Automated Identification of Competing Narratives in Political Discourse on Social Media

Sergej Wildemann[1,*], Erick Elejalde[1]

[1]L3S Research Center, Leibniz Universität Hannover, Hannover, Germany

Abstract

Social media platforms have become central to shaping political discourse, serving as arenas where narratives form and evolve, influencing public opinion. Identifying and analyzing these narratives, particularly when they compete across different political ideologies, is crucial for understanding the dynamics of modern political communication. This paper presents an unsupervised framework for identifying and characterizing competing narratives in political discourse on social media, focusing on German politicians' tweets. The framework employs a multi-stage pipeline that integrates natural language processing techniques such as topic modeling, event detection, and event linking. By forming data into coherent stories and uncovering the distinct perspectives of user communities, the system is able to detect the key competing narratives, highlighting the divergent framings and conflicts surrounding trending political topics. Two case studies on polarizing political issues demonstrate the efficacy of the methodology, showcasing its ability to uncover and analyze divergent viewpoints. The findings contribute to the broader understanding of how narratives propagate within the digital public sphere and offer insights for policymakers, social media platforms, and researchers interested in monitoring political discourse.

Keywords

Narrative Frames, Competing Narrations, Social Media, Story Generation, Political Ideologies

1. Introduction

Social media has become a central arena for political discourse, shaping public opinion. Platforms like X (formerly Twitter) are key outlets for politicians to communicate directly with citizens [1]. Social media allows politicians to craft and directly disseminate narratives that resonate with their followers and political base [2]. These narratives—structured interpretations of events and issues—often compete with those from other groups, reflecting differing political ideologies and perspectives on pressing societal matters [3]. This can be seen, for example, in contrasting framings of events related to health crises [3] or climate change [4], where some groups emphasize a scientific perspective while others focus on belief-based interpretations [5]. Understanding how these competing narratives emerge, evolve, and propagate is essential for capturing the dynamics of contemporary political communication and addressing societal challenges such as polarization, misinformation, and the shaping of public opinion.

Our work contributes to the development of computational narrative framing analysis [6]. Specifically, this paper addresses the challenge of automatically identifying and analyzing competing narratives within the political discourse of German politicians on Twitter. For this, we propose a novel framework grounded in natural language processing (NLP) and clustering techniques. By focusing on political narratives in the social media context, we tap into a naturally polarized environment with frames that develop across multiple documents (e.g., tweets) and in a (potentially) collaborative way involving multiple sources. This analysis can provide insights into how political actors shape online debate and how competing viewpoints may influence societal understanding of political events.

This research area holds promise for fostering evidence-based approaches to societal challenges. For example, promoting media literacy and critical thinking through exposure to diverse viewpoints can empower individuals to evaluate information more critically, identify biases, and make informed

decisions [7]. Moreover, understanding competing narratives can facilitate constructive dialogue by highlighting shared values and differences, enabling collaborative problem-solving. The automatic extraction and analysis of competing narratives can advance efforts to combat misinformation, improve public understanding, and support more critical engagement with media content.

2. Background and Related Work

Narrative studies have a long history across various disciplines, including literary studies, psychology, and sociology. Narratives are typically defined as representations of a connected succession of events [8] and often blend reality and fiction to influence choices [9, 10]. They serve as systems of interrelated stories that may include competing storylines [9]. Various models of narrative analysis have been proposed, each focusing on different aspects [8]. For example, models can focus on temporal ordering, textual coherence, or the social functions of narratives. Models concerned with textual coherence examine how narratives are structured and organized to create a cohesive and understandable story. Conversely, models based on references and temporal order analyze actual events and their sequencing rather than just their textual representation [8]. Finally, models focusing on social functions investigate how narratives are used in social contexts to persuade, entertain, or inform a target audience [8]. This theoretical work lays the foundations for understanding narratives and developing practical computational methods for their analysis. Our approach combines various models to automatically identify narratives that are temporally and semantically coherent while considering their social functions, especially among political actors on social media.

A crucial aspect of computational narrative understanding is the automatic extraction of narrative elements from text data. This includes identifying actors, actions, events, settings, and the relationships between them [11, 12, 13]. For example, researchers have developed methods for identifying events, linking them based on temporal and causal relationships, and representing event-based narrative structures in different domains [11]. One promising representation approach employs a route map metaphor to visualize the events and storylines within a narrative, thereby facilitating understanding of the narrative's central themes [14]. These algorithms also assist in extracting events and storylines based on coherence and coverage [14, 15]. Furthermore, techniques for semantic role labeling and semantic graphs can effectively improve the analysis of narrative frames by extracting detailed information about the roles of different actors and actions within a story [5, 13, 16]. These building blocks are essential because reliable computational techniques enable analysts to scale their investigation to large volumes of narrative data, allowing for a deeper understanding of narrative structures and their impact.

Framing understanding is another crucial aspect in shaping the reader's perception of a narrative. Computational framing analysis aims to automatically detect framing devices and understand how they are used to influence interpretation [6]. Previous works have recognized specific linguistic patterns and rhetorical techniques that signal framings, such as word choice, metaphors, and the presentation of evidence [17, 16]. Frames can also be studied through sentiment analysis, where determining the emotional tone expressed toward different actors and events reveals evoked emotions and attitudes in the reader [18, 19]. Analyzing how specific frames are associated with various actors and concepts within a narrative helps uncover potential biases and perspectives [20, 21]. Computational framing analysis has shown to be effective in distinguishing between framing across different sources, such as conspiracy versus mainstream media, and in unveiling media bias [20, 6]. Nonetheless, challenges remain in consistently capturing framing information across narrative chains [22]. We adopt a computational narrative framing approach that examines the framing based on the structure of the narrative (e.g., the inclusion/exclusion of events by specific sources). However, more content-driven methods of framing analysis (e.g., sentiment analysis) can complement our methodology to provide additional insights.

Finally, studying competing narratives is essential for understanding how different perspectives and interpretations of events emerge and spread. This could help to monitor community interests, to counter propaganda campaigns, and to model the spread of (mis-)information [23, 6]. The task of computational approaches in this area focuses on identifying and tracking competing narratives' evolution over time

and understanding the factors that drive their dynamics [24, 25]. For this, proposed methods should distinguish between different narratives based on variations in actors, actions, framing devices, and the overall message conveyed. Competing narratives can then be tracked by analyzing changes in language use, sentiment, framing devices, and the prominence of different actors and actions over time. Investigating how external events, cultural shifts, and other narratives influence the emergence and spread of competing narratives allows us to understand their driving factors.

3. Dataset

For our study, we focus on the narratives within the political discourse of German politicians on Twitter/X. Before collecting the corresponding tweets, the relevant politicians and their Twitter accounts must first be identified. For this task, we leverage the structured database Wikidata, which provides information on politicians' affiliations with political parties and their country of citizenship.

We filter Wikidata entities by the *occupation* property matching "politician". Subsequently, the *country of citizenship* is utilized to filter German politicians. Their party affiliation is provided by the *member of political party* property. However, this property may list multiple parties, as politicians may change their affiliation over time. In certain instances, the current party is marked by a qualifier. Otherwise, we consider the first party listed, ignoring any affiliation with an end date. Extracting Twitter usernames from Wikidata with the above filters yielded 1,324 accounts belonging to 1,300 politicians.

Using Twitter's API, we collected all tweets of these politicians from January 1st, 2022, to June 24th, 2023. The resultant dataset comprised 189,850 tweets from 786 accounts, excluding retweets, replies, and quotes. For presentation purposes, we translated the tweets into English using the NLLB model [26]. Otherwise, all the analyses are conducted on the original content (i.e., original language).
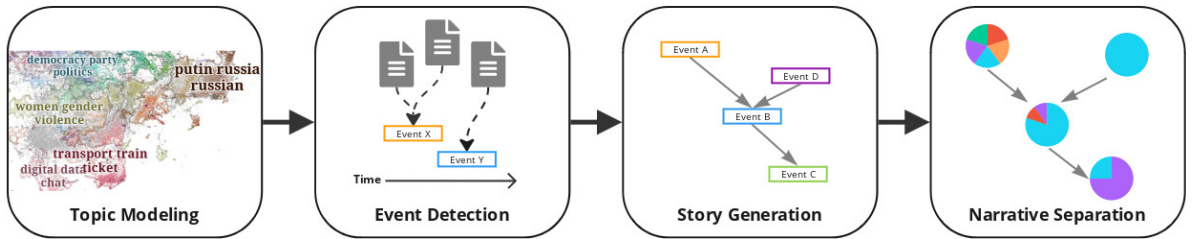
4. Methodology

Figure 1: Competing Narrative Extraction Pipeline: Separation of user viewpoints within stories of a topic.

We operationalize narratives as semantically and temporally coherent stories represented by structured interpretations of events and issues, providing meaning and context to political communication. Starting with the social media posts, we sought to identify these stories and the competing narratives of different user groups within them. Given a single tweet's brevity and limited scope, we follow a corpus-level approach to identify narratives rather than trying to extract them from individual tweets.

Our methodology consist of a multi-stage pipeline (see Fig. 1). After identifying general topics within the data, we then detect events that form the building blocks of our stories. Events are linked into coherent stories from which we extract different viewpoints, which we refer to as competing narratives.

Note that, during the narrative extraction process, we disregard tweet authors' political affiliations, as this information is unique to our dataset and may not be available in other contexts. However, it will be used during the evaluation to assess the effectiveness of our method to identify competing narratives.

We are publicly releasing the full source code for our data processing pipeline[1] to promote open science and ensure reproducibility.

---

[1]https://github.com/fjen/competing-narratives

Table 1

Excerpt of Topics Identified in the Dataset

Topic	Documents	Users	Keywords	Topic	Documents	Users	Keywords
0	11516	558	putin, russia, russian, war, ukraine	12	2836	413	climate, climate protection, protection, climate change, change
1	5898	519	going, good, morning, day, know	13	2819	431	european, eu, europe, european parliament, parliament
2	5237	505	transport, train, ticket, traffic, mobility	14	2506	309	water, food, forest, agriculture, nature
3	4477	485	ukraine, tanks, war, weapons, ukrainian	15	2303	435	greens, green, spd, fdp, party
4	3990	439	vaccination, vaccine, corona, pandemic, coronavirus	16	2142	427	qatar, football, congratulations, cup, world cup
5	3859	446	tax, inflation, debt, budget, discharge	17	1974	279	iran, regime, iranian, prison, prisoners
6	3789	460	children, education, child, wage, minimum	18	1932	307	nuclear, nuclear power, power, power plants, plants
7	3697	448	women, gender, violence, men, queer	19	1788	388	democracy, party, politics, left, reform
8	3360	462	berlin, great, concert, festival, today	20	1734	381	chancellor, scholz, merkel, olaf, olaf scholz
9	3165	420	anti, israel, semitism, anti semitism, victims	21	1606	258	39 39, sk, 39, people, war
10	3084	345	digital, data, chat, chat control, control	22	1554	258	china, chinese, taiwan, xi, beijing
11	3026	416	energy, gas, renewable, electricity, price	23	1472	297	asylum, refugees, migration, immigration, refugee

4.1. Topic Modeling

Given the diversity of topics in the dataset, the first step is to identify individual topics and remove noise. An unsupervised topic model is trained to group documents based on their semantic similarity. Here, we choose an embedding-based approach, as it is more robust to the short and noisy nature of tweets than traditional topic models like LDA. In particular, we rely on the modular BERTopic framework [27], which essentially clusters documents based on their semantic embeddings and generates text representations for each topic. Notably, we use Sentence-BERT with the `paraphrase-multilingual-MiniLM-L12-v2` model [28] for embedding generation and set a minimum cluster size of 500 documents to avoid overly fine-grained topics.

After disregarding outliers (i.e., not clustered tweets), we identify 57 topics in the dataset containing 108,063 documents. A significant portion of the documents are labeled as outliers (43%), which is expected given the nature of the data. As shown in Table 1, the remaining documents are grouped into topics that successfully capture the main discussion strands during the observation period with examples like the war in Ukraine, the COVID-19 pandemic, or Europe's energy crisis.

4.2. Event Detection

We define an event as a set of documents that discuss the same issue in close temporal proximity [15]. This definition implies that documents within an event share similar identifying information, such as keywords, actors, and/or locations. However, semantic similarity alone might lead to similar but distinct discussions being grouped together. Thus, the temporal aspect is crucial for clearly separating events.

For the extraction of events, we borrow parts of the proposed Multi-TimeLine Summarization (MTLS) approach by Yu et al. [15]. Although we use the same definition of events, their work is centered on (larger) news articles, whereas we focus on (shorter) tweets. Consequently, rather than relying on individual sentences, we utilize the entire tweet as the fundamental unit of analysis.

Tweets are clustered into events using the affinity propagation (AP) algorithm [29]. It operates on an affinity matrix S_1, where $S_1(i,j)$ represents the similarity between tweets i and j based on a linear combination of their semantic similarity and temporal proximity:

$$S_1(i,j) = \alpha \cdot S_{temp}(i,j) + (1 - \alpha) \cdot S_{cos}(i,j) \tag{1}$$

where S_{cos} is the cosine similarity of the tweet embeddings and S_{temp} is the temporal similarity defined by an exponential function (Eq. 2) with a decay factor λ:

$$S_{temp}(i,j) = \frac{1}{\exp^{\lambda \cdot |t_i - t_j|}} \tag{2}$$

Our time unit is in days and we set λ to 0.05 as suggested by Yu et al.. Also, α is set to 0.4 to balance the similarity measures. For example, we obtain 155 events for the *"energy"* topic (Topic 11 in Table 1).

4.3. Story Generation

Next, we group events into coherent stories, creating sequences of events that are meaningfully related. Since we are working at the document level and not at the sentence level, we cannot use a procedure similar to MTLS for event linking, which is based on sentence co-occurrence. Our approach is based on an event similarity graph where we filter the most significant edges to partition the events into stories.

Event Similarity Measure Since each event contains multiple tweets, we must define the semantic similarity measure on the inter-event level. The AP algorithm used to cluster those tweets also returns documents at the cluster centers as representatives of the events [15]. However, these centers may not cover the full breadth and depth of the event discussion. The same applies if calculating centroids of the tweet embeddings. Thus, we calculate semantic event similarity using the Sliced Wasserstein Distance (SWD) of the tweet embeddings within event pairs. Unlike cosine similarity, which focuses on the angular difference between pairs of document embeddings, SWD measures the distance between entire embedding distributions. It does this by slicing the multidimensional embedding space along various directions, computing Wasserstein distances in these 1D projections, and then averaging the results. Additionally, given that we average 15.9 tweets per event, SWD is computationally more efficient than calculating pairwise document distances. Finally, the SWD of the two distributions is transformed to the semantic similarity via the exponential decay function: $S_{sem} = \exp^{-\text{SWD}(e_1, e_2)}$.

Similarly to event detection, an affinity matrix S_2 is created based on the similarity of event embeddings and a penalty term, temporal distance, to avoid connecting temporally distant events, which in turn could lead to improbable stories spanning the entire time range. We define the event similarity $S_2(e_1, e_2)$ between a pair of events e_1 and e_2 as:

$$S_2(e_1, e_2) = S_{sem}(e_1, e_2) \cdot \exp^{-\gamma \cdot |t_{e_1} - t_{e_2}|} \tag{3}$$

where S_{sem} is the similarity of the event embeddings, and t_{e_i} is the timestamp of the representative document within an event. The temporal penalty decreases the similarity score as the time difference between the events increases, with the decrease rate controlled by the γ parameter. We set $\gamma = 0.01$.
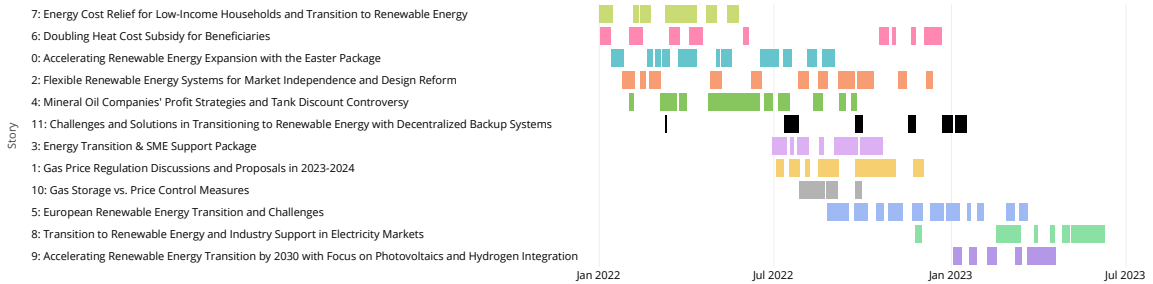

Figure 2: Extracted Stories with Events of Topic 11 (*energy, gas, renewable, electricity, price*). Each row represents a story. Bars in each story correspond to events, whose width is defined by the events' start and end dates.

Forming Storylines To create coherent storylines, we leverage the events similarity matrix S_2. We construct a directed graph $G = (V, E)$, where V represents the set of events and E the edges between them. Each edge e_{ij} between events i and j is assigned a weight based on their similarity score $S_2(i, j)$. The edge direction is determined by the chronological order of the events given by the representative documents, ensuring that each event has a pathway to all subsequent events. To maximize coherence, we retain only the highest-weighted outgoing edge for each event. This helps us focus on the main storyline avoiding diverging substories. Finally, to separate the storylines, we apply the Leiden algorithm [30] for community detection, with each partition representing a different story.

Figure 2 illustrates the outcome of the story extraction process for topic 11 (*energy, gas, renewable, electricity, price*), which resulted in the creation of 12 stories from 155 events.

4.4. Separation of Narratives

To distinguish the different viewpoints expressed in the extracted stories, we aim to identify users who share similar stances and are likely to contribute to the same narratives. Since politicians often hold aligned opinions on various topics, such as party policies, we adopt a global (i.e., not restricted to one particular topic) approach to user grouping. This method further ensures that even users less active within a given topic are appropriately categorized based on their overall trace of tweets.

Using a user's posts, we aim to create an embedding that captures their overall position. Text embeddings have been shown to successfully quantify the degree of political bias in texts [31] and identify ideological placement [32, 33]. Moreover, they can highlight differences in how political groups discuss specific policy issues [34]. To determine each user's overall stance, we first average the document embeddings of their posts within each single topic. This gives us a topic-specific representation of the user's perspective. Next, we calculate the user's relative position within that topic by subtracting the topic centroid from the user's average embedding. This step helps to situate the user's stance relative to the broader discourse on that issue. Finally, we average these topic-specific user embeddings across all topics. This global embedding represents the user's overall political stance across various issues.

To unravel the diverse perspectives shaping the narratives within a story, we use the commonalities of users involved in terms of their global embeddings. For this, we cluster the user embeddings (via HDBSCAN) to form distinct communities of users who share comparable political views. With multiple users contributing to a story, we can now split the story posts according to their respective communities. As a result, each event within a story is enriched by the perspectives of the various user communities, reflecting the collective voices that contribute to the discourse. *The structure represented by a given community's contribution to one particular story constitutes a competing narrative within that story.*

5. Results

We present the results based on applying the proposed framework to the dataset of German politicians' tweets. We first evaluate the division of user positions to assess the overall coherence and consistency of the narratives within each user community. Then, we illustrate, through two examples, how the same story or event can be differently perceived and discussed by opposing groups. Our analysis shows how the proposed framework can help reveal distinct narratives emerging from the political discourse. We validate our approach by showing how supporters of different (politically opposed) parties emphasizing divergent interpretations of the same events are segregated by our method into separate narratives.

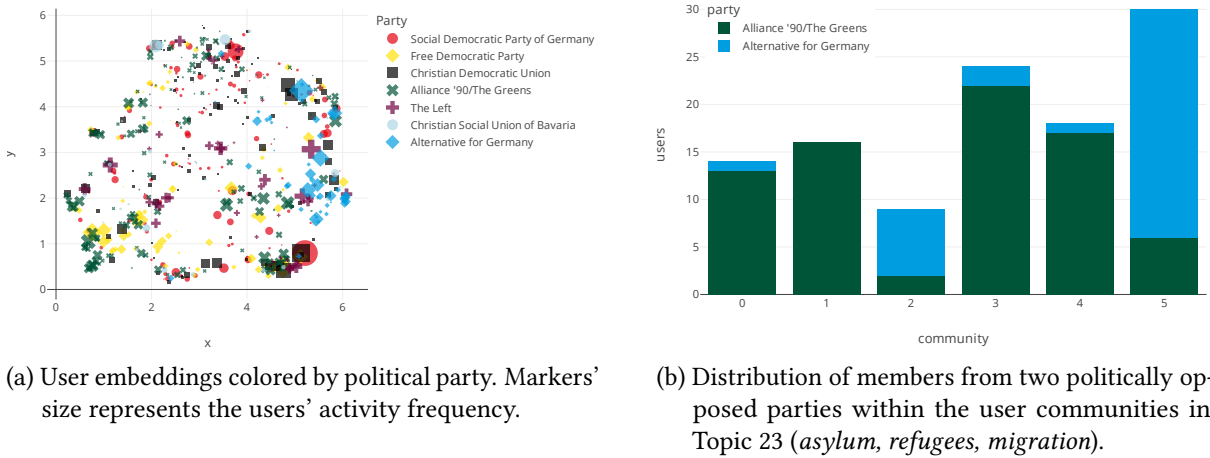### 5.1. Evaluation of Uncovered User Communities

(a) User embeddings colored by political party. Markers' size represents the users' activity frequency.

(b) Distribution of members from two politically opposed parties within the user communities in Topic 23 (*asylum, refugees, migration*).

Figure 3: Forming Political Communities with Similar Stances through Global User Embeddings.

(a) Story 10 within the directed graph of events

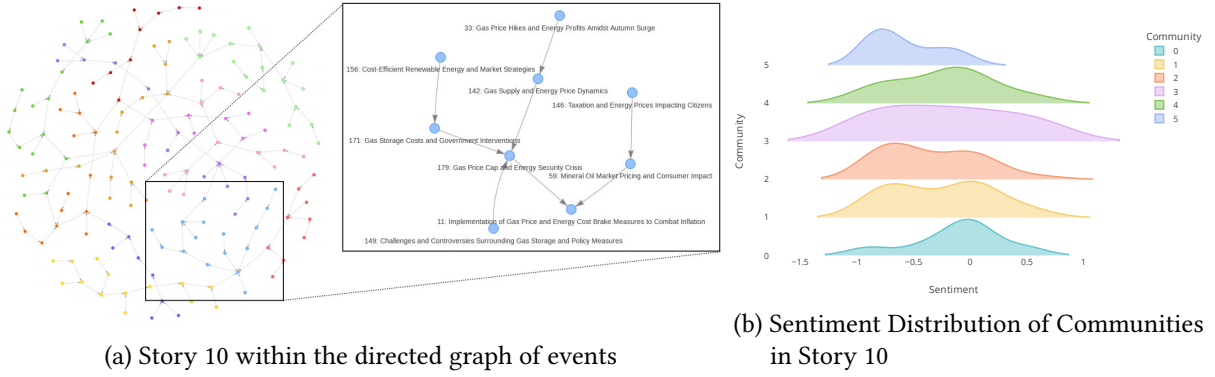(b) Sentiment Distribution of Communities in Story 10

Figure 4: Story 10 (*Gas Storage vs. Price Control Measures*) in Topic 11 (*energy, gas, renewable, electricity*)

Since the identification of distinct narratives within each story relies on user communities, we first verify the validity of the user classification. In our dataset, each user represents a German politician. Using Wikidata, we extracted information about their current party memberships. This allows us to assess the alignment of the identified communities with their respective political parties.

Figure 3a shows a 2D projection of the user embeddings generated in Section 4.4. We focus our analysis on the main parties currently represented in the German parliament with 758 users in our dataset. The figure indicates a good characterization of politicians based on their party affiliations. For example, The Greens party members form a cluster mostly to the left side of the embedding space, reflecting their shared ideological positions. In stark contrast, the right-wing Alternative for Germany (AfD) party occupies a cluster clearly separated and opposed to the Greens, highlighting the deep ideological divides between these two political factions. Other major parties, such as the CDU/CSU and SPD, are positioned more centrally, with their members distributed across the embedding space in a manner that also corresponds to their relative ideological stances. Further, we can observe an alignment of vocal members of the liberal party FDP with some of the Greens positions (bottom left).

The generated user communities (see Section 4.4) should express different viewpoints within each topic. We do not expect a clean separation of the individual parties given the heterogeneous views present within the larger, more central parties. Instead, for validation purposes, we focus our analysis on parties with relatively polarized and objectively opposing positions, namely The Greens and the AfD. These two parties appear well-separated within the identified communities, enabling the downstream use for narrative separation. Figure 3b exemplifies the six identified communities and their party affiliations composition for topic 23 (*asylum, refugees, migration*) when considering the two selected parties. Note, however, as implied by the embeddings overview, that these groups additionally include other members from multiple other parties, covering the sizeable political landscape.

5.2. Narratives Around Germany's 2022 Energy Crisis

As an insightful example for competing narratives, we focus on the discussions surrounding the energy crisis in Germany due to gas supply issues with Russia in 2022, which is captured by Topic 11 (*energy, ...*). From the story graph in Figure 4a, we select the debate over "*Gas Storage vs. Price Control Measures*". This story links the discussion of a planned gas price surcharge due to higher purchase costs with the later adopted "gas price cap" by the ruling "traffic light" (Ampel) coalition. Further references are made to increased storage costs, related taxation, and the excess profits of energy companies.

To delineate the differences in the viewpoints of the extracted communities and to identify contrasts, we first use sentiment analysis. Figure 4b shows the sentiment distribution within them, with documents classified using XLM-T [35]. Comparing the dominant communities (i.e., most posts in the story), 1 and 5, the latter uses noticeably more negative language, suggesting different framing.

Further, large language models allow us to summarize the community voices in order to demonstrate tangible differences in storytelling. The dominant positions taken by both communities are summarized

> "Community 1 predominantly expresses urgent concern over rising energy costs, particularly focusing on gas prices, which have spurred calls for immediate relief measures such as the removal of the Gas Umlage (surcharge) and implementation of a Gas Price Cap or Break. There is consensus on the necessity for robust governmental intervention to mitigate economic strain caused by these escalating expenses. Many advocate for suspending fiscal constraints like the debt brake to facilitate necessary spending during this crisis, emphasizing that such financial flexibility is essential to provide timely support to households and businesses. Additionally, there's a strong demand for long-term solutions including transitioning away from fossil fuels towards sustainable energy alternatives, with some suggesting more radical actions like the nationalization of energy companies to secure public interest over corporate profits. Overall, while immediate relief measures are widely supported, there is also recognition of the need for strategic planning and investment in renewable technologies to address future challenges."

> "Community 5 predominantly expresses strong opposition to the proposed Gas Surcharge, viewing it as economically burdensome, poorly conceived, and socially unfair. A dominant sentiment is that this measure was hastily implemented and crafted under influence from profit-driven entities rather than considering public welfare, highlighting a significant disconnect between government actions and citizens' interests. There's widespread support for the decision to halt the Gas Surcharge and replace it with measures like a Gas Price Brake, which are perceived as more equitable solutions aimed at alleviating the financial strain on households while ensuring economic stability without violating fiscal constraints such as the debt brake. The discourse also underscores frustration with governmental inefficiency and lack of transparency in decision-making processes, advocating for clearer communication and effective crisis management strategies that prioritize public interest over political or corporate gains. Overall, there is a call for more prudent and socially responsible policy-making that adequately addresses energy costs while safeguarding both individuals' financial well-being and broader economic health."

Figure 5: Generated summaries of community voices 1 and 5 within Story 10 (*Gas Storage vs. Price Control Measures*) in Topic 11 (*energy, gas, renewable, electricity*)

in Figure 5. Here, we prompt the phi-4 model [36] for all posts per community, gaining insights without sifting through the content.

In the ongoing debate over rising energy costs, the two communities find common ground in their advocacy for financial relief for households and their opposition to the proposed surcharge, but their narratives diverge as they each shed light on different matters. Community 1 supports government intervention to relieve the economy and also mentions long-term strategies such as the transition to renewable energy or, more radically, the nationalization of energy companies and the lifting of the debt brake. In contrast, Community 5 is mainly skeptical about the current government's approach and questions its intentions, but fails to mention any far-reaching solutions. The negative sentiment observed is also reflected in the individual posts, which are characterized by strong language. For example, the crisis is attributed a purely political origin and the main players in the government are defamed, while the planned gas surcharge is described as "truly miserable", unconstitutional and insane.

5.3. Law on Better Residence Opportunities for Migrants

Table 2
Diverse community positions within an Event in Story "Modern immigration law and civil equality in Germany" of Topic 23 (asylum, refugees, migration). [translated from German]

Community	0	1	2	3	4	5
Central Post	As #Ampel, we stand for a paradigm shift in #migration policy. We are opening up more legal ways to come to Germany and work here. Because we have a shortage of workers everywhere.	It is true. We need more regulated immigration and less illegal immigration. That includes: anyone who does not have a right of residence in Germany for reasons of political persecution or humanitarian emergency cannot stay in the country.	German SMEs are so desperately in need of new employees that your concerns fade into the background. Of course, we must consistently prevent migration into the social systems. But our country must be a country of immigration.	The Bundestag has just debated our new #ResidenceOpportunitiesAct for the first time. Tolerated people who have been afraid of deportation for years now have prospects in Germany. Good for people, good for companies!	The Residence Opportunities Act gives people who have lived in Germany for many years more planning security for training or a job.	We need a deportation offensive. #Germany must no longer be a haven for mentally conspicuous "lone offenders". The #security of citizens must have priority. We owe this to the many victims of migration policy since 2015.

144

In addition to unpacking the community positions in the overall context of a story, we want to address the different facets of an event. For this example, we switch to the polarizing debate on migration in Topic 23, where the planned facilitation of the citizenship of migrant workers is discussed.

Table 2 shows the extracted representative tweets per community within an event. The perspectives on work migration and swifter integration are divided into generally positive statements from Communities 0, 3, and 4, positive with limitations from 1 and 2, and general rejection from Community 5. The main similarities are the recognition of the need for more regulated immigration and the desire to address the labor shortage in Germany. However, their approaches differ, with some advocating for more open and inclusive policies, while others prioritize security and deportation. Different framings are employed to support the narratives. For example, Communities 0 and 2 frame the issue in terms of economic needs, while Communities 1 and 5 focus on security and preventing illegal immigration. Conversely, Communities 3 and 4 highlight the positive impact of the new Residence Opportunities Act.

Dissecting community involvement within each event in a story can also give analysts insights into the narrative framing at a more structural level. For example, we can track which events different narratives try to emphasize by looking at their community distribution and which events communities select to skip altogether. Moreover, analysts can study the timing and order in which competing narratives cover the events in the story to try to uncover internal dynamics and driving factors.

6. Conclusions and Future Work

This work introduces an unsupervised framework for identifying and analyzing competing political narratives on social media, focusing on German politicians' tweets. We uncovered competing narratives by automatically organizing topics and events into coherent stories and leveraging politician embeddings to identify distinct community perspectives. We validate our method through illustrative examples related to the Energy Crisis and Migration Policy, demonstrating the framework's effectiveness in separating competing narratives within highly polarized political discourse. This approach to large-scale social media analysis offers valuable insights, underscoring the potential of automatic narrative framing to monitor how political actors promote divergent interpretations of events on social media. These findings provide valuable insights for policymakers, social media platforms, and researchers seeking to better understand political discourse, address polarization, and foster more balanced public discussions.

It is essential to highlight that the development of computational approaches for analyzing competing narratives relies on suitable datasets and evaluation metrics. Existing datasets for narrative extraction often lack annotations related to competing narratives, making it challenging to train and evaluate models that can effectively distinguish and analyze them [37]. Furthermore, developing robust metrics to assess the quality and coherence of extracted narratives and the accuracy of framing analysis remains an open challenge [38, 6]. In this work, we rely primarily on qualitative evaluation. However, we see potential to further strengthen our approach by correlating the linking of events to stories with the user streams. This would allow us to measure and potentially enhance the coherence of the narratives extracted. Initial experiments in this direction were promising, helping identify more coherent storylines that better reflect the framing and propagation of narratives across different political communities.

Acknowledgments

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude 3 Haiku and Grammarly in order to: Grammar and spelling check, paraphrase, and reword. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] F. Gilardi, T. Gessler, M. Kubli, S. Müller, Social media and political agenda setting, Political communication 39 (2022) 39–60.

[2] B. C. Silva, S.-O. Proksch, Politicians unleashed? political communication on twitter and in parliament in western europe, Political science research and methods 10 (2022) 776–792.

[3] E. Jing, Y.-Y. Ahn, Characterizing partisan political narrative frameworks about covid-19 on twitter, EPJ data science 10 (2021) 53.

[4] A. S. Ross, D. J. Rivers, Internet memes, media frames, and the conflicting logics of climate change discourse, Environmental communication 13 (2019) 975–994.

[5] M. Reiter-Haas, B. Klösch, M. Hadler, E. Lex, Framing analysis of health-related narratives: Conspiracy versus mainstream media, arXiv preprint arXiv:2401.10030 (2024).

[6] M. Reiter-Haas, B. Klösch, M. Hadler, E. Lex, Computational narrative framing: Towards identifying frames through contrasting the evolution of narrations., in: Text2Story@ ECIR, 2024, pp. 129–135.

[7] R. Axelrod, J. J. Daymude, S. Forrest, Preventing extreme polarization of political attitudes, Proceedings of the National Academy of Sciences 118 (2021) e2102139118.

[8] E. G. Mishler, Models of narrative analysis: A typology, Journal of narrative and life history 5 (1995) 87–123.

[9] R. E. Page, Stories and social media: Identities and interaction, Routledge, 2013.

[10] M. D. Jones, M. K. McBeth, A narrative policy framework: Clear enough to be wrong?, Policy studies journal 38 (2010) 329–353.

[11] B. F. Keith Norambuena, T. Mitra, C. North, A survey on event-based news narrative extraction, ACM Computing Surveys 55 (2023) 1–39.

[12] I. Mani, Computational modeling of narrative, Springer Nature, 2022.

[13] D. Metilli, V. Bartalesi, C. Meghini, et al., Steps towards a system to extract formal narratives from text, in: Text2Story@ ECIR, 2019, pp. 53–61.

[14] B. F. Keith Norambuena, T. Mitra, Narrative maps: An algorithmic approach to represent and extract information narratives, Proceedings of the ACM on Human-Computer Interaction 4 (2021) 1–33.

[15] Y. Yu, A. Jatowt, A. Doucet, K. Sugiyama, M. Yoshikawa, Multi-timeline summarization (MTLS): improving timeline summarization by generating multiple summaries, in: C. Zong, F. Xia, W. Li, R. Navigli (Eds.), Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, Association for Computational Linguistics, 2021, pp. 377–387. URL: https://doi.org/10.18653/v1/2021.acl-long.32. doi:10.18653/V1/2021.ACL-LONG.32.

[16] Y. Otmakhova, S. Khanehzar, L. Frermann, Media framing: A typology and survey of computational approaches across disciplines, in: L.-W. Ku, A. Martins, V. Srikumar (Eds.), Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Bangkok, Thailand, 2024, pp. 15407–15428. URL: https://aclanthology.org/2024.acl-long.822/. doi:10.18653/v1/2024.acl-long.822.

[17] S. Liu, L. Guo, K. Mays, M. Betke, D. T. Wijaya, Detecting frames in news headlines and its application to analyzing news framing trends surrounding us gun violence, in: Proceedings of the 23rd conference on computational natural language learning (CoNLL), 2019, pp. 504–514.

[18] N. D. Gitari, Z. Zuping, H. Damien, J. Long, A lexicon-based approach for hate speech detection, International Journal of Multimedia and Ubiquitous Engineering 10 (2015) 215–230.

[19] N. Pitropakis, K. Kokot, D. Gkatzia, R. Ludwiniak, A. Mylonas, M. Kandias, Monitoring users' behavior: anti-immigration speech detection on twitter, Machine Learning and Knowledge Extraction 2 (2020) 11.

[20] F. Hamborg, A. Zhukova, B. Gipp, Automated identification of media bias by word choice and labeling in news articles, in: 2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL), IEEE, 2019, pp. 196–205.

[21] S. Wildemann, C. Niederée, E. Elejalde, Migration reframed? a multilingual analysis on the stance shift in europe during the ukrainian crisis, in: Proceedings of the ACM Web Conference 2023, 2023, pp. 2754–2764.

[22] J. Mendelsohn, C. Budak, D. Jurgens, Modeling framing in immigration discourse on social media, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2021, pp. 2219–2263.

[23] F. Hamborg, K. Donnay, B. Gipp, Automated identification of media bias in news articles: an interdisciplinary literature review, International Journal on Digital Libraries 20 (2019) 391–415.

[24] K. K. Bandeli, M. N. Hussain, N. Agarwal, A framework towards computational narrative analysis on blogs., in: Text2Story@ ECIR, 2020, pp. 63–69.

[25] K. A. Schmid, A. Z "ufle, D. Pfoser, A. Crooks, A. Croitoru, A. Stefanidis, Predicting the evolution of narratives in social media, in: Advances in Spatial and Temporal Databases: 15th International Symposium, SSTD 2017, Arlington, VA, USA, August 21–23, 2017, Proceedings 15, Springer, 2017, pp. 388–392.

[26] N. team, M. R. Costa-jussà, J. Cross, O. cCelebi, M. Elbayad, K. Heafield, K. Heffernan, E. Kalbassi, J. Lam, D. Licht, J. Maillard, A. Sun, S. Wang, G. Wenzek, A. Youngblood, B. Akula, L. Barrault, G. M. Gonzalez, P. Hansanti, J. Hoffman, S. Jarrett, K. R. Sadagopan, D. Rowe, S. L. Spruit, C. Tran, P. Y. Andrews, N. F. Ayan, S. Bhosale, S. Edunov, A. Fan, C. Gao, V. Goswami, F. Guzm'an, P. Koehn, A. Mourachko, C. Ropers, S. Saleem, H. Schwenk, J. Wang, No language left behind: Scaling human-centered machine translation, ArXiv abs/2207.04672 (2022).

[27] M. Grootendorst, Bertopic: Neural topic modeling with a class-based tf-idf procedure, arXiv preprint arXiv:2203.05794 (2022).

[28] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2019. URL: https://arxiv.org/abs/1908.10084.

[29] B. J. Frey, D. Dueck, Clustering by passing messages between data points, science 315 (2007) 972–976.

[30] V. A. Traag, L. Waltman, N. J. Van Eck, From louvain to leiden: guaranteeing well-connected communities, Scientific reports 9 (2019) 1–12.

[31] J. Gordon, M. Babaeianjelodar, J. N. Matthews, Studying political bias via word embeddings, in: A. E. F. Seghrouchni, G. Sukthankar, T. Liu, M. van Steen (Eds.), Companion of The 2020 Web Conference 2020, Taipei, Taiwan, April 20-24, 2020, ACM / IW3C2, 2020, pp. 760–764. URL: https://doi.org/10.1145/3366424.3383560. doi:10.1145/3366424.3383560.

[32] L. Rheault, C. Cochrane, Word embeddings for the analysis of ideological placement in parliamentary corpora, Political Analysis 28 (2020) 112–133. doi:10.1017/pan.2019.26.

[33] E. Elejalde, L. Ferres, E. Herder, On the nature of real and perceived bias in the mainstream media, PloS one 13 (2018) e0193765.

[34] M. J. Annika Fredén, D. Saynova, Word embeddings on ideology and issues from swedish parliamentarians' motions: a comparative approach, Journal of Elections, Public Opinion and Parties 0 (2024) 1–22. doi:10.1080/17457289.2024.2433979.

[35] F. Barbieri, L. Espinosa Anke, J. Camacho-Collados, XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond, in: Proceedings of the Thirteenth Language Resources and Evaluation Conference, European Language Resources Association, Marseille, France, 2022, pp. 258–266. URL: https://aclanthology.org/2022.lrec-1.27.

[36] M. Abdin, J. Aneja, H. Behl, S. Bubeck, R. Eldan, S. Gunasekar, M. Harrison, R. J. Hewett, M. Javaheripi, P. Kauffmann, J. R. Lee, Y. T. Lee, Y. Li, W. Liu, C. C. T. Mendes, A. Nguyen, E. Price, G. de Rosa, O. Saarikivi, A. Salim, S. Shah, X. Wang, R. Ward, Y. Wu, D. Yu, C. Zhang, Y. Zhang, Phi-4 technical report, 2024. URL: https://arxiv.org/abs/2412.08905. arXiv:2412.08905.

[37] B. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, S. Nunes, A survey on narrative extraction from textual data, Artificial Intelligence Review 56 (2023) 8393–8435.

[38] P. Ranade, S. Dey, A. Joshi, T. Finin, Computational understanding of narratives: A survey, IEEE Access 10 (2022) 101575–101594.