THE UNIVERSITY OF CHICAGO

Emotional Catalysts and Dynamic Networks: A Computational Analysis of Mobilization in Digital Activism

By

Huanrui CHEN

May 2025

A paper submitted in partial fulfillment of the requirements for the Master of Arts degree in the Master of Arts in Computational Social Science

Faculty Advisor: Zhao Wang

Preceptor: Ali Sanaei

Contents

1 Introduction **2**

2 Literature Review **7**

2.1 Networks and Social Mobilization 7

2.2 Emotional Contagion and Focusing Events 9

2.3 Network Perspective on Emotion and Mobilization 11

3 Data and Methods **12**

3.1 Research Setting 12

3.2 Data Sources 13

3.3 Data Annotation 14

4 Results **19**

4.1 RQ1: Reshaping the Emotional Landscape of Digital Discourse 19

4.1.1 Method: Operationalizing Emotional Change through Volume, Intensity, and Compositional Measures 19

4.1.2 Results: Sustained Surge in Participation and Elevated Emotional Intensity 21

4.1.3 Analysis: Structural Turnover and the Affective Logic of New Participation Drove Emotional Change 24

4.2 RQ2: Structural Transformation of the information diffusion 32

4.2.1 Method: Regression of Emotion–Diffusion Linkage and Retweet Network Analysis 32

4.2.2 Results: Emotional amplification and Structural Centralization in Information Diffusion 33

4.2.3 Analysis: New Participants and Influential Users Reshaped Information Diffusion 37

4.3 RQ3: Mechanisms of Cross-Community Diffusion 38

4.3.1 Method: Community Detection and Edge Classification 38

4.3.2 Results: Weak Ties Transmit Anger Across Communities, Strong Ties Mobilize Collective Action 41

4.3.3 Analysis: Affective Compensation Facilitated Outrage, Reciprocal Trust Supported Motivation 46

5 Conclusion **48**

Abstract

The murder of George Floyd served as a 'focusing event," an unexpected and tragic incident that captured public attention and reoriented ongoing discourse. In communication theory, such events often catalyze rapid agenda shifts and heightened collective expression, particularly within the digital public. This study investigates both structural and content-level transformations in information diffusion on Twitter before and after the event. Results indicate that: (1) user engagement intensified post-event; tweets exhibited increased expressions of anger and sadness, and users who entered the conversation afterward were more responsive to emotionally charged narratives; (2) the underlying social network became more centralized, shifting from decentralized user-to-user interactions to an influencer-driven broadcast model; and (3) despite this centralization, weak unidirectional ties across community boundaries played a critical role in bridging structural holes, enabling the wide dissemination of grievance-oriented, high-intensity emotional content, often framed diagnostically. Building on these findings, the study identifies a dual-pathway model of affective diffusion: Weak ties maximize emotional reach by spreading high-arousal messages across communities, while strong ties consolidate motivational framing. Together, these findings suggest that focusing events amplify user engagement and fundamentally reshape the structural and affective dynamics of online information diffusion, embedding emotional expression within the network's topology.

Keywords: Digital Activism; Collective Action; Network Analysis; Emotional Contagion; Media Effect

---

1 Introduction

Social media platforms have increasingly become vital spaces for expressing opinions, emotions, and social concerns, especially during moments of societal crisis. In such critical junctures, individuals use digital spaces to access information and participate in forming collective consciousness and social mobilization through emotional expression, affective resonance, and the construction and reproduction of narratives. Scholars have long recognized the role of social media in amplifying public emotions, facilitating group participation, and shaping antagonistic or solidaristic public narratives within political discourse (Bennett & Segerberg, 2012; Freelon et al., 2018). At the same time, platform interaction mechanisms and algorithmic recommendations further enhance the visibility of emotionally charged content, shaping both the prioritization of issues and emotional tendencies in the digital public sphere (Highfield, 2016; Papacharissi, 2015). As a result, social media has evolved beyond a mere channel for information dissemination into an arena for collecting public emotion and generating shared social imaginaries. Beyond emotional mobilization, scholars have engaged in deeper debates over the role of digital networks in shaping collective action. On one hand, some researchers emphasize the democratizing potential of digital networks, arguing that social media lowers participation barriers, accelerates information diffusion, and facilitates decentralized collaboration, thereby enhancing public engagement and collective efficacy (Earl & Kimport, 2011; Shirky, 2008). From this perspective, social media is seen as an *empowerment tool* that can swiftly aggregate individual efforts into broad, flexible frameworks of networked action. On the other hand, critics contend that digital mobilizations are often fragmented and short-lived, lacking stable organizational structures and sustained resource investment, which limits their capacity to generate enduring social change (Gladwell, 2010; Morozov, 2011). They further emphasize that while social media can rapidly ignite public emotions, such transient emotional surges rarely translate into deep, long-term collective commitments.

These divergent views are underpinned by contrasting theoretical frameworks. Resource mobilization theory attributes collective action to the accumulation of organizational resources and the availability of political opportunities (McCarthy & Zald, 1977). In contrast, emotional and cultural approaches highlight how affective resonance and cultural framing shape actors' perceptions of their circumstances and inform the scripts of collective action (Jasper, 2011; Polletta & Jasper, 2001). Together, these perspectives illustrate the complex dynamics of collective action within digital networks: while technological platforms lower connectivity thresholds and amplify emotional mobilization, participants continuously oscillate between affective surges and mobilization fatigue, renegotiating the boundaries of participation and the rhythms of action within a fragmented public sphere.

To better understand these dynamics, the concept of focusing events offers a critical analytical lens for examining the interplay between emotional expression and network transformation. Defined as sudden, high-impact public incidents, focusing events concentrate public attention and trigger intense emotional resonance, disrupting the fragmented flow of everyday information and transforming individualized emotional experiences into collective expressions and interactive momentum (Birkland, 1998; Kingdon, 2011). Unlike the gradual diffusion of routine topics, focusing events provoke widespread emotional responses accompanied by significant shifts in network structures due to their immediacy and symbolic significance.

To capture how these changes unfold, this study proposes the concept of *affective-structural coupling*, highlighting the reciprocal relationship between emotional dynamics and network reconfiguration. Social media platforms serve as critical venues for tracing the evolving interplay between affective resonance and community connectivity. During focusing events, shifts in emotional expression and network reorganization often occur simultaneously, exhibiting synchronicity and mutual reinforcement. However, existing scholarship typically examines these dimensions in isolation, focusing either on emotional content (Bail, 2016) or network structure (González-Bailón & Wang, 2016), with limited attention to their intertwined evolution within specific event contexts. Addressing this gap, this study integrates insights from multiple theoretical traditions. Frame alignment theory (Snow et al., 1986) helps explain the rapid formation of cognitive consensus during emotionally charged periods; research on emotional contagion (Kramer et al., 2014) illuminates how emotions propagate through networks; and network diffusion models (Centola & Macy, 2007) provide valuable perspectives on relational transformations. Rather than treating these perspectives separately, this study adopts an integrative approach to trace their dynamic interaction in the context of focusing events. Through this framework, the research seeks to understand the co-evolution of emotional expression and network structure during critical moments, capturing a multidimensional picture of collective expression and interactional dynamics within digital spaces.

To empirically examine the framework of affective-structural coupling, this study focuses on a paradigmatic focusing event in the digital sphere, analyzing how it catalyzed emotional resonance and reshaped network structures. Focusing events are distinguished by their capacity to concentrate social attention, trigger intense emotional reactions, and rapidly accelerate collective action through social media platforms (Birkland, 1998; Kingdon, 2011). Prior research has primarily examined emotional content (Bail, 2016) or network structures (González-Bailón & Wang, 2016) in isolation without fully illuminating how these two dimensions interact during critical events to produce large-scale mobilization effects. To address this gap, this project takes the murder of George Floyd in May 2020 as our empirical case. Although police violence against Black communities in the United States had previously sparked significant public attention, such as the cases of Michael Brown, Eric Garner, Tamir Rice, and Breonna Taylor (Lebron, 2017; Taylor, 2016), the George Floyd incident triggered an unprecedented wave of public protest. Within weeks, an estimated 15 to 26 million people participated in demonstrations nationwide (Buchanan et al., 2020), while solidarity networks flourished across digital platforms, cutting across diverse communities. This case provides an ideal setting to observe the whole arc of emotional escalation and network reconfiguration, offering critical insights into how focusing events animate collective expression and restructure digital mobilization networks.

On the data side, this project collected and analyzed a large-scale dataset of Twitter activity from January 2019 to December 2021, comprising 141,319 original tweets and over 5 million retweets related to racial justice and police violence. Twitter offers distinct advantages as a research platform: its openness and widespread use of hashtags facilitate the diffusion of movement discourse, and its central role in racial justice activism has been well documented (Freelon et al., 2016). Methodologically, this study employs generative large language models (LLMs) to automate the annotation of tweet content, accomplishing both emotion classification and framing analysis tasks. Compared to traditional keyword matching or sentiment lexicon approaches, generative LLMs demonstrate superior capability in capturing nuanced emotional expressions and discursive frames within context, especially when processing sarcasm, metaphor, or ambiguous statements (Benford & Snow, 2000; Mohammad & Kiritchenko, 2015). In addition, this project applied Louvain community detection algorithms (Blondel et al., 2008) to map the retweet network structure. It combined this with network diffusion models (Centola & Macy, 2007) to characterize the dynamic evolution of network topology before and after the event. Particular attention was paid to changes in network centralization, tie strength, and patterns of cross-community connections. Through this integrated research design, this project not only capture how the focusing event served as a catalyst for both emotional and structural transformations but also systematically investigate how fluctuations in collective emotion and network reconfiguration jointly shaped the pathways of digital mobilization. This approach allows this study to better understand the complex mechanisms underpinning collective action in the digital age.

The analysis demonstrates that while protest-related discourse was widespread across digital platforms, the evolution of emotional expression and network structure did not unfold uniformly but were significantly shaped by specific event-driven junctures. Notably, during the focusing event, expressions of anger surged and became a critical driver of information diffusion. This emotional shift reflected the intensification of public sentiment and created favorable conditions for reconfiguring the network structure. Further network analysis indicates that the prominence of central influencers increased markedly during the event, while newly established weak ties bridged previously isolated communities. This structural transformation suggests that the scope and speed of information flow are contingent upon the coupling between emotional activation and network centralization. Fluctuations in collective emotion provided the momentum for network reorganization, expanding the boundaries of information dissemination.

Moreover, the discourse surrounding the focusing event achieved notable expansion beyond its initial supporter communities, gradually permeating broader public discussions. This diffusion of issue framing not only enhanced the visibility of movement narratives but also fostered shared understandings of the event's significance across different communities, thereby deepening and broadening the flow of information. Thus, shifts in emotional expression, network structure, and discursive diffusion in digital space did not occur in isolation but rather co-evolved as interconnected and mutually reinforcing processes. This study successfully traced the dynamic interplay between affective shifts and structural transformations by integrating generative large language models with network analysis. The findings further reveal how these processes are interwoven with the pathways of discourse diffusion, collectively shaping the trajectories of information flow during periods of heightened social mobilization.

Building on this framework, this study systematically examines the dynamic transformations within digital spaces across three analytical dimensions: the evolution of emotional expression, the restructuring of network topology, and the expansion of discursive boundaries. First, the analysis traces changes in the intensity and composition of emotional expression before and after the event, with particular attention to the shift from sadness to anger and its prominence within the discourse. Second, it investigates the structural transformation of the retweet network, focusing on the increasing concentration of influence among central actors and the formation of new weak ties between communities to assess how network reconfiguration shapes pathways of information diffusion. Third, it explores the expansion of discursive boundaries by observing how movement frames penetrate beyond their original communities and reach a broader public sphere.

The study integrates emotion and framing labels generated by large language models with network structural analysis to achieve this. By synchronously tracking emotional fluctuations and network changes, this approach captures the co-evolution of affective and structural dynamics during the progression of the focusing event. Rather than merely describing patterns of change, the analysis emphasizes the coupling mechanisms between emotion and structure: whether surges in emotional intensity coincide with increasing network centralization and whether the formation of weak ties facilitates the diffusion of emotional narratives into new communities.

To guide this study on the dynamic interplay between emotional expression and network structure during a focusing event, the following research questions are proposed:

- **RQ1: How does a focusing event shape the emotional landscape of digital discourse?**
  Specifically, this study examines whether the murder of George Floyd amplified or suppressed different types of emotional expression, and how new participants in the conversation responded to varying emotional narratives.

- **RQ2: In what ways does a focusing event reconfigure the topology of online communication networks?**
  This question explores whether the focusing event shifted the balance between decentralized peer-to-peer communication and more centralized, broadcast-style structures in the dissemination of information, and how such transformations manifested in prac-

tice.

- **RQ3: What structural mechanisms facilitate or inhibit the cross-community diffusion of emotional narratives following a focusing event?**
  The analysis aims to identify broader structural conditions, such as bridging ties, clustering, and network centralization,which determine whether and how emotional narratives transcend their initial community boundaries.

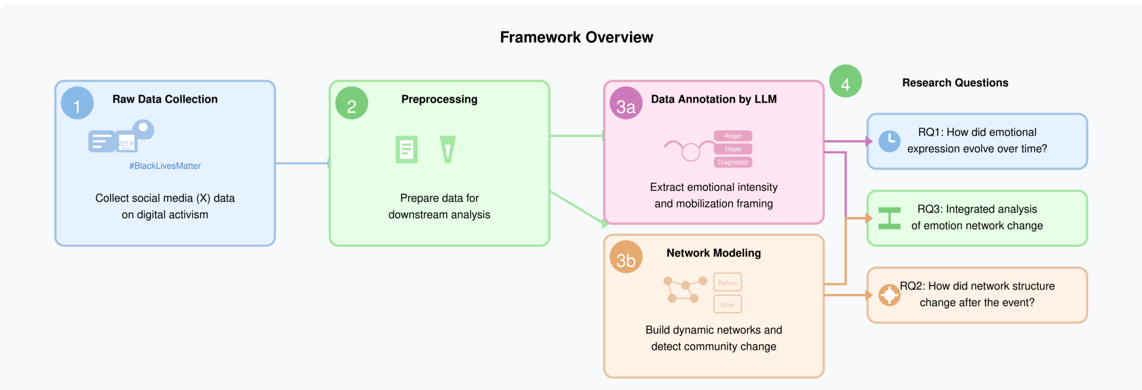

Figure 1: Research Design and Analytical Workflow

2 Literature Review

2.1 Networks and Social Mobilization

In the study of social mobilization, relational networks are widely recognized as a crucial factor influencing participation in and diffusion of collective action. Granovetter's (Granovetter, 1973) theory of strong and weak ties laid the foundation for network analysis, emphasizing that strong ties, characterized by close and enduring interactions, foster high levels of trust and responsibility among participants, providing critical social support to mitigate the risks associated with activism (McAdam, 1986; Snow & Benford, 1988). In high-cost collective actions, such trust and sense of belonging effectively reduce participants' hesitation and strengthen their identification with and commitment to group objectives (Kitts, 2000; Passy, 2003).

By contrast, due to their lower frequency of interaction and limited emotional depth, weak ties have traditionally been regarded as secondary pathways in mobilization processes (Diani, 2003; Klandermans, 1997). Although weak ties are valuable for disseminating information and broadening the reach of potential mobilization, they lack the strong social bonds and willingness to undertake risks essential for sustaining high-risk, long-term collective action (McAdam & Paulsen, 1993).

However, with the widespread adoption of social media, scholars have begun to reconsider the functional boundaries of weak ties. Social platforms transcend geographical and community constraints, fostering large-scale, loosely structured user networks (Hampton et al., 2011; Rainie & Wellman, 2012). Although most connections on these platforms are typical weak ties, research has shown that such expansive networks constitute crucial channels for information diffusion (Bakshy et al., 2012). Gil de Zúñiga et al. (Gil de Zúñiga et al., 2012) and Valenzuela et al. (Valenzuela et al., 2009) further argue that social media, by lowering participation thresholds, grants ordinary users greater opportunities to engage in public affairs, allowing weak ties to play a foundational role in the early stages of mobilization.

Particularly during the initial outbreaks of sudden events or emergent social issues, weak ties enable the rapid dissemination of information, facilitating swift exposure for social movements (González-Bailón et al., 2013; Vromen et al., 2015). Moreover, the interactive affordances of social media platforms endow weak ties with multidimensional possibilities for engagement. These platforms enable previously marginalized users to transcend community boundaries and access broader flows of information and public discourse (Meraz & Papacharissi, 2013). Studies have found that even without substantive personal connections, social media's visibility and interactive design can still stimulate users' interest in participating in public conversations (Barberá et al., 2015; Valenzuela, 2013). This mechanism suggests that weak ties on social media facilitate information transmission and foster emotional identification with public issues and initial willingness to act (Bode, 2017; Halupka, 2014).

Nevertheless, scholarly debates remain unresolved regarding whether weak ties can sustain deeper forms of mobilization in collective action contexts. On one hand, some studies emphasize that weak ties within social media environments enable rapid information diffusion and issue visibility through their expansive reach (González-Bailón, 2020; Theocharis et al., 2015). On the other hand, critics argue that while weak ties offer clear advantages during the early stages of mobilization, their lack of depth and sustained interaction renders them insufficient to sustain complex and enduring mobilization processes independently (Bimber et al., 2005; Halupka, 2014).

Overall, existing research has predominantly focused on the structural dimension of weak ties in information dissemination, paying less attention to how content-level factors may influence their mobilizing potential. In particular, there is limited understanding of how emotional expression activates or constrains engagement within weak-tie networks on social media platforms. This opens an important avenue for further investigating the interplay between weak ties and content diffusion dynamics. On social media, content functions as a vehicle for information transmission and an emotional catalyst that stimulates user engagement. Emotional diffusion may enhance users' resonance with public issues and responsiveness within structurally loose weak-tie networks, presenting a critical lens for understanding patterns of digital mobilization. The following section thus turns to this underexplored intersection, examining how emotional content propagates through social media network structures to transform widespread but loosely connected networks into more vibrant spaces of public participation.

2.2 Emotional Contagion and Focusing Events

Emotional Contagion is widely recognized as a critical driver of information flow within social media. Classical psychological research has long demonstrated that emotions can rapidly spread among individuals through unconscious social contact, a phenomenon known as emotional contagion (Hatfield et al., 1994). In the digital era, the interactive affordances of social media platforms, including real-time engagement, algorithmic curation, and visual content formats, have further accelerated the speed and reach of emotional transmission (Ferrara & Yang, 2015). Experimental studies provide robust empirical evidence for these dynamics. Kramer et al. (Kramer et al., 2014), in a large-scale experiment conducted on Facebook, found that the emotional valence of content delivered by the platform significantly influenced users' emotional expressions, triggering cascading effects across their networks. Similarly, Stieglitz and Dang-Xuan (Stieglitz & Dang-Xuan, 2013) observed that emotionally charged content, particularly posts with clear attitudinal positions, was more likely to capture user attention and interaction, thereby increasing information visibility.

Different types of emotions exhibit distinct patterns of diffusion across social media networks. Berger and Milkman (Berger & Milkman, 2012) found that high-arousal emotions such as anger, excitement, and awe possess more significant diffusion potential than low-arousal emotions like sadness or calmness. This advantage partly stems from the physiological and psychological activation triggered by these emotions and users' anticipation of social feedback in the form of likes and shares within the social media environment. Similarly, Brady et al. (Brady et al., 2017) observed that moral-emotional language significantly increased information sharing, allowing emotionally charged content to spread rapidly. Furthermore, some studies argue that social media platforms' design features amplify emotional expression. Benski and Fisher (Benski & Fisher, 2014) suggested that platform architectures encourage users to engage in public discussions through emotionally expressive modes of communication. Extending this argument, Papacharissi (Papacharissi, 2015) posited that social media provides a space for emotional expression and the construction of public narratives, making emotion a central force in shaping the dynamics of the digital public sphere.

However, it is important to note that the effects of emotional diffusion are not uniform. Some research cautions that emotional transmission may remain confined within existing communities, leading to "echo chamber," effects which inhibit the broader cross-community dissemination of content (Barberá et al., 2015).

In this context, scholars have increasingly recognized that specific societal events often serve as critical triggers for emotional contagion. Birkland (1997) defines such occurrences as "focusing events," referring to sudden and highly visible incidents that attract widespread public attention. By their abruptness and symbolic resonance, these events concentrate emotional expression among users and propel related issues to the forefront of public discourse. Unlike the routine flow of information, focusing events can temporarily dissolve community boundaries, significantly accelerating the speed and scope of information diffusion (McAdam & Sewell, 2001). For instance, during major natural disasters, social media users often rapidly converge in discussions concerning the unfolding crisis, expressing concern and sympathy, which generates pronounced peaks in emotional diffusion (Fraustino et al., 2012). Similarly, during the Australian bushfires, Alam et al. (Alam et al., 2020) observed that emotions of fear and anger dominated conversations on social media, coinciding with a dramatic surge in information dissemination.

However, focusing events do not necessarily guarantee widespread emotional diffusion. Wolfe (Wolfe, 2012) argues that the effectiveness of emotional mobilization depends on multiple factors, including the nature of the event, how the media frame it, and the level of public attention before its occurrence. Binder's (Binder, 2020) research on climate disasters further demonstrates that even with similar disaster intensities, the degree of emotional expression on social media varies significantly across regions. These findings suggest that while the relationship between focusing events and emotional diffusion is broadly observed, its specific manifestations remain highly context-dependent. Nevertheless, existing studies have predominantly focused on describing the correlation between focusing events and emotional spread without systematically understanding how emotional diffusion operates across different network environments, particularly within social media platforms' dynamic and expansive ecosystems.

In summary, existing research has demonstrated a close connection between emotional diffusion on social media and the occurrence of focusing events. Such events often act as critical nodes for emotional dissemination, triggering widespread and rapid affective expression among users and serving as pivotal moments for the emergence of collective emotion in digital spaces. However, current scholarship has primarily concentrated on identifying general patterns of emotional diffusion and the activation of emotions during focusing events. There remains a lack of systematic description of how these two dynamics, emotional dissemination and focusing events, interact concretely in social media mobilizations and how emotional diffusion shapes network relationships and stimulates public participation. This gap presents an important opportunity for advancing our understanding of the interplay between emotional diffusion and collective mobilization in social media environments, particularly in exploring how emotions are activated within networks to foster engagement with public issues.

2.3 Network Perspective on Emotion and Mobilization

Despite the valuable insights offered by existing research on the relationship between emotional diffusion and social mobilization in social media environments, much of this scholarship has approached the topic from fragmented perspectives. Prior studies have separately examined the general patterns of emotional dissemination (Brady et al., 2017), the mobilization potential of weak-tie networks (Bakshy et al., 2012; Gil de Zúñiga et al., 2012), and the emotional triggering effects of focusing events (Fraustino et al., 2012). While these strands of research collectively underscore the importance of both affective dynamics and network structures in shaping collective mobilization, the interconnections between these factors remain underexplored.

This study seeks to systematically examine the intertwined processes of emotional expression and network dynamics during focusing events. Rather than isolating the effects of either emotion or network structure, the focus is on how these two dimensions become mutually embedded to construct an ephemeral yet powerful mobilization environment within digital spaces. In other words, instead of treating affective expression and relational ties as independent variables, this research conceptualizes them as dynamically intertwined components of an evolving network. During critical social events, they jointly shape patterns of information flow and collective response.

Adopting this nested perspective allows for moving beyond reductionist analyses of individual factors toward a more comprehensive understanding of how emotions and relationships co-produce the architecture of digital mobilization. Empirical exploration of this perspective remains limited. Most studies focus either on how focusing events trigger emotional peaks (Gong et al., 2023), or how weak ties facilitate the diffusion of information to new audiences (Bakshy et al., 2012). However, few have systematically depicted how these two forces combine and evolve dynamically throughout an event. Some scholars have attempted to link network structure and information diffusion (e.g., (Tufekci, 2017)) yet often overlook the complex role of emotions within these networks. Similarly, although research on emotional contagion highlights how affect circulates within social networks (Kramer et al., 2014), it rarely considers how these emotions leverage network structures to achieve broader dissemination and resonance.

To understand this intertwined process, this study conceptualizes emotional expression and network structure as two interdependent dimensions nested within a dynamic relational system. Rather than operating in isolation, they jointly interact fluidly during focusing events to reshape information flow and collective response pathways. Emotional expression is not merely a surface-level feature of information diffusion but also a driving force behind network reconfiguration. High-intensity emotions, such as anger or sadness, activate user participation, stimulating the intensification of internal network links and expanding external connections. At the same time, transformations in network structure recursively shape the trajectories and reach of emotional narratives, influencing the extent of their resonance and diffusion. In other words, digital mobilization could be understood as a mutually reinforcing process between affective dynamics and structural evolution, rather than a linear sequence dominated by a single factor. Emotional peaks triggered by focusing events are rapidly amplified through platform algorithms and community interaction patterns, generating new connections and clustering effects within the network. Simultaneously, these evolving network configurations determine the boundaries of emotional diffusion, allowing certain narratives to break free from their original community confines and extend into broader spheres of public discourse.

3 Data and Methods

3.1 Research Setting

This study investigates how focusing events reshape the emotional and structural dynamics of digital discourse by analyzing Twitter conversations surrounding the murder of George Floyd. In political communication, focusing events, defined as sudden, high-impact incidents that capture public attention, often trigger large-scale shifts in issue salience, emotional expression, and discursive framing (Birkland, 1998). The core research question driving this study is how emotional content spreads in response to such events and how that content interacts with the rhetorical and structural logics of political mobilization.

To address this question, the study focuses on Twitter for both theoretical and empirical reasons. The platform's public architecture and real-time communication affordances make it particularly well-suited for tracking large-scale emotional expression and information diffusion in response to unfolding events. Previous work has shown that Twitter plays a central role in digital activism, particularly in the Black Lives Matter (BLM) movement, by enabling decentralized participation and hashtag-driven visibility (Freelon et al., 2016). These features make Twitter an ideal setting for studying affective and strategic communication during protest cycles.

To unpack how emotional and rhetorical content interact with structural diffusion mechanisms, the study applies a multi-layered analytical framework. First, it employs emotion annotation based on six basic emotions (anger, fear, sadness, disgust, surprise, joy), capturing both the type and intensity of affective expression. Second, to examine how discourse mobilizes action, the study uses Snow and Benford's (Snow & Benford, 1988) typology to classify each tweet's dominant framing strategy—diagnostic, prognostic, motivational, or irrelevant. This dual annotation strategy allows the study to distinguish between diferent expression of mobilization, which often coexist but serve different discursive purposes.

Finally, to link this content-level analysis with the dynamics of information flow, the study constructs directed retweet networks before and after the Floyd incident. In these networks, each edge represents a flow of information, enabling identification of tie strength (reciprocal or unidirectional), community boundaries, and bridging structures. These structural attributes are essential for understanding how emotionally and rhetorically charged content moves across digital publics, whether it stays within ideological clusters or diffuses broadly through weak-tie bridges. Together, this setting supports an integrated investigation of how focusing events restructure digital discourse through the interplay of emotion, mobilization framing, and network topology.

3.2 Data Sources

The dataset comprises two primary components: original tweets and retweet networks associated with the BLM movement. Data were collected using a third-party Twitter API through a comprehensive query strategy designed to capture the broadest possible representation of BLM-related discourse. The core sampling criteria focused on tweets containing the hashtags #BLM or #BlackLivesMatter, widely recognized as central digital anchors of the movement. This approach yielded an initial corpus of 141,319 original tweets, collectively generating 5,867,481 retweets and involving over 2.3 million unique users.

For each original tweet, this study collected rich metadata, including the author's unique identifier, tweet content, timestamp, and the full list of retweeter IDs. These data enabled the construction of a detailed retweet network, where each retweet constitutes a directed edge from the original author to the retweeter. Table 1 summarizes the structure of the raw dataset, capturing both the textual and interactional layers essential for our analysis. To capture temporal dynamics, this study structured our analysis around two 12-month windows: the year preceding the Floyd incident (May 2019 – May 2020) and the year following it (May 2020 – May 2021). While shorter time frames might better isolate immediate reactions, this broader temporal scope enables this study to analyze not only short-term spikes in engagement but also the sustained evolution of emotional tone, narrative strategies, and network architecture over time.

Table 1: Structure of Raw Tweet-Level Dataset

Field	Description
`tweet_id`	Unique identifier for each original tweet
`author_id`	Unique user ID of the original tweet's author
`text`	Full text content of the tweet, including hashtags, mentions, and emoji
`created_at`	Timestamp of when the tweet was posted (UTC)
`retweeter_ids`	List of user IDs who retweeted the original tweet (used to construct the retweet network)

Together, these data provide a high-resolution view of public discourse across a critical three-year period. They serve as the empirical foundation for analyzing how the George Floyd event, as a focusing event, reshaped the emotional, rhetorical, and structural contours of a contemporary social movement within the digital sphere.

3.3 Data Annotation

To annotate emotional content at scale, this study adopted the Six Basic Emotions model from psychology, categorizing human emotional expression into six universal types: anger, fear, disgust, sadness, surprise, and joy (Ekman, 1992). Beyond simple positive-negative sentiment polarity, this framework enables finer-grained distinctions between negative emotions, which is especially valuable in political contexts where different emotions (e.g., anger versus sadness) have distinct behavioral implications. Each tweet was annotated for emotion type and intensity, with intensity scored on a continuous scale from 0.0 (none) to 1.0 (extreme expression).

To handle large-scale annotation efficiently and consistently, this study employed the DeepSeek-V3 generative language model, guided by a structured prompt engineering strategy. Following OpenAI's official prompt design principles (OpenAI, 2024), three specific strategies were implemented to optimize annotation quality:

1. **Clear task objective:** The prompt began by explicitly stating the goal of classifying emotion types and scoring their intensity.

2. **Structured output requirements:** The model was instructed to use a numeric scoring format for each emotion, clearly specifying intensity levels.

3. **Concrete examples and reasoning:** To aid the model's understanding, fully annotated examples covering each of the six emotions were provided, complete with explanations of the reasoning behind the classification.

For example, the following case was included in the prompt materials:

> "In Louisville, Black Lives Matter organization is purchasing vacant homes for low-income families to promote stability in the community and fight gentrification."

The model was instructed as follows:
This tweet describes a positive community action, which brings hope and joy. The tone is constructive and uplifting, reflecting joy as the dominant emotion. Additionally, there is a mild element of surprise, since such generous actions may be unexpected in the broader public narrative. There are no clues of negative emotions such as anger or fear, so those are scored as 0.0.

Including the reasoning alongside the examples helped improve the model's interpretive accuracy during annotation. In practice, this study observed that providing this kind of explicit reasoning significantly optimized the model's ability to generalize to unseen examples across the dataset.

To assess the reliability of the automated annotations, this study conducted manual validation on a stratified random sample of 100 tweets. Specifically, the sample was stratified to ensure a balanced representation across all six emotion categories, randomly selecting approximately equal numbers of tweets for each emotion to avoid skewed evaluation. The comparison between the model's predictions and human judgments showed 87% agreement for emotion classification. Agreement was determined by matching the dominant emotion assigned by both the model and the human coder. Importantly, this evaluation focused solely on the classification of the dominant emotion, without assessing the exact numeric intensity scores assigned to each emotion. Annotations were considered matching as long as the primary emotion labels aligned, regardless of minor variations in intensity rating.

To further quantify intercoder reliability while accounting for chance agreement, Cohen's Kappa was computed. Given the stratified sampling across six emotion categories, Cohen's Kappa was approximately 0.84, indicating a very good agreement according to the commonly used interpretation scale (Landis & Koch, 1977).

A closer examination of annotation discrepancies revealed two primary types of errors. The first type arose from tweets expressing blended emotional signals, where multiple emotions such as sadness and anger co-occurred with comparable intensity scores. For example, the tweet "I do not like violence and am very sensitive to it, but I feel not sharing what has been going on or how outraged I am by the things the government in my country are doing would be a betrayal of what I believe in. Please be safe. Black Lives Matter." was classified by the model as anger. However, a human coder could plausibly interpret sadness or fear as the dominant emotion, given the emphasis on vulnerability and distress. Such cases highlight the potential difficulty of assigning a single dominant label when emotional expressions are nuanced and overlapping.

The second type of error stemmed from the presence of figurative or rhetorical language, such as sarcasm or ironic framing. In the tweet ”So it looks like Black Lives Matter is spending their $10.6 billion to promote twerking for Martin Luther King!?”, the model assigned anger as the dominant emotion. Nevertheless, a human reader might alternatively perceive surprise or disgust as more appropriate, depending on how the sarcasm is interpreted. But it is also important to note that this type of ambiguity presents challenges not only for automated models but also for human annotators, particularly when limited context is available. These observations suggest that future improvements could involve allowing for multi-label emotion annotations and try to introduce detailed context, capturing cases where multiple emotions are simultaneously salient rather than forcing a single dominant classification.

In addition to emotion annotation, this study also classified each tweet according to its political mobilization frame. While prior research has shown that high-arousal negative emotions such as *anger* play a crucial role in protest dynamics (**jasper2011emotions**; Goodwin et al., 2001), emotional expression alone is insufficient to explain why protest-related messages possess mobilizing power. In other words, expressions of anger or sadness may represent mere emotional catharsis, or they may be embedded in more structured narratives that strategically call for action. Without accounting for this distinction, an analysis focused solely on emotional tone risks overlooking the deeper political intent and organizational function of protest discourse.

To address this limitation, this study adopts the collective action framing typology developed by Snow and Benford (1988), which remains a foundational framework in the study of social movements. According to this framework, protest messages can be categorized into three basic frame types: *diagnostic* frames, which identify a social problem and assign blame to specific actors; *prognostic* frames, which propose strategies or solutions; and *motivational* frames, which issue moral or emotional calls to action aimed at mobilizing participants. Together, these three functions form the “meaning work” of protest discourse and remain highly relevant in digitally mediated activism.

Given the informal and polysemous nature of social media language, this study also includes an additional “irrelevant” category to account for tweets that mention protest-related hashtags (e.g., #BlackLivesMatter) but do not convey substantive political content, such as celebrity news or promotional materials unrelated to collective action.

By jointly annotating both emotional categories and framing types, this study aims to capture structure, affect, and discourse. For instance, two tweets may both express anger, but one may merely vent frustration, while the other explicitly attributes blame to law enforcement and urges followers to participate in protests. Without disaggregating their rhetorical structures, researchers might mistakenly treat expressive outbursts and strategic mobilization as equivalent phenomena.

To guide the annotation task, each frame type was defined using standardized criteria and illustrated with real examples drawn from manually labeled tweets:

- **Motivational Frame:** Issues a moral or emotional call to collective action, emphasizing the urgency and necessity of engagement. It does not offer specific solutions.
Example: "Tonight, we must rise up now, for our people and the rights we've lost."

- **Diagnostic Frame:** Identifies and attributes a social problem without attempting to prompt direct action.
Example: "Police use excessive force against marginalized communities."

- **Prognostic Frame:** Proposes specific solutions or forward-looking strategies, distinguishing it from diagnostic and motivational frames.
Example: "We need legal reforms to ensure police accountability."

- **Irrelevant:** Mentions protest-related terms but lacks substantive political messaging.
Example: "Michael B. Jordan is People's Sexiest Man Alive in 2020! The actor is known for starring in 'Black Panther' and for supporting the Black Lives Matter movement."

Operationally, the same generative language model used for emotion annotation (DeepSeek-V3) was employed for frame classification, supported by prompt engineering. Prompts included clearly defined objectives, explicit output formats, and representative tweet examples with corresponding reasoning. This combination of examples and explanations enhance the model's generalization in handling nuanced language tasks, enabling consistent and scalable annotation of protest discourse across the dataset.

Manual validation of 100 randomly selected tweets, stratified to include a balanced representation of all frame categories, yielded a 74% agreement rate between model predictions and human coding for framing classification. Agreement was defined as alignment on the dominant framing function. To further assess intercoder reliability while accounting for chance agreement, Cohen's Kappa was computed. Given the stratified sampling across four framing categories, the estimated Cohen's Kappa was approximately 0.65, indicating a substantial level of agreement according to the commonly used interpretation scale (Landis & Koch, 1977).

Further examination of misclassifications revealed the potential type of error. Discrepancies often arose from tweets blending diagnostic and motivational elements. For example, the tweet "Black Lives Matter fell off after the election because they accomplished their goal. Asian Lives Matter will now rise to counter anti-Chinese sentiment in the US." was labeled by the model as motivational. While the diagnostic frame is dominant, highlighting the rise of anti-Asian sentiment as a social problem, the final phrasing ("rise to counter") could also be interpreted as implying a motivational call to action. Such cases illustrate the challenge of assigning a single dominant frame when tweets simultaneously present a problem and hint at a response.

Finally, while the dataset benefits from transparent and validated annotation procedures, this study remains attentive to potential limitations. Twitter users do not represent the general population, and highly active accounts or automated bots may influence emotional expression and framing patterns. This study applied activity-based heuristics to flag accounts with exceptionally high posting or retweeting frequencies to address this. These were then manually inspected for characteristics commonly associated with automation, such as repetitive posting behavior and minimal engagement diversity. Additionally, although structured prompts and reasoning-based examples are used, some ambiguity inevitably remains in interpreting complex or sarcastic messages. In particular, sarcasm, coded language, and culturally specific idioms posed annotation challenges, as the model occasionally misinterpreted these nuances despite the prompt guidance. In addition to rhetorical complexity, another notable source of ambiguity arose from tweets exhibiting blended emotional signals. Many tweets simultaneously conveyed multiple emotions, such as anger, sadness, and fear, with comparable intensity, making it difficult to identify a single dominant emotional category. In such cases, both the model and human annotators sometimes diverged in selecting the primary label. These limitations are acknowledged and considered in the interpretation of findings. Future work could incorporate chain-of-thought annotation strategies to more explicitly capture annotator reasoning and manage complex cases (Wei et al., 2022).

In addition to content-level variables, it is also essential to capture the structural dimension of discourse to understand how messages circulate across users and communities. While emotion and framing annotations reflect what is being said, the retweet network provides a complementary perspective on how information flows, identifying the relational architecture through which narratives gain visibility and reach. To operationalize these structural dynamics, this project constructed directed retweet networks for both the pre-event and post-event periods. In these networks, each node represents a unique Twitter user, and each directed edge represents a retweet action, from the original tweeter to the retweeter, indicating the flow of information. For example, if user A posts a tweet and user B retweets it, a directed edge is drawn from A to B. This directional structure is analytically important, as it enables the reconstruction of information pathways, identifying the originators of discourse and the amplifiers extending its reach.

4 Results

4.1 RQ1: Reshaping the Emotional Landscape of Digital Discourse

4.1.1 Method: Operationalizing Emotional Change through Volume, Intensity, and Compositional Measures

To systematically examine whether and how the murder of George Floyd reshaped the emotional structure of digital protest discourse, this study conceptualizes "emotional change" as a multidimensional and temporally dynamic process. Rather than focusing on short-lived emotional reactions, this project analyzes how collective affect evolved across three interrelated dimensions: the scale of emotional participation, the intensity of affective expression, and the structural composition of expressed emotions. These dimensions reflect different layers of transformation, from micro-level emotional articulation to macro-level behavioral shifts, and together constitute a comprehensive framework for capturing systemic changes in the digital emotional landscape.

The dataset consists of all tweets related to the Black Lives Matter (BLM) movement posted between May 2019 and May 2021. Tweets were collected using a combination of keyword and hashtag-based queries (e.g., #BlackLivesMatter, #BLM), and subsequently filtered to remove duplicates, non-English content, and irrelevant or spam-like material. Each tweet was annotated for emotional content using a transformer-based language model fine-tuned for affective classification in short-form social media text. The model identifies the dominant emotion expressed in a tweet based on Ekman's (1992) six basic emotion categories, anger, sadness, joy, fear, disgust, and surprise, and assigns a corresponding confidence score representing the intensity of that emotional signal. These scores were normalized to a continuous [0, 1] scale and treated as proxies for emotional strength.

At the behavioral level, the first step of the analysis assessed whether the event produced a large-scale increase in public expression. To capture this, the monthly volume of BLM-related tweets was calculated and plotted as a time series (Figure 2). This measure serves as a behavioral indicator of public attention and expressive engagement. Covering a two-year window before and after the event, the time series enables comparison between routine, low-frequency periods and the post-event spike, allowing for empirical evaluation of whether the George Floyd incident constitutes a focusing event in line with Birkland's theory (Birkland, 1998).

Beyond participation volume, this study examines whether the intensity of emotional expression experienced structural shifts. All tweets were assigned dominant emotion labels and intensity scores (defined by the classifier's confidence). These intensity values were aggregated daily and smoothed using a 7-day moving average to visualize long-term trends while minimizing short-term noise. The resulting time series (Figure 3) represents the average daily emotional intensity, enabling assessment of whether the emotional landscape transitioned from irregular surges to a more stable, elevated baseline. This provides a basis for evaluating whether the post-event discourse was characterized by sustained affective activation rather than episodic bursts, indicating potential affective institutionalization.

In addition to overall intensity, this study also investigates whether the emotional composition of the discourse, the distribution of dominant emotion types, underwent a systematic reconfiguration following the event. While emotional volume and intensity capture the scale and force of affective expression, compositional analysis reveals which specific emotional registers gained or lost prominence over time, thereby shedding light on shifts in the dominant emotional tone of collective discourse.

To assess these compositional dynamics, all tweets were grouped into two time periods: the pre-event period (before May 25, 2020) and the post-event period (on and after May 25, 2020). Within each period, the relative frequency of each dominant emotion category was calculated and visualized using proportional distribution charts (Figure 4). This allows for direct comparison of the affective profile of the discourse across two distinct political contexts. Additionally, to account for potential changes in emotional intensity within each category, the analysis tracks temporal trends in the average intensity of the three most prominent emotions, anger, sadness, and joy, across the full timeline. These trajectories are plotted in Figure 5, providing a dynamic account of how particular emotional modes intensified or stabilized over time.

Taken together, this methodological strategy enables an assessment of how the George Floyd incident reshaped the emotional dynamics of BLM-related discourse. The analysis of monthly tweet volume (Figure 2) directly addresses whether the event generated a large-scale surge in expressive participation, thereby testing its qualification as a focusing event. The measurement of daily average emotional intensity (Figure 3) evaluates whether emotional expression became more stable and elevated in the aftermath, providing evidence of potential shifts in the affective baseline of digital discourse. Finally, the examination of emotional composition and intensity trajectories (Figures 4 and 5) reveals whether the relative salience and expressive strength of specific emotions underwent systematic reconfiguration following the event. By triangulating across these layers, the study captures the scale of public engagement and the structural features of emotional change, namely, whether affective expression intensified, stabilized, and realigned in content. These methods provide the empirical basis for evaluating whether the George Floyd incident marked a temporary emotional disruption or catalyzed a more durable transformation in the digital emotional regime.

4.1.2 Results: Sustained Surge in Participation and Elevated Emotional Intensity

The results reveal a pronounced and durable increase in the volume and affective intensity of BLM-related discourse following the George Floyd incident. These changes mark not only a short-term reaction to a crisis, but a fundamental shift in the emotional rhythm and scale of digital protest engagement. Figure 2 presents the monthly volume of BLM-related tweets between May 2019 and May 2021. Monthly tweet counts remained consistently low throughout the pre-event period, typically under 1,000. This pattern indicates a relatively stable but low level of engagement before the event. In stark contrast, tweet volume increased sharply in June 2020, the month immediately following Floyd's murder, surpassing 60,000 posts. This more than sixtyfold increase constitutes a statistically and substantively significant disruption in the communication landscape. While tweet counts declined slightly in subsequent months, the baseline level of activity remained markedly higher than pre-event levels. This post-event stabilization suggests that the incident not only triggered an immediate surge in public attention but also catalyzed a lasting expansion in the volume of political discourse related to BLM.

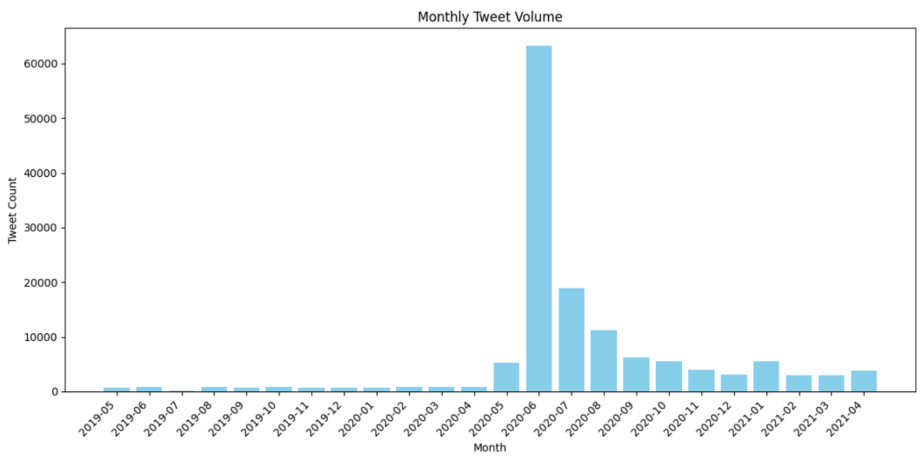

Figure 2: Monthly Tweet Volume of BLM-Related Posts (May 2019 – May 2021)

In addition to behavioral participation, the emotional tone of discourse exhibited a marked shift in intensity and temporal patterning. Figure 3 plots the average daily intensity of the dominant emotion expressed in each tweet. Prior to the event, emotional intensity was characterized by instability: sharp spikes were often followed by immediate reversals, producing an erratic affective landscape with no clear baseline. This volatility is consistent with emotional responses tied to routine fluctuations or minor events. However, beginning in late May 2020, emotional intensity rose substantially and stabilized at a consistently higher level. The post-event period is distinguished by two features: (1) elevated average intensity levels sustained over several months, and (2) reduced short-term volatility, suggesting that high-intensity emotional expression became a routinized feature of the discourse.

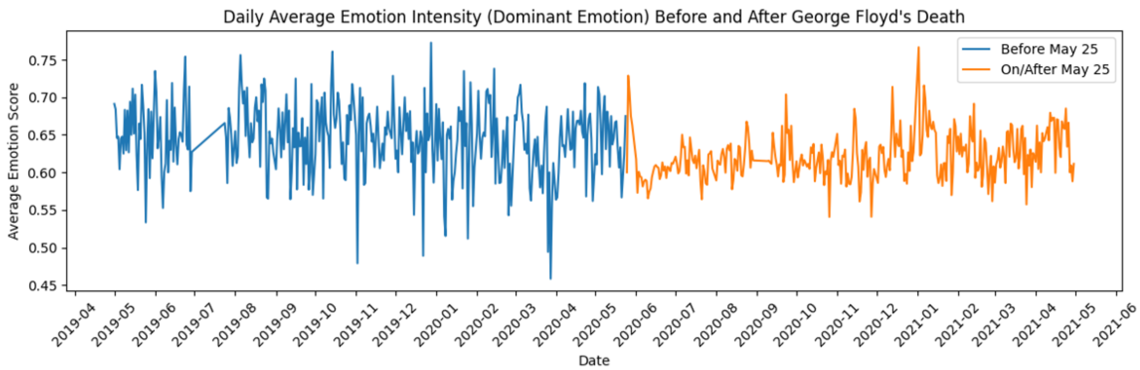

Figure 3: Daily Average Intensity of Dominant Emotion per Tweet

This stabilization is particularly notable given the inherent noise and reactivity typical of social media platforms. Rather than returning to a pre-crisis emotional baseline, the discourse remained in a state of sustained affective activation. This pattern suggests a qualitative transformation in the emotional environment: emotional expression ceased to be episodic and took on the characteristics of a persistent, high-arousal communication regime.

In addition to changes in scale and intensity, the emotional composition of BLM-related discourse shifted substantially after the George Floyd incident. Figure 4 compares the proportional distribution of dominant emotions before and after the event. During the pre-event period, emotional expression was relatively balanced: anger accounted for 43.1% of tweets, closely followed by joy at 41.9%, while sadness, fear, disgust, and surprise together constituted a minor share. This distribution suggests that, although grievance-related affect was prominent, it was tempered by a significant presence of positive emotional content. The post-event distribution, however, reveals a distinct reconfiguration. Anger became the dominant emotion, rising to 51.7% of all tweets, while joy declined to 34.3%. Sadness, though still secondary, increased from 8.2% to 10.6%, indicating a modest rise in expressions of grief or loss. Other negative emotions, such as fear and disgust , remained relatively infrequent, suggesting that the emotional shift was not a generalized turn toward negativity, but rather a selective intensification of particular affective modes. These patterns signal a reorientation of emotional discourse around anger as a central mobilizing affect, consistent with theories emphasizing the role of anger in protest escalation and blame attribution.

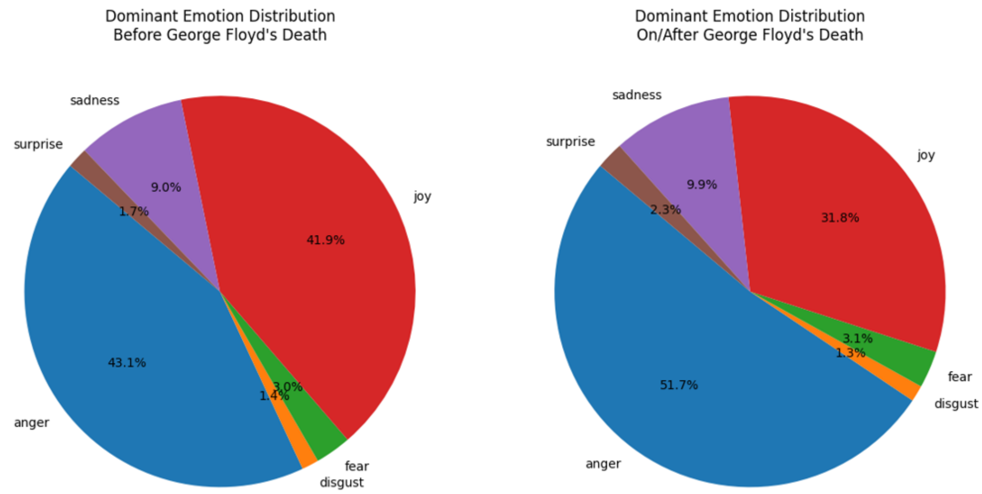

Figure 4: Distribution of Dominant Emotions Before and After the George Floyd Event

Further evidence of emotional realignment is provided by the intensity trajectories of the three most salient emotions: anger, sadness, and joy. Figure 5 presents the average daily intensity of these emotions over the full study period. Prior to the event, emotional intensities were moderate and tightly clustered, with all three emotions fluctuating below 0.4 on the normalized intensity scale. This pattern indicates a discourse environment where no single emotion dominated expressive force, and affective attention was likely distributed across multiple concerns or episodic issues. After May 25, 2020, the emotional landscape became more stratified. Anger intensity increased sharply rising above 0.5 immediately following the event and remained elevated for several months, even as overall participation declined. This sustained intensity suggests that anger became not only more prevalent, but more deeply embedded in the discourse's expressive structure. In contrast, the average intensity of joy declined and showed greater day-to-day volatility, indicating a less stable role in the post-event affective environment. Sadness rose modestly in both frequency and intensity, but did not match anger's prominence or endurance. Together, these results underscore that anger was not merely a transient reaction but rather emerged as a persistent emotional anchor within post-event discourse.

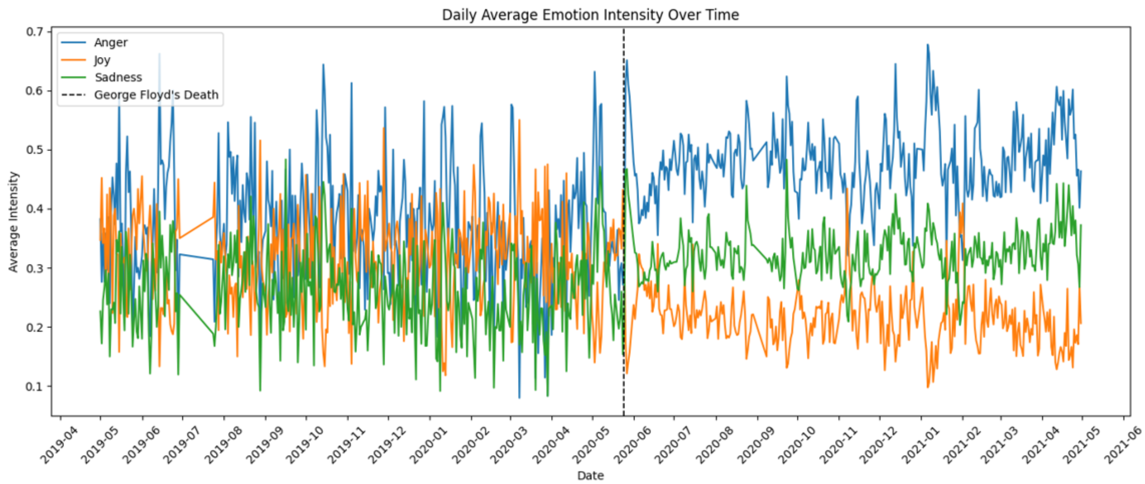

Figure 5: Daily Average Intensity of Anger, Sadness, and Joy

Taken as a whole, the results demonstrate a threefold affective transformation following the George Floyd incident: a surge in the scale of participation, a stabilization of emotional intensity at a higher baseline, and a compositional shift that positioned anger as the dominant emotional mode. These findings provide descriptive evidence that the event did not merely provoke short-lived emotional reactions but restructured the expressive norms of digital protest discourse in a sustained and selective manner.

4.1.3 Analysis: Structural Turnover and the Affective Logic of New Participation Drove Emotional Change

The pronounced increase in the volume and intensity of BLM-related discourse in the wake of the George Floyd incident offers strong empirical support for conceptualizing the event as a "focusing event" in the sense articulated by (Birkland, 1998). As shown in Figure 2, tweet volume remained low and stable throughout the pre-event period, then surged over sixtyfold in June 2020 and stabilized at an elevated baseline. This pattern suggests that the incident acted not merely as a temporary trigger of interest, but as a structural change to the rhythms of political expression on social media. Focusing events are defined by their capacity to disrupt existing discursive routines and concentrate public attention through their moral clarity, symbolic resonance, and high media visibility. The George Floyd incident met these criteria. The widely disseminated video footage, which visibly documented state-sanctioned racial violence, transformed abstract issues of structural racism into emotionally salient and publicly shareable content.

However, focusing events do more than direct attention; they reshape the affective pre-conditions for participation. As shown in Figure 3, the emotional intensity of tweets did not simply spike in the immediate aftermath, but stabilized at a consistently elevated level throughout the post-event period. Prior to the incident, emotional intensity was marked by volatility, frequent surges followed by equally sharp declines, suggesting that emotional expression was episodic, contingent on smaller news cycles or localized events. After May 25, 2020, this volatility gave way to a routinized pattern of high-intensity affective engagement. This affective stabilization is theoretically significant. Drawing on emotional contagion theory (Hatfield et al., 1994), it is clear that the sudden influx of emotionally charged content, particularly high-arousal emotions such as anger and outrage, spread rapidly through social media networks, facilitated by algorithmic amplification and high retweet activity. Yet unlike short-lived affective waves, the post-event stability of emotional intensity suggests a shift in the expressive norms of the discourse itself. This is consistent with what (Goodwin et al., 2001) describe as the institutionalization of affect: a process whereby emotional expression becomes durably embedded within the communicative structure of a movement, rather than remaining reactive or contingent.

In short, the data presented in Figures 2 and 3 demonstrate that the George Floyd incident functioned as more than a transient moral shock. It transformed the scale and structure of public political discourse, creating a sustained attention economy in which affective intensity was no longer an exceptional occurrence but a normative baseline. This shift established the conditions under which more specific emotional configurations, such as the dominance of anger, could emerge and consolidate.

While the initial behavioral and affective surge created the conditions for sustained emotional engagement, the George Floyd incident also restructured the composition of affect within BLM-related discourse. As shown in Figure 4, the dominant emotions prior to the incident were relatively balanced between anger (43.1%) and joy (41.9%), with sadness, fear, and other categories occupying only marginal shares. This balance shifted decisively in the post-event period, where anger rose to 51.7% of all dominant emotions, surpassing joy and becoming the most prevalent affective tone. Sadness also experienced a modest increase, while joy declined. This transition is further evidenced by the intensity trajectories of these three emotions in Figure 5. Before the incident, average emotional intensities were low and fluctuated in a narrow range, reflecting a discourse marked by fragmented affective expression. After May 25, 2020, anger surged and sustained that elevated level over an extended period. Joy, in contrast, declined in intensity and showed increased volatility, while sadness rose modestly but remained secondary. These shifts point toward fundamentally reordering emotional hierarchies within the protest discourse.

The ascendance of anger as both the most frequent and the most intensely expressed emotion aligns with prior work on the political function of anger in moments of moral crisis.

Unlike sadness, which may elicit inward reflection, or fear, which can produce avoidance, anger externalizes blame and demands redress, making it especially suited for mobilization in contentious political environments (Gamson, 1992; van Zomeren et al., 2012). Anger helps articulate a shared sense of grievance, identify perpetrators of injustice, and morally justify calls for action. In this context, its rise to dominance is not incidental but reflects a strategic affective alignment between emotional tone and the demands of mass protest.

Importantly, this shift should not be understood as a zero-sum replacement of positive emotions. As noted in Figure 4, joy remained a substantial presence even after the event. At first glance, this persistence may seem contradictory, especially given the overall rise in negative emotional expression. However, closer inspection reveals that many joy-labeled tweets were tied to celebrations of collective action, local protest victories, or expressions of solidarity. These findings suggest that joy did not reflect detachment from political grievance but rather functioned as a mobilizing affect in its own right, reinforcing group cohesion and sustaining engagement in outrage.

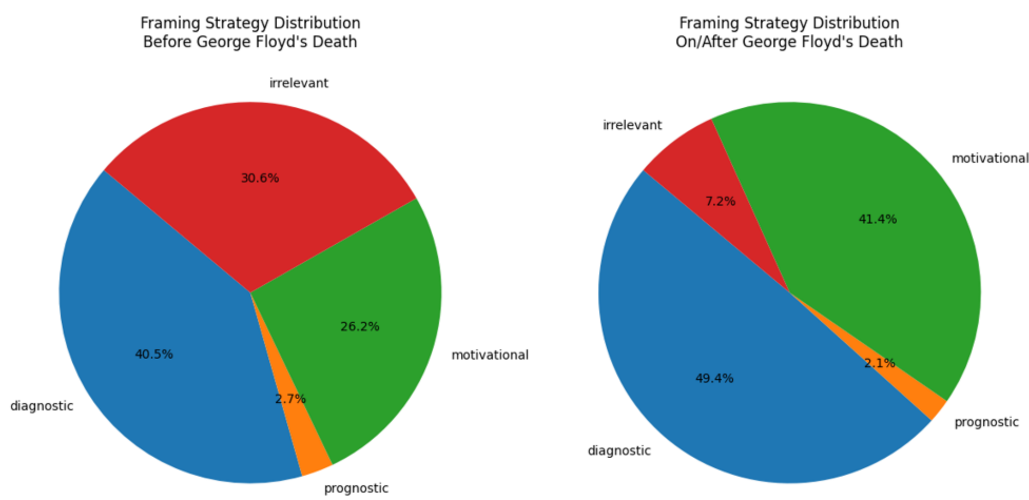

Figure 6: Distribution of Political Framing Strategies Before and After the Event

The findings presented in Figure 6 and Figure 7 indicate that the emotional transformations triggered by the George Floyd incident were not only structural and compositional, but also strategically coordinated with the discursive functions of political messaging. Following the incident, BLM-related discourse shifted toward emphasizing diagnostic and motivational framing: identifying sources of injustice and issuing calls to action. Simultaneously, distinct emotions became differentially aligned with these functions, forming stable patterns of emotion-frame coupling.

Emotion & Framing Strategy Association

Emotion	Diagnostic	Irrelevant	Motivational	Prognostic
Anger	73.4%	2.7%	22.7%	1.1%
Disgust	82.5%	10.4%	6.4%	0.6%
Fear	77.8%	4.1%	14.8%	3.3%
Joy	2.2%	14%	80.2%	3.7%
Sadness	60.3%	3.6%	34.4%	1.7%
Surprise	52.3%	35.6%	11.1%	1.1%

Note: Darker blue indicates higher percentage values

Figure 7: Emotion and Framing Strategy Association

Anger, the most prevalent post-event emotion, was overwhelmingly associated with diagnostic frames. As shown in Figure 7, 73.4% of anger-dominant tweets constructed narratives of blame and systemic harm, while another 22.7% issued direct motivational appeals. This distribution is theoretically consistent with the dual role of anger in protest communication: it sharpens boundaries between in-groups and perceived antagonists, while simultaneously galvanizing mobilization by presenting grievances as morally urgent and actionable (Gamson, 1992; van Zomeren et al., 2012). In this sense, anger clarified political stakes and escalated participation, anchoring the affective core of diagnostic discourse while energizing strategic messaging. In contrast, joy, despite being less prevalent, exhibited a radically different framing profile. Over 80% of joy-dominant tweets were associated with motivational frames, and only 2.2% with diagnostic ones. This suggests that joy functioned primarily to reinforce collective identity, celebrate symbolic victories, and sustain morale, rather than to define adversaries or attribute systemic blame. The endurance of joy in this affective environment may appear paradoxical. Still, its motivational alignment reveals an important function: it complemented anger by fostering emotional continuity and solidarity, providing participants with affective reinforcement amid the intensity of grievance-driven mobilization.

Sadness occupied a more ambivalent position. While 60.3% of sadness-dominant tweets were associated with diagnostic frames, a significant 34.4% also contributed to motivational framing. This dual alignment suggests that sadness played a more complex communicative role, reflecting grief and loss while reinforcing a moral imperative to act. This pattern is particularly evident in discourse centered on mourning victims, which often invoked their memory not as a passive gesture, but as a foundation for renewed calls to justice.

The coupling of emotion and framing strategy thus reveals a significant degree of affective-discursive alignment in the aftermath of the event. Emotions were not randomly distributed across rhetorical forms; they were selectively activated and expressed in ways that advanced specific political functions. These findings resonate with the broader literature on emotional culture in social movements, emphasizing that affect is not merely reactive but deeply embedded in the strategic repertoire of collective action (Goodwin et al., 2001).

The results suggest that the affective shift triggered by the George Floyd incident was not only structural and expressive but also strategically articulated. Anger anchored the discursive construction of injustice, joy sustained mobilization through affirmational appeals, and sadness enriched the moral depth of collective mourning. In this configuration, emotional expression functioned not simply as a byproduct of protest, but as a communicative infrastructure through which digital publics organized meaning, mobilized action, and sustained engagement over time.

While the preceding analysis has demonstrated that emotions such as anger, joy, and sadness became differentially aligned with political framing strategies, a closer inspection of their temporal evolution reveals an additional layer of transformation. Specifically, Figure 5 shows that the emotional impact of the George Floyd incident was not only immediate but extraordinarily persistent. The average intensity of anger surged in late May 2020 and remained elevated for over a year, accompanied by a corresponding decline in joy and a muted rise in sadness. Such a prolonged affective divergence is unusual in the context of digital discourse. Existing literature on emotional dynamics in crisis communication suggests that collective emotional responses on social media typically exhibit short-lived spikes rather than sustained changes. Emotional reactions to natural disasters, terrorist attacks, and political scandals have been shown to fade within days or weeks as attention shifts and emotional arousal subsides (Ferrara & Yang, 2015; Stieglitz et al., 2018). Even within the domain of political mobilization, surges in emotional intensity, such as those surrounding elections or high-profile debates, tend to normalize relatively quickly. Against this backdrop, the year-long elevation of anger and the concurrent shifts in other emotional baselines stand out as an empirical anomaly that merits further scrutiny.

Two potential explanations are considered. The first posits that the observed divergence may be an artifact of sample imbalance. Although the primary analysis relied on daily mean intensity scores, the number of posts per day increased dramatically following the event (see Figure 2). This discrepancy could artificially magnify post-event trends if larger sample sizes amplified rare but high-intensity expressions. To test this hypothesis, a normalized resampling procedure was conducted to ensure comparable daily sample sizes before and after the event. The resulting trends are visualized in Figure 8, which confirms that although some fluctuations were attenuated, the divergence in emotional intensity, particularly the sustained elevation of anger, remained robust.

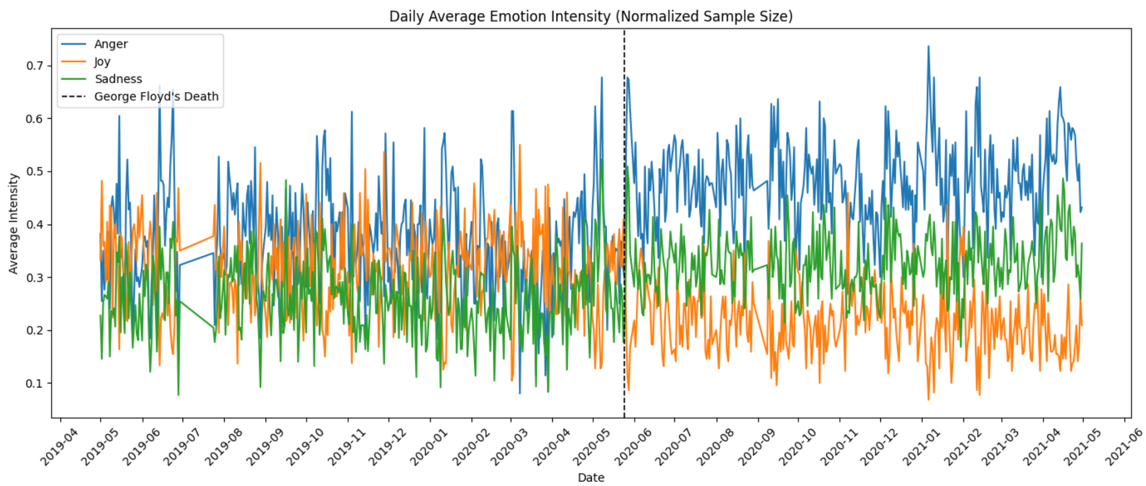

Figure 8: Daily Average Emotion Intensity with Normalized Sample Size

This result suggests that the persistent post-event divergence is not merely a byproduct of data density, but reflects a more profound structural shift in the affective dynamics of BLM-related discourse. Nevertheless, normalization alone does not explain the underlying mechanism: who sustained this elevated emotional intensity, and why? To address this question, the following section introduces a user-level comparison that distinguishes between new participants, users who first posted after the George Floyd incident, and existing users who were already active before the event. This distinction allows this study to examine whether the long-term affective divergence was broadly shared across the digital public or disproportionately driven by newly mobilized participants.

To further unpack the mechanism behind the sustained affective divergence, this analysis disaggregates post-event users into two categories: (1) *old users*, defined as those who had participated in BLM-related discourse prior to the George Floyd incident, and (2) *new users*, whose first appearance in the dataset occurred only after the event. By comparing emotional intensity trends across these two groups, we can assess whether the observed divergence was a shared shift among long-standing participants or primarily driven by newly mobilized voices.

Figure 9 presents the pre-event baseline (left) and the post-event emotional trends of old users (right). It shows that while anger increased following the event, the overall shift was moderate and less volatile than in the aggregate trends discussed earlier. In contrast, Figure 10 displays the emotional trajectory of new users after the incident. Here, the affective divergence is much more pronounced: anger intensity rose sharply and remained elevated, while joy declined and sadness increased more steadily. These findings suggest that newly activated users were disproportionately responsible for the sharp, prolonged divergence in emotional expression.

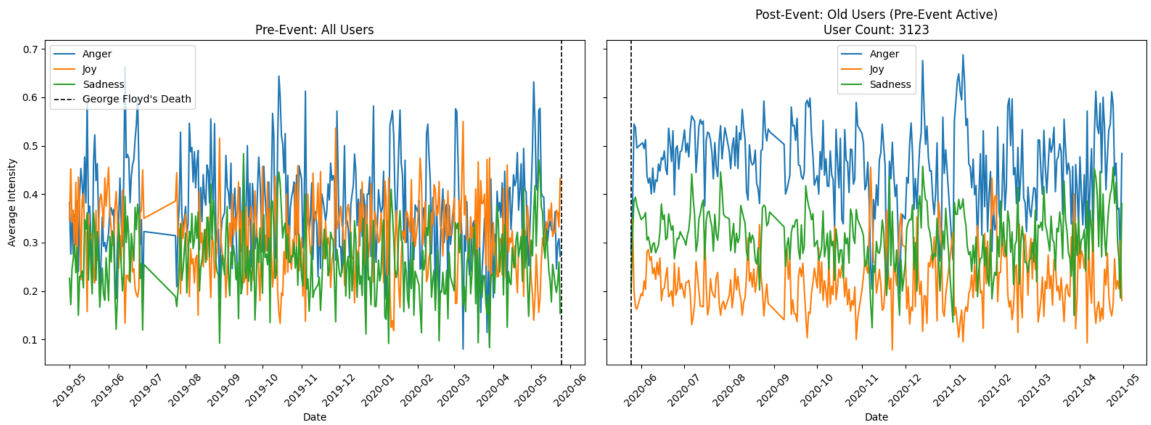

Figure 9: Emotion Intensity of Old Users (Pre-Event Active)

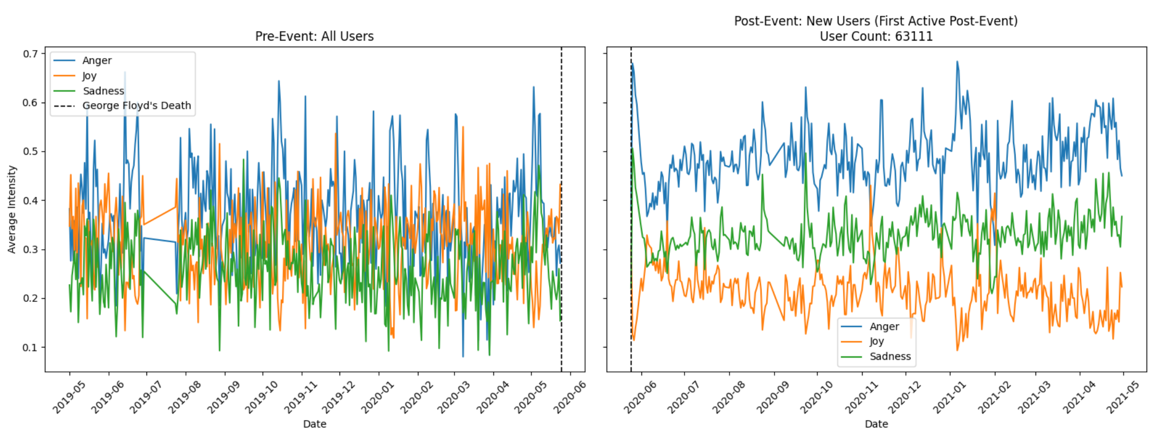

Figure 10: Emotion Intensity of New Users (First Active Post-Event)

To quantify these differences more rigorously, this study applies *Dynamic Time Warping (DTW)*, a method for measuring the similarity between two time series that may vary in speed or alignment. DTW is particularly suitable for analyzing digital emotional trajectories, as it accommodates misalignments and local warping in temporal patterns while preserving structural signals (Berndt & Clifford, 1994). Formally, the DTW distance between two time series $A = (a_1, a_2, ..., a_n)$ and $B = (b_1, b_2, ..., b_m)$ is defined as:

$$DTW(A, B) = \min_{w \in W} \sum_{(i,j) \in w} d(a_i, b_j) \tag{1}$$

where W is the set of all possible warping paths and $d(a_i, b_j)$ is typically the Euclidean distance between points a_i and b_j. In this case, the pre-event emotional trajectory is treated as the baseline. DTW distances are computed between it and each of the two post-event groups (old vs. new users), separately for anger, joy, and sadness.

Table 2: DTW Distances Between Pre-Event and Post-Event Emotional Trajectories

Emotion	**Pre vs. Old Users**	**Pre vs. New Users**	**Difference**
Anger	1.343	1.674	+0.331
Joy	1.176	1.295	+0.119
Sadness	0.953	1.118	+0.165

As shown in Table 2, new users consistently exhibit greater DTW distances from the pre-event baseline than old users do. This is most evident for *anger* (1.674 vs. 1.343) and *joy* (1.295 vs. 1.176), with a smaller but still noticeable difference for *sadness* (1.118 vs. 0.953). These results confirm that new users did not merely adopt existing emotional norms—they reshaped them. Their participation coincided with a structural reconfiguration of affective expression, marked by elevated anger and reduced joy, which came to define the post-event emotional landscape.

While the DTW results confirm that new users diverged more substantially from the pre-event emotional baseline than old users, it remains necessary to theorize why such differences emerged. Several mechanisms may account for this divergence. First, newly activated users likely entered the discourse under heightened emotional conditions. Focusing events such as the George Floyd incident do not merely attract attention; they redefine the emotional entry threshold. In contrast to long-standing participants who may have gradually habituated to the rhetorical norms of BLM discourse, new users were mobilized by a specific moral shock that intensified their affective baseline at the moment of entry. This distinction reflects what Papacharissi (Papacharissi, 2015) describes as affective synchronization: newly mobilized participants tend to align emotionally with the dominant affective flow of the triggering event, shaping their expressive behavior accordingly. Second, the asymmetry may also reflect differences in political socialization. Existing users are more likely to be movement insiders with prior exposure to BLM frames and norms, which may regulate their emotional expression through mechanisms of emotional discipline and collective identity reinforcement (Polletta, 2006). New users, in contrast, lack this history and may draw more heavily on raw affect and viral cues in shaping their posts, leading to more volatile and emotionally extreme patterns. Third, algorithmic amplification may further exacerbate this asymmetry. Social media platforms often reward emotional extremity, particularly high-arousal negative conten, with increased visibility (Brady et al., 2017). New users, especially those seeking resonance or community validation, may therefore be more susceptible to emotional mimicry and expressive escalation, amplifying affective divergence from the movement's historical norms. Taken together, these explanations suggest that the observed divergence is not merely a compositional artifact but the potential result of layered dynamics involving recruitment conditions, discursive learning, and platform incentives. Focusing events do not just change who participates, they can reshape the affective and strategic character of participation itself. However, confirming the causal mechanisms behind these hypotheses will require further experimental or longitudinal research.

4.2 RQ2: Structural Transformation of the information diffusion

4.2.1 Method: Regression of Emotion–Diffusion Linkage and Retweet Network Analysis

Having established that the George Floyd incident marked a turning point in the emotional tenor of BLM-related discourse, this section investigates whether it also altered the dynamics of information diffusion. Specifically, it asks whether emotionally intense content became more likely to reach broader audiences and whether the structural topology of the digital network supporting this diffusion underwent systematic change. To address these questions, this study operationalizes information diffusion in two complementary ways: first, as the visibility of individual tweets (measured by retweet count), and second, as the structural characteristics of the broader retweet network. The former captures micro-level audience responsiveness to emotional cues, while the latter reveals macro-level shifts in how influence and attention are distributed across participants. Combining these two perspectives allows this study to test whether emotionally charged messages gained more traction after the event and whether the underlying mechanisms of diffusion, namely, the network architecture, became more centralized and less reciprocal. At the micro level, the relationship between emotional intensity and diffusion reach is evaluated by binning tweets into deciles based on their predicted emotional intensity scores (ranging from 0.0 to 1.0, in 0.1 intervals). The average retweet count is computed for each bin, and this analysis is conducted separately for the pre-event and post-event periods. This enables a comparative study of whether the focusing event altered the functional link between affective salience and visibility. To enhance interpretability, the resulting trends are visualized as a two-line plot (Figure 11), and linear regression lines are fitted to both time periods to estimate baseline shifts and slope changes.

At the macro level, this study constructs retweet networks separately for the pre-event and post-event periods. In each network, nodes represent unique Twitter users, and directed edges represent retweet actions flowing from the original author to the retweeter. These networks are used to capture how information circulated through the BLM discourse community and whether this circulation structure changed in response to the George Floyd incident. To characterize these networks, three sets of metrics are computed. First, basic scale indicators, such as the number of nodes (users) and edges (retweet links), are used to assess the overall magnitude of participation and interaction. These metrics are visualized in Figure 12, providing a baseline for the network's expansion over time. Second, the concentration of information flow is assessed by calculating the proportion of total retweets generated by the top 1% most-retweeted accounts. This metric, presented in Figure 13, captures the extent to which influence is centralized among elite users. Third, the strength of relational ties is evaluated by calculating the proportion of "strong ties", defined as reciprocated retweet relationships, within the network. The relative change in these structural indicators is further summarized in Figure 14, which displays the growth rates of users, diffusion links, and strong ties between the two periods. Together, these micro- and macro-level methods enable an analysis of how emotion intensity interacts with visibility and how the broader structure of digital mobilization may become more or less centralized and asymmetric following a triggering event.

4.2.2 Results: Emotional amplification and Structural Centralization in Information Diffusion

To examine whether emotional intensity influenced the reach of tweets and whether the George Floyd incident altered this effect, this study computed the average retweet count for tweets within each 0.1-width bin of emotional intensity (ranging from 0.0 to 1.0). This analysis was conducted separately for the pre-event and post-event periods. As shown in Figure 11, the relationship between emotional intensity and average retweet count exhibits three notable shifts after the event.

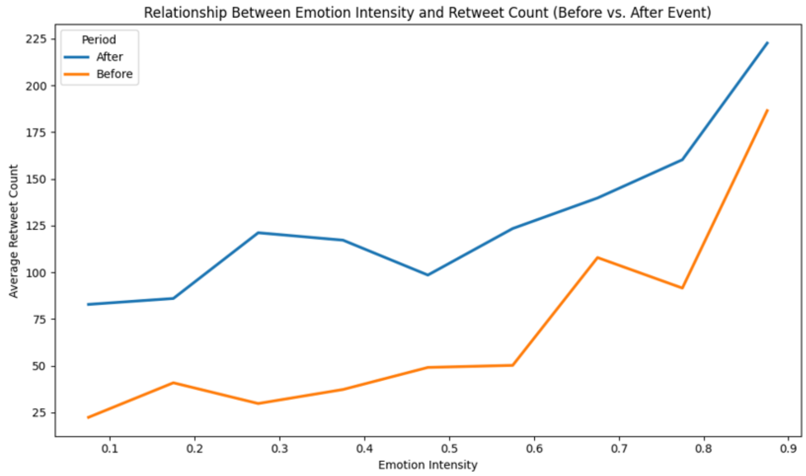

Figure 11: Relationship Between Emotion Intensity and Retweet Count (Before vs. After Event)

First, diffusion has a substantial baseline upliftacross all intensity levels. Tweets with low emotional salience (e.g., intensity ∼0.2) more than doubled in average retweet count after the event. This reflects a general expansion of audience attention after the George Floyd incident, suggesting that the communicative environment became more permeable to all BLM-related content, regardless of affective strength. Second, the slope of the intensity–diffusion relationship steepens considerably. Whereas tweets with moderate emotional intensity (e.g., ∼0.5) previously received around 50 retweets, their post-event counterparts attracted more than 120 on average. This steepening indicates that emotional expression became more effective at capturing attention, possibly due to algorithmic amplification mechanisms prioritizing emotionally resonant content. Third, the activation threshold for emotional diffusion appears to have lowered. Before the event, only tweets with high intensity (>0.7) achieved significant diffusion, while in the post-event period, even moderately emotional tweets gained wide visibility. This structural shift implies a recalibration in the attention economy, where emotional expressiveness has become a more potent and accessible resource for diffusion. These findings indicate that emotional intensity functioned not merely as a stylistic feature of protest discourse but as a key amplifier of visibility, one whose effectiveness was significantly enhanced following the focusing event. The George Floyd incident not only elevated the baseline attention to BLM content but also heightened the diffusion efficiency of emotional communication, laying the foundation for large-scale mobilization.

The emotional amplification effects described above occurred within a rapidly expanding and structurally transforming communication network. To understand the scope and nature of this transformation, this study constructed two directed retweet networks: one for the month prior to the George Floyd incident (April 25 – May 24, 2020), and another for the month following it (May 26 – June 25, 2020). Each node represents a Twitter user, and each directed edge represents a retweet from one user to another, capturing the flow of information. As shown in Figure 12, the post-event network experienced a dramatic expansion. The number of active users rose from approximately 172,000 to over 2.25 million, while the number of retweet links increased from around 207,000 to nearly 5 million. This over twentyfold increase reflects the mass digital mobilization triggered by the event.

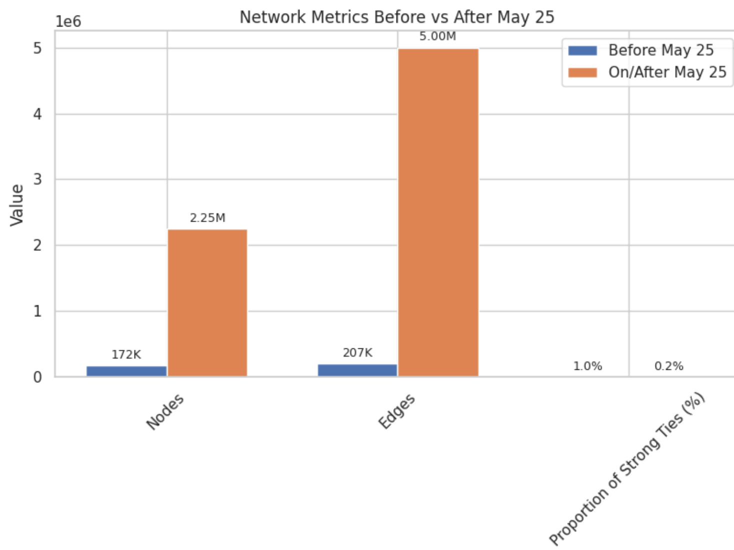

Figure 12: Network Metrics Before vs. After May 25

However, this growth was not simply a matter of scale. The internal structure of the network also shifted significantly. As Figure 13 shows, the concentration of retweet activity became markedly more unequal: before the event, the top 1% of users accounted for about 83% of all retweets, but after the event, that figure surged to over 95%. This sharp centralization suggests that information circulation became increasingly dominated by a small group of highly visible accounts.

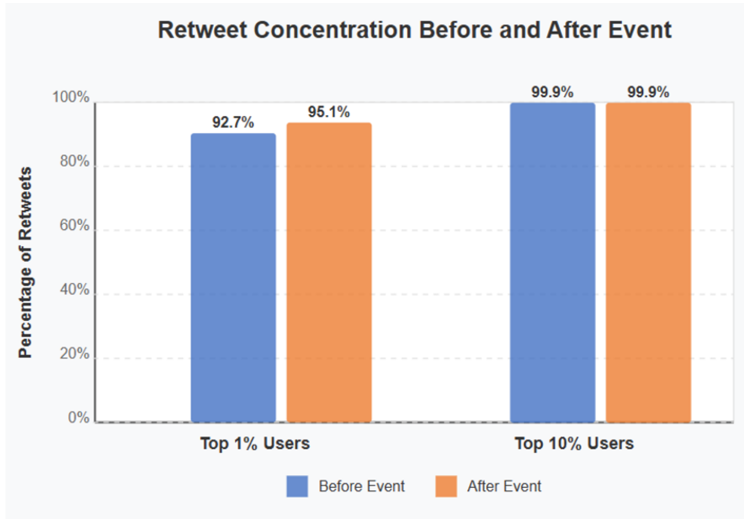

Figure 13: Retweet Concentration Before and After the Event

Further analysis confirms this trend. As visualized in Figure 14, although the number of users increased by over 1,200%, the number of retweet links grew by an even larger 2,317%. Yet, the proportion of strong ties—defined as reciprocal retweet relationships—plummeted from 1.0% to just 0.2%. This indicates that while participation broadened, mutual engagement thinned: users were more likely to amplify central accounts than to interact within peer networks.

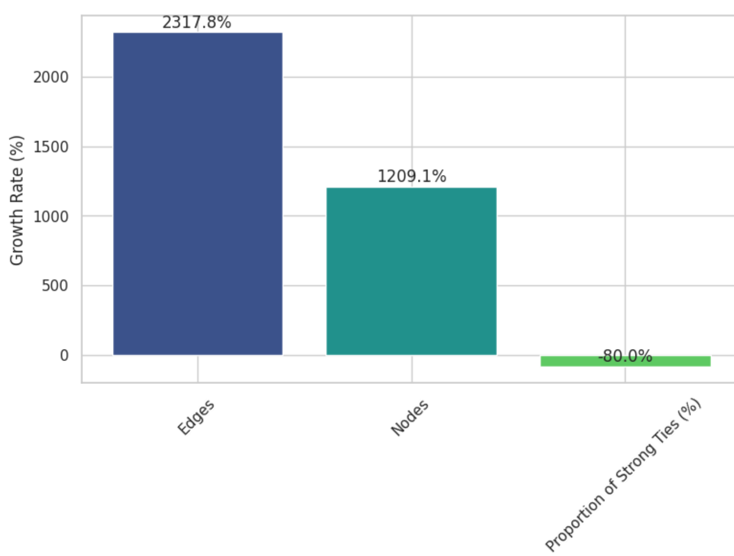

Figure 14: Growth Rate of Network Metrics

In sum, the retweet network evolved into a highly centralized, broadcast-like structure, wherein a small elite of users disseminated emotionally resonant content to a broad but relatively passive audience. This transformation increased the efficiency of emotional diffusion and reduced the relational cost of participation, enabling users to engage in the movement through low-effort amplification rather than reciprocal ties or two-way communication.

4.2.3 Analysis: New Participants and Influential Users Reshaped Information Diffusion

The results presented in Figure 11 indicate that emotional intensity became a significantly stronger predictor of message visibility after the George Floyd incident. While emotionally charged messages already enjoyed higher diffusion before the event, the amplification effect intensified in the post-event environment: even moderately emotional tweets began attracting disproportionately high levels of engagement. This structural recalibration of audience responsiveness aligns with the broader emotional shifts discussed in RQ1 as shown in the DTW analysis (Figure 2), new users diverged more sharply from the pre-event emotional baseline than old users, particularly in their elevated expression of anger and diminished display of joy. This influx of affectively synchronized participants (Papacharissi, 2015), mobilized under morally charged and emotionally saturated conditions, transformed the expressive norms and the receptive thresholds of the public sphere. The digital audience became more emotionally attuned and the discourse became more emotionally saturated.

This transformation likely reflects a compounding feedback loop: new users, entering with heightened emotional baselines, produced more intense content; this content, in turn, received greater amplification via algorithmic visibility dynamics (Brady et al., 2017), reinforcing both the expressive and diffusion logics of high-arousal participation. The resulting environment privileged emotionally expressive posts and emotionally expressive users, especially those aligned with the dominant affective flow triggered by the event.

Moreover, this shift was not merely expressive but infrastructural. The emotional activation of newly mobilized users reshaped the expressive landscape of BLM discourse and also fundamentally restructured the network architecture through which information circulated. As shown in Figures 12 and 13, the network experienced a dramatic expansion in size after the George Floyd incident, both in terms of participating users and retweet interactions. Information diffusion no longer relied on distributed relational networks. It became centralized around a small group of affective broadcasters whose messages were repeatedly amplified by newly mobilized audiences. Specifically, efore the event, the BLM retweet network exhibited characteristics of moderate decentralization: retweets were more evenly distributed, and a non-negligible share of interactions occurred between mutually engaged users. Following the event, this pattern changed dramatically. The proportion of strong ties—edges formed through reciprocal or repeated interactions—declined from 1.0%

to just 0.2%, and over 95% of all retweets became concentrated in the top 1% of accounts (Figure 13). This suggests that newly activated users did not engage in mutual discourse or peer deliberation; they overwhelmingly amplified messages from elite nodes, typically verified users, prominent activists, or institutional accounts. This transformation aligns with what scholars have described as the broadcast logic of digital mobilization (Bennett & Segerberg, 2012). Rather than forming deliberative communities or horizontal alliances, large-scale digital movements often rely on affective dissemination through high-visibility hubs. In this configuration, emotional content does not spread through relational ties, but through repeated retweets of emotionally resonant messages from central actors. The result is a core-periphery topology: a handful of visible nodes define the narrative, while a mass of peripheral users engage primarily through amplification. This structural pattern is reinforced by the emotional logic discussed earlier. New users, entering with heightened affective intensity, did not engage through reciprocal discussion; they contributed by boosting emotionally charged content from already-visible figures. This lowered the threshold for participation and simplified the communicative labor required to join the movement. In effect, retweeting replaced deliberation as the dominant mode of engagement. A structural dominance of centralized broadcast thus matched the expressive dominance of emotion.

This shift also signals a change in the mobilization of digital platforms. In the wake of the George Floyd incident, the communicative architecture of Twitter privileged one-way flows from influential accounts to emotionally primed audiences. This configuration can generate rapid attention and visibility, highlighting the power of emotional amplification. Together, the transformation of expressive norms and structural topology suggests that focusing events do not simply inject emotional content into stable networks; they can reconstitute the mobilization infrastructure by changing who participates, what they express, and how their expressions circulate.

4.3 RQ3: Mechanisms of Cross-Community Diffusion

4.3.1 Method: Community Detection and Edge Classification

To investigate the structural mechanisms of cross-community information diffusion in the wake of the George Floyd incident, this study first constructs a directed retweet network from the post-event dataset, where each node represents a unique Twitter user and each edge denotes a retweet relationship from one user to another. To capture the meso-level structure of discourse, the network is partitioned into communities using the Louvain method for community detection (Blondel et al., 2008). This algorithm identifies clusters of nodes by optimizing a modularity function Q, which evaluates the density of links within communities relative to links between communities. Formally, modularity is defined as:

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j) \tag{2}$$

where A_{ij} is the adjacency matrix of the graph, k_i and k_j are the degrees of nodes i and j, m is the total number of edges, and $\delta(c_i, c_j)$ is an indicator function equal to 1 if nodes i and j are assigned to the same community and 0 otherwise. The Louvain method begins by assigning each node to its own community, and then iteratively aggregates communities to maximize the overall modularity score. This approach is computationally efficient and well-suited for large-scale networks such as those derived from Twitter data, where users self-organize into dense interaction clusters.

After community detection, this study categorizes each edge based on two orthogonal criteria: *tie strength* and *community boundary*. Tie strength is operationalized using reciprocity: an edge is defined as a *strong tie* if it is reciprocated (i.e., both $u \to v$ and $v \to u$ exist in the graph), and as a *weak tie* otherwise. Community boundaries are determined by comparing the detected community assignments of the source and target nodes. If both nodes belong to the same community, the edge is labeled as *intra-community*; if they belong to different communities, it is labeled as *bridge*. This yields four types of retweet interactions:

	Within community	**Between communities**
Strong (two-way)	**Strong intraconnection**	**Strong bridge connection**
Weak (one-way)	**Weak intraconnection**	**Weak bridge connection**

Table 3: Structural edge type classification

This classification allows the study to examine how different types of network ties contribute to the transmission of emotional content, particularly across community boundaries, which are assumed to be potential sites of discursive divergence or convergence.

Once structural labels were assigned to each retweet edge, the next analysis stage examined whether emotional content systematically varies across structural categories. Given that retweets do not typically introduce new textual content, the emotional attributes of each edge were inferred directly from the original tweet being retweeted. This design assumes that the affective and discursive signals conveyed by a retweet are transmitted primarily through the content of the source message. Accordingly, each edge inherits three key annotations from the source tweet: (1) the dominant emotion, selected from a six-category taxonomy of basic emotions (anger, joy, sadness, fear, disgust, surprise); (2) the emotion intensity, defined as a normalized scalar ranging from 0 to 1 reflecting the confidence level of the emotion classifier; and (3) the framing strategy, classified as either diagnostic, motivational, or none, based on an independent classification pass.

These emotion-related variables were embedded as edge-level attributes in the graph, enabling the integration of structural and expressive dimensions into a unified network representation. Based on this representation, the study operationalized a three-part analytical framework to explore the structural mechanisms of affective diffusion. Each component of this framework was designed to assess a distinct facet of the emotional dynamics shaping inter-community information flow.

First, a macro-level assessment of community integration was conducted to contextualize these edge-level patterns. Two metrics were calculated from the community-labeled network: (1) the total number of communities detected using the Louvain algorithm, and (2) the number and proportion of communities connected to at least one other community via a weak-tie bridge. This measure captures the extent to which emotionally annotated weak ties function as structural integrators in the post-event network.

Second, to assess whether structurally peripheral or bridging connections are more likely to transmit emotionally salient content, the study compared the distribution of emotion intensity across the four previously defined edge types: weak intra-community, strong intra-community, weak bridge (cross-community, unidirectional), and strong bridge (cross-community, reciprocal). For each edge, the emotion intensity score was extracted and grouped by edge category. The resulting distributions were visualized using boxplots, and descriptive statistics were computed for each group. This analysis directly tests the hypothesis that weak ties, especially those linking different communities, are more or less likely to carry emotionally charged content, thereby discussing their mobilization potential. To test whether mean emotion intensity varied significantly across edge types, a one-way analysis of variance (ANOVA) was performed. ANOVA evaluates the null hypothesis that all group means are equal by comparing the between-group variance to the within-group variance. The test statistic is calculated as:

$$F = \frac{MS_{between}}{MS_{within}} = \frac{\sum_{i=1}^{k} n_i(\bar{X}i - \bar{X})^2/(k-1)}{\sum i = 1^k \sum_{j=1}^{n_i}(X_{ij} - \bar{X}_i)^2/(N-k)} \tag{3}$$

where k is the number of groups (here, k=4), ni is the number of observations in group i, Xi is the mean emotion intensity in group i, X is the grand mean, and N is the total number of edges with emotion annotations. A significant F-statistic indicates that at least one group mean differs from the others. To further identify which specific group differences were statistically significant, a post-hoc pairwise comparison was conducted using Tukey's Honestly Significant Difference (HSD) test. Tukey's HSD controls for family-wise error rate in multiple comparisons and computes a standardized test statistic:

$$q = \frac{|\bar{X}_i - \bar{X}_j|}{\sqrt{MSE/n^*}} \tag{4}$$

where MSE is the mean squared error from the ANOVA model and n is the harmonic mean of the sample sizes for the two groups being compared. The q-statistic is evaluated against the studentized range distribution to determine significance. Together, the ANOVA and Tukey HSD tests allow for robust inference about whether emotion intensity varies systematically across structural connection types, with particular attention to whether different types of ties are more or less likely to carry emotionally charged content than other connection types.

Third, the compositional structure of emotional content was examined by analyzing the distribution of dominant emotions and framing strategies across edge types. For each structural category, the proportion of edges carrying a given dominant emotion (e.g., anger, joy) or rhetorical frame (e.g., diagnostic) was calculated. These proportions were visualized using grouped bar charts, allowing for comparative assessments of which emotional and discursive patterns are more common in particular edge types. This analysis is particularly useful for identifying whether specific emotions and strategic framings are disproportionately transmitted through different types of edges. If so, this would suggest that affective content is not evenly distributed, but structurally sorted in ways that prioritize the cross-group diffusion of emotionally and rhetorically mobilizing narratives.

These three methods, intensity comparison, emotion and framing composition, and community integration, allow for a layered examination of how emotional salience interacts with structural constraints in the diffusion of protest discourse. By linking micro-level expressions to micro- and macro-level network features, this section operationalizes a structural-affective perspective on digital mobilization, testing whether specific forms of emotional communication are structurally privileged in traversing social boundaries following a focusing event.

4.3.2 Results: Weak Ties Transmit Anger Across Communities, Strong Ties Mobilize Collective Action

To contextualize the diffusion of emotional content within the retweet network, this section first examines the structural configuration of communities and the distribution of edge types linking them. Applying the Louvain algorithm for modularity optimization identified 2,024 distinct communities in the post-event network. Despite this high level of fragmentation, most communities were not isolated: as shown in Figure 15, 1,775 communities (approximately 87.7%) were connected to at least one other community via a weak-tie bridge. This suggests that cross-community retweeting was not an incidental occurrence, but rather a pervasive feature of the network's structure.

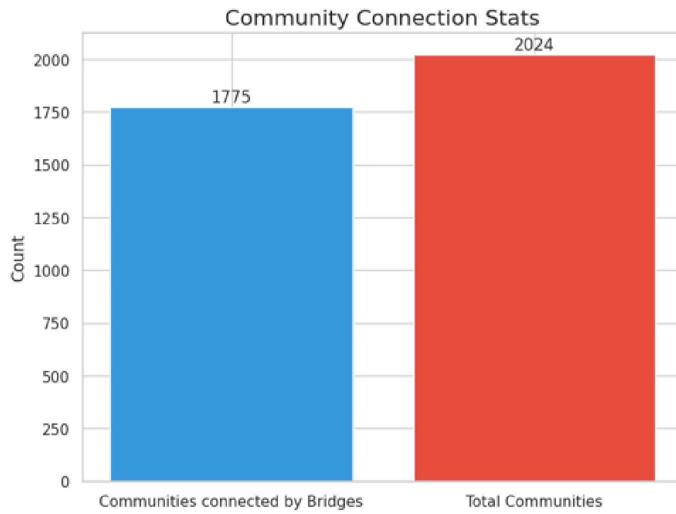

Figure 15: Community Connectivity via Weak Bridges

In addition to community-level connectedness, the analysis also classified all edges into one of four structural types, strong intra-community, weak intra-community, strong bridge, and weak bridge, based on reciprocity and community membership. As shown in Figure 16, weak intra-community ties accounted for the vast majority of emotionally annotated retweet connections (80.8%), followed by weak bridges (19.1%). Strong ties, whether intra- or inter-community, were exceedingly rare, together comprising less than 0.2% of the total. This distribution indicates that unidirectional retweeting, especially among users within the same community, is the dominant mode of affective information flow. The presence of weak bridges highlights their structural importance in linking otherwise segregated communities.

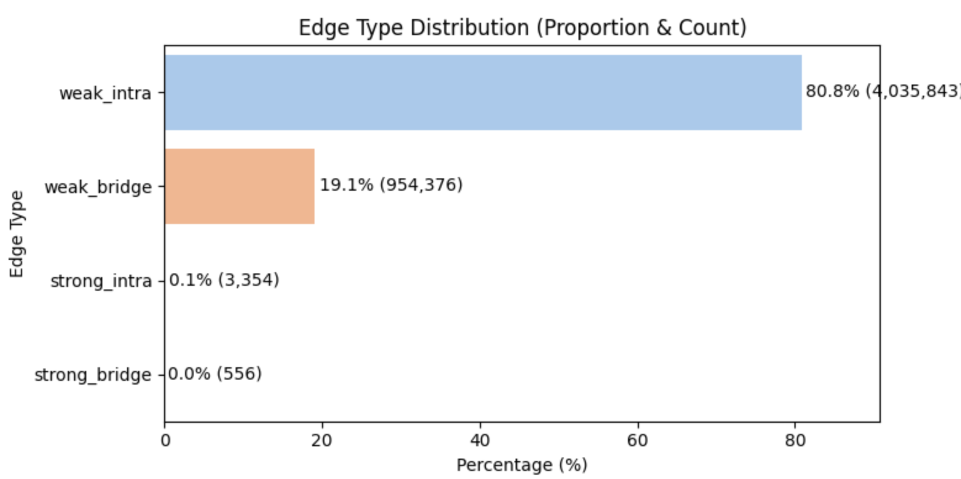

Figure 16: Distribution of Edge Types by Proportion and Count

These descriptive results provide critical context for the analyses that follow. The widespread presence of weak-tie bridges between communities suggests that the network was structurally conducive to cross-group diffusion. Moreover, the dominance of weak ties in general implies that most emotional content was not exchanged through mutual, deliberative interactions, but rather through unidirectional amplification.

To examine whether different structural connection types are associated with distinct emotional patterns, Figure 17 displays the distribution of dominant emotions across the four edge types. As expected, *anger* dominates in all categories, particularly in weak ties, indicating that high-arousal negative emotions are widely used in both intra- and inter-community diffusion. Notably, strong bridge edges exhibited a reversed pattern, where *joy* (43.3%) surpassed anger (41.0%). This inversion may suggest a distinctive mobilization role of reciprocated ties that span community boundaries, possibly aligned with motivational framing strategies a hypothesis further examined in subsequent sections.

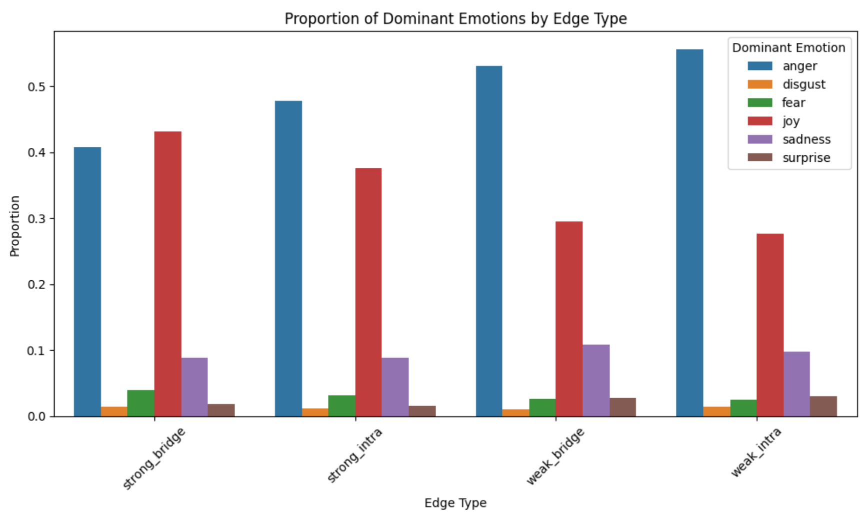

Figure 17: Distribution of Dominant Emotions by Edge Type

Figure 18 presents a boxplot comparing emotion intensity across the four edge types. The results show that both weak intra-community and weak bridge ties carry higher average emotion intensity than their strong counterparts. This pattern suggests that unidirectional ties—whether intra- or inter-community—are more likely to transmit emotionally salient content. One interpretation is that weaker structural ties require higher affective resonance to sustain engagement, especially when lacking reciprocal reinforcement. Conversely, strong ties may exhibit lower intensity due to higher baseline trust or sustained relational familiarity, reducing the need for emotionally charged content.

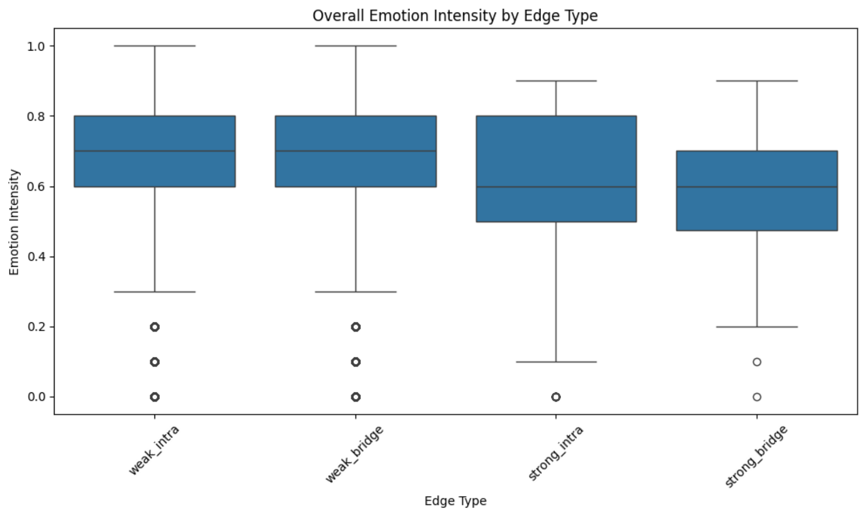

Figure 18: Distribution of Emotion Intensity by Edge Type

To statistically assess these differences, a one-way analysis of variance (ANOVA) was conducted with edge type as the independent variable and emotion intensity as the dependent variable. The test revealed a highly significant effect, $F(3, N) = 223.29$, $p < .001$, indicating that emotional intensity varies systematically by structural category. Table 4 summarizes the mean, standard deviation, and sample size for each edge type.

Table 4: ANOVA Summary and Descriptive Statistics by Edge Type

Edge Type	**Mean**	**SD**	**N**	**Significance**
Strong Bridge	0.608	0.192	556	***
Strong Intra	0.584	0.189	3,354	***
Weak Bridge	0.634	0.182	954,376	***
Weak Intra	0.629	0.185	4,035,843	***

Note. $F(3, N) = 223.29$, $p < .001$. *** indicates $p < .001$ in all pairwise comparisons.

To validate the robustness of these group-level differences, a post-hoc Tukey HSD (Honestly Significant Difference) test was performed. This method accounts for multiple comparisons by controlling the family-wise error rate and identifies which group differences are statistically significant. As shown in Table 5, all pairwise comparisons between edge types yielded significant differences ($p < .05$ or $p < .001$). Larger differences were observed between strong and weak ties, and between bridge and intra-community edges, suggesting that both tie strength and community boundary status play roles in shaping the intensity of transmitted emotions.

Table 5: Tukey HSD Pairwise Comparisons of Emotion Intensity

Group 1	**Group 2**	**Mean Diff**	**95% CI**	**Significance**
Strong Intra	Strong Bridge	0.0238	[0.0022, 0.0453]	*
Strong Intra	Weak Intra	0.0448	[0.0249, 0.0647]	***
Strong Intra	Weak Bridge	0.0498	[0.0298, 0.0697]	***
Strong Bridge	Weak Intra	0.0211	[0.0129, 0.0292]	***
Strong Bridge	Weak Bridge	0.0260	[0.0179, 0.0341]	***
Weak Intra	Weak Bridge	0.0050	[0.0044, 0.0055]	***

Note. * $p < .05$, *** $p < .001$.

Figure 19 examines the distribution of framing strategies across the four edge types to complement the analysis of dominant emotions. Diagnostic frames, which focus on problem diagnosis and the attribution of blame, were especially prevalent among weak ties, comprising 51.6% of weak bridge edges and 55.3% of weak intra-community connections. This mirrors the previously observed dominance of anger in these categories and suggests that weak ties are channels for emotional intensity and grievance-oriented discourse that identifies perceived injustices and their causes. Such diagnostic framing may be crucial in framing shared grievances and fostering cognitive alignment across users and communities.

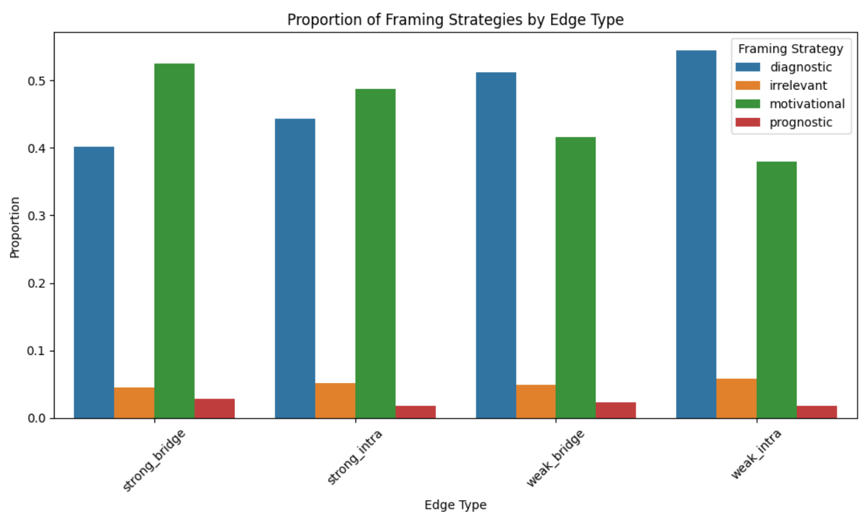

Figure 19: Distribution of Framing Strategies by Edge Type

In contrast, motivational frames, which aim to inspire collective efficacy, urgency, and the will to act—were notably more frequent in strong ties. In particular, strong bridge edges exhibited the highest proportion of motivational framing (52.6%), slightly surpassing diagnostic frames (40.1%). This is consistent with the earlier finding that joy was more prominent than anger among strong bridge connections, and suggests that reciprocal, cross-community ties may serve a different function: not just spreading grievances but reinforcing hopeful and action-oriented narratives. These emotionally positive and mobilizing messages may strengthen solidarity and commitment across group boundaries, indicating a dual structure where weak ties help diffuse diagnostic grievances, while strong ties consolidate motivational appeals.

4.3.3 Analysis: Affective Compensation Facilitated Outrage, Reciprocal Trust Supported Motivation

The first two results reveal the extent of community fragmentation and the dominance of weak ties in emotional diffusion. At the macro level, identifying over two thousand modularity-defined communities suggests a highly clustered communication environment. However, the fact that nearly 88% of these communities are linked by at least one weak-tie bridge indicates that this fragmentation is not absolute. Instead, the network exhibits what could be described as a permeable modularity: while discourse is locally clustered, it is structurally poised for inter-community transmission, which often takes the form of ”segmented but porous” structures where distinct communities can still coordinate via sparse but crucial connective links . This structure enhances the interpretive value of the second result: the overwhelming dominance of weak ties, especially within-community unidirectional edges (80.8%), followed by weak bridges (19.1%). Such a distribution reflects the broadcast logic of digital mobilization, where the majority of users amplify high-visibility messages rather than engage in reciprocal deliberation. The rarity of strong ties (both intra- and inter-community) further underscores that emotional communication in this context is asymmetric, mainly non-dialogic, and it flows from prominent sources to dispersed audiences with minimal mutual engagement. Crucially, the presence of weak bridges, despite their minority status in absolute terms compared to the weak intraconnection, plays an outsized functional role. These ties are structurally necessary for cross-group affective diffusion. Without them, even emotionally charged messages would remain trapped within community silos. The combination of a high number of weak intra-community ties and the structural permeability created by weak bridges suggests a diffusion regime that is locally dense but globally connected, optimizing for scale and speed of amplification rather than for relational trust or consensus-building.

These structural features challenge classical assumptions from collective action theory which emphasize strong ties as the drivers of mobilization (McAdam, 1986). In contrast, the post-event Twitter network observed here aligns more closely with theories of networked affect (Papacharissi, 2015), where emotionally charged content spreads through thin but wide pathways, shaping collective orientations without requiring deep relational embedding.

The third result provides empirical support for the hypothesis that emotionally intense content is not evenly distributed across the network, but varies systematically by structural tie type. Specifically, weak ties within and across communities consistently exhibit higher average emotion intensity than their strong-tie counterparts. This finding is statistically robust, confirmed through ANOVA and Tukey HSD post-hoc tests.

This pattern suggests a functional distinction in the different types of ties' roles in digital mobilization. Weak ties function as affective amplifiers, transmitting high-arousal messages likely to trigger attention and engagement across larger audiences. This is especially pronounced in weak bridges, which connect communities and do so through emotionally charged, unreciprocated content. From a network communication perspective, this confirms that weaker structural connections must compensate for the absence of relational reinforcement through increased expressive salience. That is, when users lack trust, familiarity, or prior interaction, as is typical across weak ties, the affective charge of a message becomes critical to its propagation. Conversely, on average, strong ties, particularly those within the same community, show lower emotional intensity. One interpretation is that these ties operate under higher baseline trust or mutual understanding, which reduces the need for extreme affect to achieve attention or alignment. Users in such relationships may share context or prior consensus, enabling communication through more measured emotional registers.

The fourth and fifth results, on dominant emotion and framing strategy distributions, add nuance to this intensity-based account by highlighting how types of emotions and rhetorical structures vary by connection. The dominance of anger in weak ties (both intra- and inter-community) supports the hypothesis that diagnostic framing, focused on naming injustices and attributing blame, travels most effectively through sparse, unidirectional channels. Anger is a mobilizing but divisive emotion, often effective at drawing attention but less effective at sustaining cooperation (Brady et al., 2017).

In contrast, strong bridge connections exhibit a reversed emotional pattern, with joy slightly surpassing anger and motivational framing dominating over diagnostic. This combination implies a different communicative function: instead of only broadcasting grievance, strong bridges may serve to consolidate solidarity and encourage collective efficacy. Their reciprocal nature allows for mutual validation, while their cross-community position enables them to frame action-oriented narratives capable of fostering commitment across divides. These ties, though rare, appear to play a cohesive oriented role in movement discourse, in contrast to the reactive, attention-driving function of weak ties.

These findings suggest the emergence of a dual-pathway model of affective diffusion in the context of post-event digital activism. Rather than operating as a uniform emotional field, the network reveals a patterned alignment between structure and affect. Weak ties, particularly weak bridges connecting distinct communities, function as key conduits for high-intensity emotional content, especially messages framed diagnostically to highlight grievances and attribute responsibility. These ties, though typically unidirectional, are crucial in facilitating the wide dissemination of emotionally resonant narratives. Their structural position allows them to reach broader audiences and initiate engagement across boundaries, making them particularly effective in catalyzing initial awareness and emotional activation.

Though typically unidirectional, these ties In contrast, strong ties, especially those bridging across community boundaries, are more likely to carry messages framed in motivational terms, emphasizing hope, efficacy, and collective resolve. These ties are often reciprocal and situated within contexts of sustained engagement, which may support the reinforcement of emotionally positive and action-oriented narratives. This suggests that strong ties contribute not only to the transmission of mobilizing sentiment but also to the integration and maintenance of cross-group solidarity, helping to sustain collective momentum beyond the initial emotional spark. This alignment between structure and emotion underscores a potential insight: affective expression in digital protest networks is not merely an attribute of content, but is structurally conditioned by the pathways through which it travels. Tie strength and community configuration jointly shape the reach of information and its emotional character. These findings extend networked publics and emotional mobilization theories by demonstrating that emotional salience is differentially distributed across network topologies, with distinct structural roles for diagnostic outrage and motivational solidarity. In doing so, the study offers a more nuanced account of how emotions circulate in digitally mediated movements, highlighting not only what is felt but also where and how it is transmitted.

5 Conclusion

This study offers an integrated examination of how focusing events, such as the murder of George Floyd, reshape both the emotional and structural dimensions of digital mobilization. Following the George Floyd incident, this study observed a sharp and sustained increase in the volume of political discourse, accompanied by a pronounced rise in anger- and sadness-laden narratives. Importantly, the influx of newly participating users after the event was disproportionately engaged with emotionally intense messages, suggesting that affective resonance played a key role in drawing attention and participation. In parallel, the structure of the retweet network shifted from a relatively decentralized configuration to a more centralized, influencer-driven model, where a smaller set of highly connected nodes gained prominence in disseminating content.

Building on these findings, this study further reveals that emotional surges and structural shifts are not isolated processes, but mutually reinforcing dynamics. The escalation of collective anger and sadness energized information flows across the network, while the evolving network topology, particularly the emergence of weak-tie bridges, facilitated the cross-community spread of emotionally charged narratives. This alignment between affective expression and structural configuration suggests that digital protest networks do not merely reflect what is being felt, but actively organize how and where such emotions are transmitted. Rather than functioning as a neutral backdrop, the network architecture amplifies certain types of affective content, embedding emotional salience into the pathways of digital mobilization.

From a theoretical perspective, these findings contribute to bridging two longstanding strands of literature that often remain disconnected. While previous studies have emphasized either the role of emotional narratives in mobilization (Brady et al., 2017; Hatfield et al., 1994) or the structural affordances of weak ties in expanding information reach (Bakshy et al., 2012; Rainie & Wellman, 2012), this research demonstrates that these mechanisms are not independent. In the context of a focusing event, emotional content and network structures co-evolve, shaping what messages are transmitted and how they travel across fragmented digital spaces. Notably, the results suggest that weak ties on social media platforms do more than transmit information, they carry emotional energy across communities, where internal strong ties embed and reinforced these narratives.

Methodologically, this project advances mixed computational approaches by integrating large-scale language model annotation with network topology analysis. Leveraging structured prompt design and manual validation, the study achieves high-quality annotations for both emotional expression and discursive framing, capturing subtle patterns in politically charged discourse. Coupled with retweet network reconstruction, this design offers a systematic lens to observe how expressive shifts and relational structures unfold simultaneously in response to exogenous shocks.

Empirically, the findings challenge conventional assumptions about the limitations of weak ties in sustaining engagement. In the high-velocity environment of digital mobilization, weak-tie bridges function not merely as peripheral conduits of information but as active agents that facilitate emotional resonance across community boundaries. This hybrid role of weak and strong ties highlights a layered diffusion process: weak ties enable rapid exposure to new audiences, while strong ties within communities stabilize and deepen emotional commitments. Such insights enrich existing models of online collective action by foregrounding the interaction between network structure and emotional dynamics.

However, several limitations should be acknowledged. First, although the study captures large-scale emotional and structural shifts, it does not employ a controlled counterfactual design. Without a defined control group or pre-registered experimental conditions, causal inferences regarding the effects of the focusing event remain suggestive rather than definitive. Second, the network analysis is confined to retweet interactions, omitting other relational signals such as replies, mentions, or follower networks, which may offer additional perspectives on discourse diffusion and engagement pathways. Third, the current network model primarily observes single-step interactions between users. The inability to reconstruct deeper, multi-hop cascades limits fully capturing the layered dynamics of information spread, especially for identifying how content moves through complex chains of intermediaries over time.

Despite these constraints, this study advances the understanding of digital mobilization by revealing how emotional intensity and network architecture co-construct the pathways of collective expression. The findings underscore the importance of viewing emotional dynamics and network structures as intertwined processes, particularly in focusing events that activate latent social tensions. Future research could extend this approach by incorporating multi-modal interaction data, developing causal identification strategies, and reconstructing deeper network hierarchies to illuminate further the mechanisms driving large-scale digital mobilization.

Data and Code Availability Statement

The Twitter (X) data collected and analyzed in this study have been deposited at Here. In accordance with Twitter's privacy policy, this dataset do not share raw tweet text or user information. Instead, processed and anonymized data, including tweet IDs, emotional and framing annotations, and network metadata, are made available to support replication. All analysis code is included in the same GitHub repository.

Here is the GitHub repository.

References

Alam, F., Ofli, F., & Imran, M. (2020). Cicero: Analyzing crisis communication patterns in social media during australian bushfires. *Proceedings of the 2020 International Conference on Social Media & Society*, 46–56.

Bail, C. A. (2016). Combining natural language processing and network analysis to examine how advocacy organizations stimulate conversation on social media. *Proceedings of the National Academy of Sciences*, *113*(42), 11823–11828.

Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012). The role of social networks in information diffusion. *Proceedings of the 21st international conference on World Wide Web*, 519–528.

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, *26*(10), 1531–1542.

Benford, R. D., & Snow, D. A. (2000). Framing processes and social movements: An overview and assessment. *Annual Review of Sociology*, *26*(1), 611–639.

Bennett, W. L., & Segerberg, A. (2012). The logic of connective action: Digital media and the personalization of contentious politics. *Information, Communication & Society*, *15*(5), 739–768. https://doi.org/10.1080/1369118X.2012.670661

Benski, T., & Fisher, E. (2014). Emotional performances on social media: The role of affective practices in constructing collective action frames. In *The routledge companion to media and activism* (pp. 181–190). Routledge.

Berger, J., & Milkman, K. L. (2012). What makes online content viral? *Journal of Marketing Research*, *49*(2), 192–205.

Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining (KDD)*, 359–370.

Bimber, B., Flanagin, A. J., & Stohl, C. (2005). Reconceptualizing collective action in the contemporary media environment. *Communication Theory*, *15*(4), 365–388.

Binder, C. R. (2020). Climate change and emotion: Perspectives from social science. *Current Opinion in Behavioral Sciences*, *42*, 85–90.

Birkland, T. A. (1998). *After disaster: Agenda setting, public policy, and focusing events*. Georgetown University Press.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), P10008.

Bode, L. (2017). Closing the gap: Gender parity in political engagement on social media. *Information, Communication & Society*, *20*(4), 587–603.

Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, *114*(28), 7313–7318.

Buchanan, L., Bui, Q., & Patel, J. K. (2020). Black lives matter may be the largest movement in u.s. history. *The New York Times*.

Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, *113*(3), 702–734.

Diani, M. (2003). *Introduction: Social movements, contentious actions, and social networks: 'from metaphor to substance'?* Oxford University Press.

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

Ekman, P. (1992). *An argument for basic emotions* (Vol. 6). Taylor & Francis.

Ferrara, E., & Yang, Z. (2015). Measuring emotional contagion in social media. *PLoS one*, *10*(11), e0142390.

Fraustino, J. D., Liu, B. F., & Jin, Y. (2012). Social media use during disasters: A review of the knowledge base and gaps. *National Consortium for the Study of Terrorism and Responses to Terrorism*, *59*.

Freelon, D., McIlwain, C. D., & Clark, M. D. (2016). Beyond the hashtags: ferguson, blacklivesmatter, and the online struggle for offline justice. *Center for Media & Social Impact*.

Freelon, D., McIlwain, C. D., & Clark, M. D. (2018). *Beyond the hashtags: #Ferguson, #blacklivesmatter, and the online struggle for offline justice*. Oxford University Press.

Gamson, W. A. (1992). *Talking politics*. Cambridge University Press.

Gil de Zúñiga, H., Jung, N., & Valenzuela, S. (2012). Social media use for news and individuals' social capital, civic engagement and political participation. *Journal of Computer-Mediated Communication*, *17*(3), 319–336.

Gladwell, M. (2010). Small change: Why the revolution will not be tweeted. *The New Yorker*, *4*(2010), 42–49.

Gong, X., Zhang, W., & Li, X. (2023). Triggering collective emotions: How focusing events drive emotional peaks on social media. *Social Media + Society*, *9*(1), 1–13.

González-Bailón, S. (2020). *Social networks and protest participation: Evidence from the chilean student movement*. Cambridge University Press.

González-Bailón, S., Borge-Holthoefer, J., Rivero, A., & Moreno, Y. (2013). Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist*, *57*(7), 943–965.

González-Bailón, S., & Wang, N. (2016). Networked discontent: The anatomy of protest campaigns in social media. *Social Networks*, *44*, 95–104.

Goodwin, J., Jasper, J. M., & Polletta, F. (2001). *Passionate politics: Emotions and social movements*. University of Chicago Press.

Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, *78*(6), 1360–1380.

Halupka, M. (2014). Clicktivism: A systematic heuristic. *Policy & Internet*, *6*(2), 115–132.

Hampton, K. N., Sessions Goulet, L., Rainie, L., & Purcell, K. (2011). Social networking sites and our lives. *Pew Internet & American Life Project*.

Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). *Emotional contagion*. Cambridge University Press.

Highfield, T. (2016). *Social media and everyday politics*. Polity Press.

Jasper, J. M. (2011). Emotions and social movements: Twenty years of theory and research. *Annual Review of Sociology*, *37*, 285–303.

Kingdon, J. W. (2011). *Agendas, alternatives, and public policies* (Updated 2nd edition). Longman.

Kitts, J. A. (2000). Mobilizing in black boxes: Social networks and participation in social movement organizations. *Mobilization: An International Quarterly*, *5*(2), 241–257.

Klandermans, B. (1997). *The individual and the collective: Individual motivation and collective action*. Blackwell Publishing.

Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*(24), 8788–8790.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, *33*(1), 159–174.

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

McAdam, D. (1986). Recruitment to high-risk activism: The case of freedom summer. *American journal of sociology*, *92*(1), 64–90.

McAdam, D., & Paulsen, R. (1993). Specifying the relationship between social ties and activism. *American journal of sociology*, *99*(3), 640–667.

McAdam, D., & Sewell, W. H. (2001). It's about time: Temporality in the study of social movements and revolutions. In *Silence and voice in the study of contentious politics* (pp. 89–125). Cambridge University Press.

McCarthy, J. D., & Zald, M. N. (1977). Resource mobilization and social movements: A partial theory. *American Journal of Sociology*, *82*(6), 1212–1241.

Meraz, S., & Papacharissi, Z. (2013). Networked gatekeeping and networked framing on egypt. *The International Journal of Press/Politics*, *18*(2), 138–166.

Mohammad, S. M., & Kiritchenko, S. (2015). Using hashtags to capture fine emotion categories from tweets. *Computational Intelligence*, *31*(2), 301–326.

Morozov, E. (2011). *The net delusion: The dark side of internet freedom*. PublicAffairs.

OpenAI. (2024). Best practices for prompt engineering with the openai api [Retrieved from https://help.openai.com/en/articles/6654000-best-practices-for-prompt-engineering-with-the-openai-api].

Papacharissi, Z. (2015). *Affective publics: Sentiment, technology, and politics*. Oxford University Press.

Passy, F. (2003). Social networks matter. but how? *Social Movements and Networks: Relational Approaches to Collective Action*, *1*, 21–48.

Polletta, F. (2006). *It was like a fever: Storytelling in protest and politics*. University of Chicago Press.

Polletta, F., & Jasper, J. M. (2001). Collective identity and social movements. *Annual Review of Sociology*, *27*(1), 283–305.

Rainie, L., & Wellman, B. (2012). *Networked: The new social operating system*. MIT Press.

Shirky, C. (2008). *Here comes everybody: The power of organizing without organizations*. Penguin Press.

Snow, D. A., & Benford, R. D. (1988). Ideology, frame resonance, and participant mobilization. *International social movement research*, *1*(1), 197–217.

Snow, D. A., Rochford, E. B., Worden, S. K., & Benford, R. D. (1986). Frame alignment processes, micromobilization, and movement participation. *American Sociological Review*, *51*(4), 464–481.

Stieglitz, S., & Dang-Xuan, L. (2013). Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. *Journal of Management Information Systems*, *29*(4), 217–248.

Stieglitz, S., Mirbabaie, M., Ross, B., & Neuberger, C. (2018). Crisis communication on twitter during a global crisis of volkswagen–the case of "dieselgate". *Business Information Systems Engineering*, *60*(4), 243–259.

Taylor, K.-Y. (2016). *From blacklivesmatter to black liberation*. Haymarket Books.

Theocharis, Y., Lowe, W., van Deth, J. W., & Garcia-Albacete, G. M. (2015). Uses of social media in public activism: Lessons from the transition process in egypt and beyond. *Information, Communication & Society*, *18*(9), 1011–1026.

Tufekci, Z. (2017). *Twitter and tear gas: The power and fragility of networked protest*. Yale University Press.

Valenzuela, S. (2013). Unpacking the use of social media for protest behavior: The roles of information, opinion expression, and activism. *American Behavioral Scientist*, *57*(7), 920–942.

Valenzuela, S., Park, N., & Kee, K. F. (2009). Social network sites, political engagement, and young adults: Political engagement. *Journal of Computer-Mediated Communication*, *14*(4), 875–901.

van Zomeren, M., Postmes, T., & Spears, R. (2012). Collective action and the role of self-definition and self-efficacy. *Journal of Social Issues*, *68*(1), 93–110.

Vromen, A., Xenos, M., & Loader, B. D. (2015). Everyday making through facebook engagement: Young citizens' political interactions in australia. *Political Communication*, *32*(1), 72–79.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*.

Wolfe, M. (2012). Putting on agenda setting: How media coverage affects public attention to policy issues. *Policy Studies Journal*, *40*(2), 220–246.

Appendix

Appendix