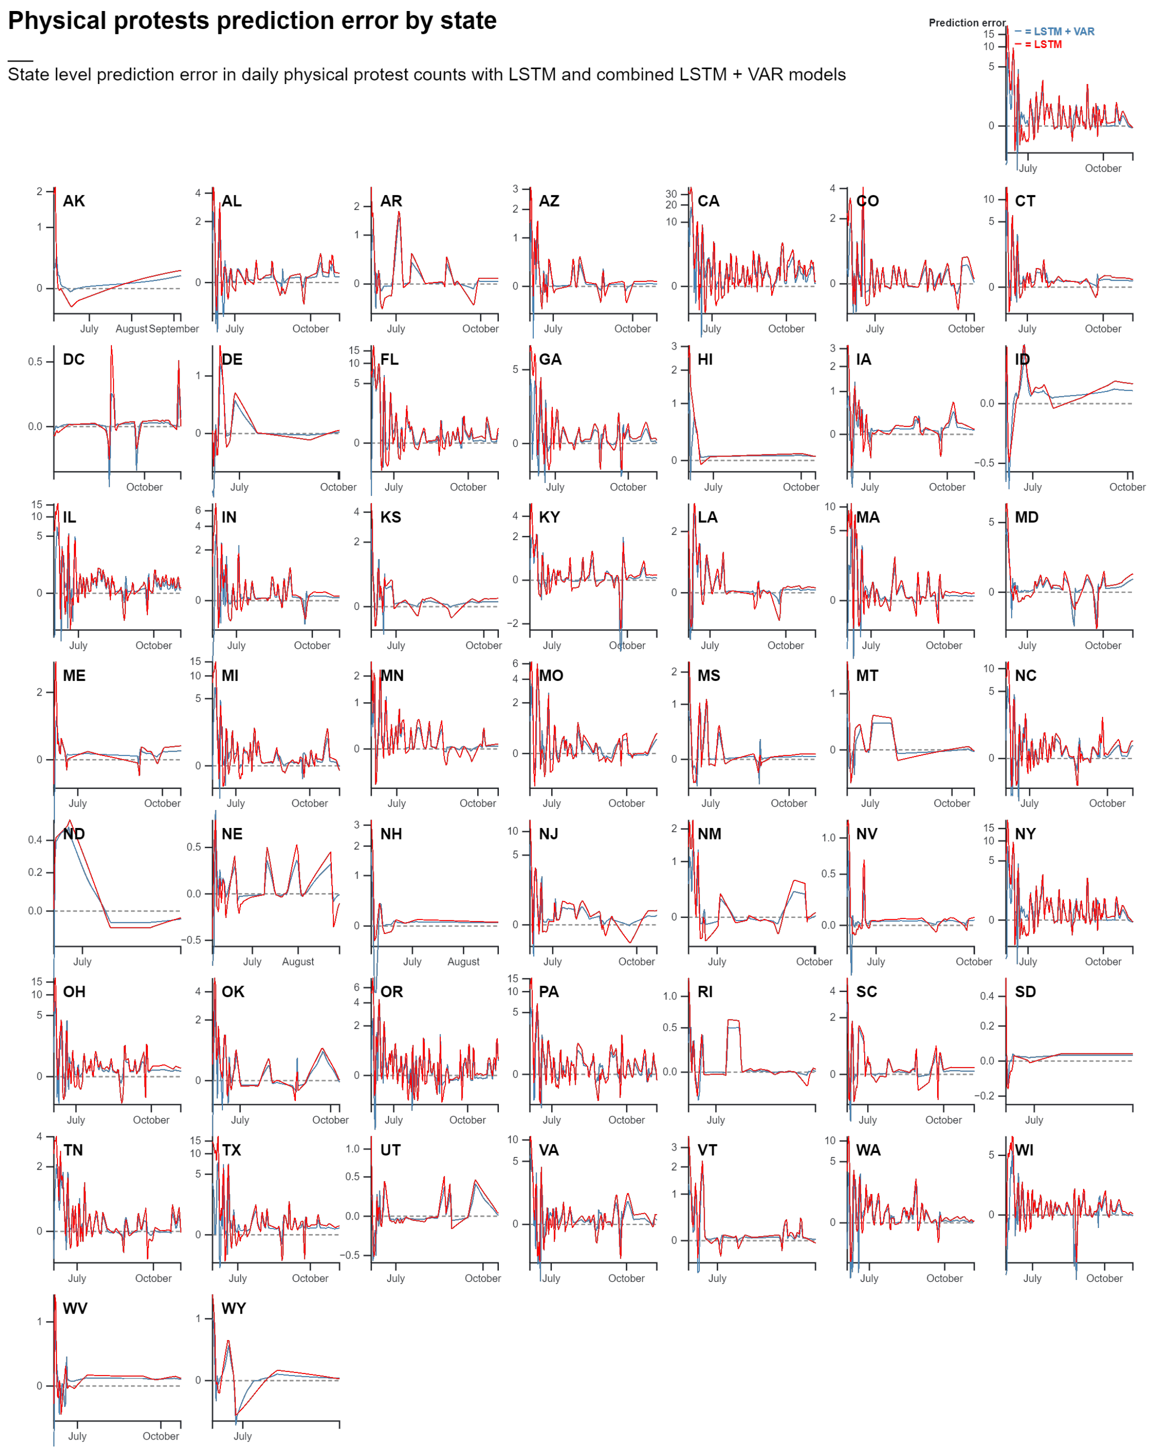

Figure 8.7: State-level physical protest count prediction error with LSTM neural network (red) and combined VAR + LSTM neural network (blue).