

Exploring Large-Scale Social Media Data Using Sentiment Analysis and Topic Modeling - A Case Study of Instagram Posts in Context of the 2024 German Government Crisis

Jessica Stieglmayer*
Universität Regensburg
Regensburg, Germany
Jessica.Stieglmayer@stud.uni-regensburg.de

Arabella Petz*
Universität Regensburg
Regensburg, Germany
Arabella.Petz@stud.uni-regensburg.de

Leo Bruckmüller
Universität Regensburg
Regensburg, Germany
Leo.Bruckmueller@stud.uni-regensburg.de

Jakob Fehle
Universität Regensburg
Regensburg, Germany
jakob.fehle@ur.de

Michael Achmann-Denkler
Universität Regensburg
Regensburg, Germany
michael.achmann@ur.de

Christian Wolff
Universität Regensburg
Regensburg, Germany
christian.wolff@ur.de

Abstract

This paper presents the findings of a study combining sentiment analysis and topic modeling on Instagram posts published by members of the German Parliament, with a particular focus on the government crisis in the fall of 2024. The analyzed dataset – collected and curated for this study – contains more than 480,000 posts, including an annotated subset of 2,000 posts for evaluation. For sentiment classification, we fine-tuned the German language model GBERT-base, and employed BERTopic to discover thematic structures within the corpus. We conducted a detailed temporal analysis around the period of the coalition’s collapse, closely integrating sentiment trends with topical developments. This case study demonstrates the potential of combining sentiment analysis and topic modeling for distant reading and the large-scale exploration of political communication on social media platforms, showing that these two methods combined can be used to identify crucial events such as incidents regarding the government crisis, the Magdeburger Christmas market attack or news about the case of Gisèle Pelicot.

CCS Concepts

• **Computing methodologies** → **Natural language processing**: *Machine learning*; • **Information systems** → *Social networks*.

Keywords

Political communication, Social media analysis, Sentiment analysis, Topic modeling, Instagram, German politics

ACM Reference Format:

Jessica Stieglmayer, Arabella Petz, Leo Bruckmüller, Jakob Fehle, Michael Achmann-Denkler, and Christian Wolff. 2026. Exploring Large-Scale Social Media Data Using Sentiment Analysis and Topic Modeling - A Case Study of Instagram Posts in Context of the 2024 German Government Crisis. In *18th ACM Web Science Conference (WebSci '26)*, May 26–29, 2026, Braunschweig, Germany.

*The first and second author contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WebSci '26, Braunschweig, Germany*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2504-3/26/05

<https://doi.org/10.1145/3795766.3799749>

Germany. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3795766.3799749>

1 Introduction

Social media platforms have become central arenas of political communication, enabling elected representatives to articulate positions, react to current events, and engage with constituents in real time [8]. While much of the computational research in this domain has focused on Twitter (now X) [10], Instagram has grown into a major venue for political communication [3, 40]. Existing studies on Instagram focus, for example, on visual content or engagement metrics [e.g. 23, 24, 35], while the platform’s textual dimension has received less scholarly attention.

While Instagram is built around visual content, captions represent the platform’s central textual component. They provide short, structured text accompanying images or videos and thus constitute a consistent basis for computational analysis. Methods such as sentiment analysis and topic modeling, frequently applied to Twitter data [17, 27], can therefore be transferred to Instagram captions to examine thematic and affective patterns in large-scale political communication. Yet sentiment analysis and topic modeling are typically applied in isolation, even though their combined use offers potential for enhancing distant reading approaches.

This paper addresses this gap by examining the textual communication of 708 members of the German Parliament during the 2021–2025 legislative period. Using a corpus of 488,917 Instagram posts collected and curated for this study, we investigate whether combining sentiment analysis and topic modeling enables a more nuanced understanding of political communication across parties, time periods, and critical events. This leads us to the following primary research question:

RQ1: How can sentiment analysis and topic modeling be combined to enhance the distant reading of social media data and improve our understanding of political communication?

To explore this overarching question, we also examine two more specific dimensions:

RQ2: Can shifts in sentiment be observed over the course of the coalition period, and can topic modeling help to identify potential causes for these changes?

RQ3: Are there detectable differences between the political parties in terms of the sentiment and thematic focus of their Instagram posts?

The 2024 German government crisis, marked by the collapse of the coalition between the SPD (social democrats), Bündnis 90/Die Grünen (green party), and the FDP (liberals), provides a particularly informative context for evaluating the combined approach. The crisis sparked intense public and political debate [41], offering a real-world test case for assessing whether changes in sentiment and topic prevalence facilitate a comprehensive understanding of the Instagram communication of 708 members of the German parliament.

Contributions

This paper makes three main contributions:

(1) We fine-tune and evaluate German transformer-based sentiment models using public datasets and evaluate the models using a manually annotated gold-standard dataset of 2,000 posts. (2) We analyze the distribution and development of topics and sentiments during the 2021–2025 legislative period by applying sentiment analysis and topic modeling to the full corpus. (3) We demonstrate how the integration of sentiment analysis and topic modeling can reveal meaningful shifts in political communication during the 2024 government crisis.

2 Related Work

This section provides an overview of how sentiment analysis and topic modeling have been applied to the examination of social networks, with a particular focus on political communication. It also discusses the key studies and findings that underpin our research.

2.1 Sentiment Analysis in Social Networks for Political Research

According to Rodríguez-Ibáñez et al. [25], there is a great interest in sentiment analysis, both in industry and academia, as it is recognized as a useful approach to assess public opinion in various domains. The authors looked at 2306 studies that dealt with sentiment analysis in social networks and named politics as the second most dominant field, with 459 publications. They further note that research in the political field predominantly investigates topics such as elections, political parties, hate speech, and specific political processes within individual countries.

Schmidt et al. [27] distinguish between studies that focus on the communication of political actors themselves and those that analyse public sentiment toward political events or figures (see e.g. [16]). Their own study is an example of the first, as they focused on Twitter (now X) posts of the seven parties, represented in the German Parliament, around the federal election 2021 in Germany. The authors collected over 58,000 tweets and compared several approaches for the sentiment analysis and their performance.

Regarding public opinion towards politicians, studies have examined whether it is possible to predict election outcomes using social media analysis [27]. Chauhan et al. [7] list many such works, for example for elections in India, the Netherlands, Greece, Germany, Malaysia, Pakistan, Taiwan and the USA. They conclude that most

studies have successfully predicted election outcomes, highlighting the growing potential of social media data and computational methods for making political forecasts.

In our study, we want to focus on sentiment shifts and differences between parties as Schmidt et al. [27] did, but also on temporal analysis to find out how certain events can influence sentiment over time [25], and whether certain topics are more associated with negative, positive or neutral sentiments.

Rodríguez-Ibáñez et al. [25] observe that, in the studies they reviewed on sentiment analysis in social networks, state-of-the-art techniques such as transformer-based systems are not used as often as traditional approaches, including dictionary-based methods, neural networks or Support Vector Machines. However, the authors note an increased use of embeddings and transformers-based methods in recent years. Similarly, Schmidt et al. [27] find that transformer-based models, such as BERT, outperform traditional machine learning approaches in sentiment classification tasks.

Thus, in this work, we aim to build on current developments in Natural Language Processing by comparing Transformer-based approaches and selecting the best classification model for our task.

2.2 Topic Modeling on Social Media Data in the Context of Politics

Topic modeling has become a popular computational method for the analysis of social media data and is applied in fields such as journalism, public health and political science [20]. Hellwig et al. [17] highlight that topic modeling has been effectively used to extract information from tweets in various political contexts. The authors themselves examined tweets connected to the 2021 German federal election and showed that the most frequently discussed topics were party-oriented policies, climate, and financial policy. Moreover, they found out that some topics were debated constantly throughout the year, while others came up at certain periods of time. Another similar example is the work of Achmann and Wolff [1], who looked at Instagram stories and posts in the context of the 2021 Federal election campaign in Germany, and explored differences between these two formats using BERTopic.

Earlier topic modeling algorithms such as Latent Dirichlet Allocation (LDA) are suitable for processing longer texts, but Laureate et al. [20] note that their ability to extract semantic meaning from shorter texts is limited. LDA is a generative probabilistic model [17] that treats documents as mixtures of latent topics with each topic represented as a probability distribution over words [36].

In recent years, models optimized for short text analysis have been introduced, e.g., a mixture of unigram models, the biterm topic model (BTM), and the dual sparse topic model (DSTM) [36]. To address a key limitation of LDA — namely its inability to account for semantic relationships between words — text embedding techniques such as BERT are increasingly used for topic modeling, because they consider words in the context of surrounding text, allowing for a deeper understanding of their meaning and capturing semantic relationships [17].

Since it is essential to choose the right topic model for getting the best possible results [36], we build on the work of Hellwig et al. [17]. The authors used BERTopic, introduced by Grootendorst [12], to analyze German tweets related to the 2021 Federal Election.

This topic modeling technique creates document embeddings using pre-trained transformer-based language models, clusters these embeddings, and then derives topic representations through a class-based TF-IDF procedure [12].

3 Data & Methods

In this section, we outline the methodological approach to answer our research questions, including a general overview of the data collection process, the corpus, the fine-tuning of appropriate models, as well as the evaluation procedures used to assess their performance.

3.1 Data Collection

To assess how members of the 20th German Parliament used Instagram, we first compiled a list of all 733 elected representatives from Wikipedia and manually enriched it with URLs to all available Instagram accounts. This list was then programmatically preprocessed for a semi-automated scraping workflow. Given that direct access to social media data has become increasingly challenging since the so-called “APIcalypse” [4, 11], we used *Zeeschuimer* [22], a Firefox extension that intercepts the browser-server communication and records all posts rendered in the browser.

Data collection took place in late February 2025. Using a scripted workflow built with Selenium, we automated the process of loading each public Instagram profile, scrolling through its timeline, and exporting all posts captured by *Zeeschuimer*. Private accounts were skipped, and occasional scrape failures (e.g., temporary loading limits or changed profile URLs) were flagged for manual review. The resulting exports were converted to CSV format and merged into a unified corpus. After removing posts published before 2021, the final dataset contains valid post data for 712 accounts (708 individual politicians). The dataset is available on request.

3.2 Corpus Overview

Our corpus comprises 488,917 captions of Instagram posts. At the time of data collection, the 733 members of Parliament were distributed across parties as follows: 207 SPD, 196 CDU/CSU, 117 Greens, 90 FDP, 76 AfD, 28 Die Linke, 10 BSW and 9 independents [38]. The distribution within the corpus is presented in Table 1.

3.3 Data Annotation

To evaluate the performance of our sentiment analysis model, we created a gold-standard dataset of 2,000 randomly sampled Instagram posts from our corpus. Four independent annotators labelled the dataset in two batches, with two annotators assigned to each batch. Annotators assigned one of three sentiment categories to each caption: positive, neutral, negative [27]. These categories, accompanied by representative corpus examples, were defined as follows:

- positive:** Posts with a predominantly positive sentiment, recognizable by words or phrases such as “looking forward”, “thankful”, or “fun”, or if the context of a post is positive, e.g., successes, productive visits, or achievements.
- negative:** Posts with a predominantly negative sentiment, recognizable by words or phrases such as “unbearable”, “appalling”, or “danger”, or if the context of a post is negative, e.g., natural disasters, bad results, or deaths.

neutral: Posts with a neutral sentiment, no sentiment, or a mix of positive and negative sentiment.

Disagreements within each batch were reviewed by the two annotators from the other batch and resolved through majority vote. Inter-coder reliability was assessed using Krippendorff’s α using the *krippendorff* python package [6]. Agreement was high across both batches (batch 1: $\alpha = 0.82$; batch 2: $\alpha = 0.91$). These values exceed typical levels reported in social media annotation studies, where average alpha values around 0.61 are common [26], and surpass those obtained in comparable political sentiment tasks [27]. Nonetheless, annotators noted that determining sentiment was frequently challenging.

The final gold-standard dataset contains 742 positive posts, 285 negative posts, and 973 neutral posts. As outlined below, this dataset was used to evaluate and select the best-performing sentiment model. This subset is available in our GitHub repository.¹

3.4 Sentiment Analysis

Sentiment analysis is implemented as a supervised text classification task in which each Instagram post is assigned one of the above labels [19]. Since it is crucial to find the right model for the task [19] we evaluated several approaches (see Table 2a): To ensure our approach remains broadly applicable, we assume a scenario without a self-annotated dataset for a first model overview (Table 2a) and fine-tuning (see Table 2b) using a public dataset. Thus, we are using our annotated gold-standard dataset only for the final evaluation and model selection.

In a first step, we compared three pre-trained transformer-based sentiment models that support German: a multilingual DistilBERT model optimized for social media text, German Sentiment BERT, and XLM-RoBERTa German Sentiment (Table 2a). Following Schmidt et al. [27], we evaluated these models on their German Politicians Twitter Sentiment test set [2], which consists of 357 tweets labeled as positive, negative, or neutral. All models achieved only moderate performance, with weighted F_1 scores between 0.38 and 0.52, which we considered insufficient for our downstream analysis of political communication.

We therefore fine-tuned a German transformer model specifically for our task. In line with Schmidt et al. [27], we fine-tuned the GBERT-base language model [9] for three-way sentiment classification using publicly available training data and Hugging Face’s Transformers library [18]. Table 2b summarizes the configurations and results of the two best-performing variants. GBERT1 was trained on 1,428 posts from the German Politicians Twitter Sentiment dataset [2], GBERT2 extended this training data with 19,400 samples from the GermEval 2017 shared task [39]. Both models were trained for four epochs with a learning rate of 2^{-5} , differing only in batch size and training data composition.

For model selection, we evaluated both fine-tuned models on our gold standard dataset of 2,000 Instagram posts. GBERT1 achieved the highest performance, with accuracy, precision, recall, and F1 all at 0.72, outperforming GBERT2 (0.67 on all metrics). We therefore used GBERT1 to assign sentiment labels to all 488,917 posts in our corpus.

¹<https://github.com/jessimaye/BundesPosts>

Table 1: German political parties and their percentage of posts within the corpus.

Party Name	Abbreviation	Political Position	Percentage in Corpus
Sozialdemokratische Partei Deutschland	<i>SPD</i>	Centre-left	30.68%
Christlich Demokratische Union/ Christlich Soziale Union	<i>CDU/CSU</i>	Centre-right	30.42%
Bündnis 90/Die Grünen	<i>Greens</i>	Centre-left	14.35%
Freie Demokratische Partei	<i>FDP</i>	Centre-right	10.31%
Alternative für Deutschland	<i>AfD</i>	Far-right	8.79%
Die Linke	<i>Left</i>	Left-wing	3.60%
Bündnis Sarah Wagenknecht	<i>BSW</i>	Left-wing to Far-left	0.90%
Fraktionslose	<i>Independentants</i>	/	0.65%

Table 2: Comparison of Sentiment Models during our Evaluation. All evaluation metrics report weighted average values.

Model	Source	Description	Acc	Prec	Rec	F1
Multilingual Sentiment	[30]	distilBERT-based, multilingual, social media optimized	0.53	0.53	0.53	0.52
German Sentiment BERT	[15]	BERT-based, trained on 1.8M German posts	0.50	0.59	0.50	0.44
XLM-R German Sentiment	[28]	XLM-RoBERTa, trained on 200k German samples	0.47	0.60	0.47	0.38

(a) Pre-trained sentiment models, evaluated on the German Politicians Twitter Sentiment [2] test set (357 examples) by Schmidt et al. [27].

Model	Fine-Tuning Data	Epochs	LR	Batch-Size	Acc	Prec	Rec	F1
GBERT1	1,428 posts from German Politicians Twitter Sentiment	4	2^{-5}	8	0.72	0.72	0.72	0.72
GBERT2	Train data of GBERT1 + 19,400 samples from Germeval 2017	4	2^{-5}	16	0.67	0.67	0.67	0.67

(b) Fine-Tuned GBERT models and evaluated on our annotated subset (2,000 samples).

3.5 Topic Modeling

To analyze the thematic structure of political communication on Instagram and to complement the sentiment analysis with contextual information (RQ1, RQ2), we employed the BERTopic framework introduced by Grootendorst [12]. As outlined in our related work section, embedding-based topic models outperform traditional probabilistic approaches on short, heterogeneous social media texts and have been successfully applied to political communication [1, 17, 20]. Building on this line of research, we adopt BERTopic for our large-scale Instagram corpus.

Specifically, we implemented a configuration optimized for large-scale data processing, as outlined by Grootendorst [14]. Following standard preprocessing procedures, including the removal of stop words, the documents (i.e., the social media posts) are transformed into dense vector representations using Sentence-BERT embeddings [14]. In line with the recommendations of Grootendorst [13], we applied the pretrained embedding model paraphrase-multilingual-MiniLM-L12-v2,² which is well-suited for multilingual content.

The resulting embeddings were cached temporarily and subsequently passed to a multilingual BERTopic model for unsupervised clustering and topic extraction. The initial number of generated

topics was reduced to 150 using the model’s built-in topic reduction functionality. Given that a substantial portion of the data was initially classified as outliers — i.e., documents not assigned to any dominant topic — we applied the `reduce_outliers` method [13] to reassign these documents to the most semantically appropriate topics.

To evaluate the coherence and contextual relevance of the topic model, we identified the ten most frequent topics separately for each sentiment (negative, neutral, positive), and compared them with corresponding word clouds. This qualitative analysis aimed to identify thematic variations across sentiment categories rather than provide a quantitative assessment. The positive word cloud was dominated by clearly affirmative terms such as “Danke” (“thanks”) and “freue” (“look forward to”), while the negative one displayed a broader distribution of prominent words with less distinct weighting.

This is also reflected by the topics: several topic-specific keywords appear within the positive word cloud that align with the top ten topics identified by the model, including “Bundestag” (“Federal Parliament”) and “Demokratie” (“democracy”), whereas the negative one is more mixed, with some words not aligned with any of the most frequent themes.

²<https://huggingface.co/sentence-transformers/paraphrase-multilingual-MiniLM-L12-v2>

3.6 Joint Sentiment and Topic Analysis

For the temporal analysis we first looked at unusual spikes within each of the three sentiments for the months leading up to, during and after the governmental crisis. As part of the distant-reading approach, graphics, such as Figure 3 or Figure 4, were used to look at large amounts of data in an abstracted way. The next step was to determine unusual topics within the time span by collecting the most prominent topic of every sentiment category for each day, resulting in three topics per day. These topics were then cross-referenced with the sentiments on the same days and vice versa, to ensure both sides indicated a noteworthy event. As the combination of topic modeling and sentiment analysis could only alert on the importance of certain dates, we referenced news archives for additional context information to identify the specific reasons for the sentiment as well as the topic. By following these steps, we were able to single out a number of interesting dates within November and December of 2024, which will be further discussed in the following chapter.

4 Results & Analysis

In this section, we present our findings, answering our research questions, with a party-specific analysis focusing on differences in sentiment and thematic emphasis across political parties (RQ3). Subsequently, we provide an in-depth examination of the government crisis in the fall of 2024, addressing whether observable shifts in sentiment occurred and to what extent topic modeling can help identify potential causes for these changes (RQ1 and RQ2).

4.1 Findings by Political Party

Across parties, clear differences emerge in the sentiment of posted content. The AfD communicates with the most negative sentiment (34% of posts), followed by independent members (32%) and Die Linke (26%). In contrast, Germany's two major parties tend toward more positive messaging: the CDU/CSU (opposition leader) has the highest percentage of positive posts (50%), followed closely by the SPD (48%) (see Figure 1).

Longitudinal trends reveal further distinctions: Monthly sentiment averages, computed by assigning values of -1 (negative), 0 (neutral), and $+1$ (positive) to each post [27], indicate that the average sentiment expressed by most opposition parties – including AfD, Die Linke, and BSW – is notably more negative, while the governing parties (SPD, Bündnis 90/Die Grünen, and FDP) consistently score above the zero line throughout the legislative period (see Figure 2). These findings align with the results reported by Schmidt et al. [27]. An interesting outlier in this context is the CDU/CSU. Although formally an opposition party, its average sentiment aligns more closely with that of the governing parties, displaying a clearly positive tone.

The overall distribution of predominant topics is also reflected at the level of individual parties, but with some noteworthy focal points. Family policy, for example, ranks among the ten most frequently addressed topics by the SPD, which also places a comparatively stronger emphasis on social finance policy than observed in the overall dataset. Within the CDU/CSU, party-specific issues emerge more prominently, often associated with positive sentiment, as does the topic of digitization.

While there is an “AfD/FDP”-topic, that focuses on these parties, both the SPD and the CDU/CSU appear as keywords in the general political topics next to other terms, demonstrating that both parties are prominent in the political currents, even though one of the two is part of the opposition.³

4.2 Temporal Developments: The 2024 Government Crisis as a Case Study

Here follows an analysis of the most prominent findings surrounding the time of the government crisis at the end of 2024 as a case study, in which the results of the sentiment analysis are intertwined with the topic modeling findings, to show that this method combination can be used to highlight impactful events based on the Instagram posts of the German Parliament members. The analysis draws on the absolute number of posts with each sentiment, on a mean-sentiment of each party, as well as the most frequently used topic for each day in the months of November and December of 2024. Each finding will then be connected to an actual event on the basis of news articles during this period.

November. From the 4th of November on, the number of posts begins to increase until the 7th with roughly 300 negative and neutral posts each, while the number of positive posts is at a local minimum of slightly more than 100. The most prominent topic on this day is “spd-cdu-bundestag-democracy”.

Between the 7th and the 16th, this topic appears as the most frequent one in negative posts on eight out of ten days – with the exception of the 9th and 10th. It also ranks highest in neutral posts on four days and in positive posts on two days. All in all, this was the most present topic in five of these ten days, underlining its importance in this period. The – in contrast to other times and topics – increased number of neutral and negative posts, suggests that posts in this topic most likely correspond to either a critical or negative contexts.

This observation coincides with notable events regarding the coalition breakup: On the first of November, a proposal paper by the FDP was leaked, demanding fundamental revision of political leading decisions, further worsening the already strained relationship of the coalition parties as it would have required major restructuring. On the 6th of November, Olaf Scholz presented his own proposal paper to counter the leaked FDP paper, but it was rejected [29].

As a consequence, the weekly Zeit [42] mentions that Christian Lindner is to be dismissed by the president. The maximum in neutral posts reflects either neutral or mixed reactions to these events. After Christian Lindner accused Olaf Scholz of purposefully breaking the coalition, he was dismissed from his duties as financial minister on the 7th of November, representing the actual break of the coalition [42]. On this day the number of negative posts reaches its peak along with the topic regarding the Bundestag and democracy, indicating that this topic, at the given time, most likely refers to the government crisis.

In accordance with the amount of negative posts on the 7th, the mean sentiment of each party is far more negative than usual, as can be seen in Figure 3. The sentiments vary from -0.9 from the BSW to -0.6 from Die Linke to almost every other party at -0.3 .

³The GitHub repository contains the most frequent topics of the respective party, as well as the most frequent topics of the positive, negative and neutral posts.




Figure 1: Sentiment distribution per party.




Figure 2: Mean sentiment of every party from January 2021 to February 2025. Each year change is marked by a grey line and the election and the Ampelbruch are marked by a red one.

Interestingly, the most neutral party is not the FDP, which has a mean sentiment of around -0.15 , but the SPD at slightly below zero. Though neutral doesn't necessarily mean truly neutral and could instead be a mixture of positive and negative. So it is possible that e.g. the members of the SPD applauded the way the chancellor handled the situation while also potentially criticizing the FDP.

On the 9th of November the most frequently addressed topic was the “racism-humans-victims-remembrance”-topic, most likely referring to the “Reichsprognomnacht”, happening in 1938 and marking the change from discrimination to systematically prosecution and murder of Jewish people and other minorities in Germany. Therefore, the atrocities committed by German people are remembered on this day each year [5].

The newspaper *Zeit* releases an article on the 15th of November, containing a detailed timetable of the coalition breakup, revealing at what point in time the FDP began to plan this collapse [42].

Accordingly, across 15th–16th of November, the coalition-breakup-topic is the most prevalent one. Consequently, the sentiment and topic distribution reflect these controversial actions, indicating when important events took place.

The last notable sentiment development in November 2024 is on the 25th, where the negative posts reach a global maximum of 300 posts. Though the number of neutral posts usually seems to develop in correlation to the negative posts, this spike in negative posts is not accompanied by a similar increase in neutral posts. The sudden increase in posts with negative sentiment can also be seen in the sentiment of each party. Here, a drastic drop to sentiments between -0.7 and 0 can be noticed.

However, this sentiment change seems to be unrelated to the government crisis and instead is, according to the topic modeling, due to the “women-violence-violence woman-girls”-topic. This reflects the fact, that the 25th of November is the international day




Figure 3: Development of the mean-sentiment of each party in November 2024.

against violence against women, the so-called orange day [31]. Furthermore, this topic is even more relevant on this day due to an announcement in the case of Gisèle Pelicot, a woman, that has been severely abused by her ex-husband as well as 50 other men, putting this topic even more in the focus of the public eye [33].

December. The first relevant spike in post numbers appears on the 16th of December, the number of negative posts rises to a, for the negative posts, global maximum, paralleled by the neutral posts. This is mirrored in the party specific mean sentiment (see Figure 4): most parties converge to a neutral to slightly negative sentiment, except for the SPD, which stays positive, and the BSW, which remains negative.

The topic modeling shows that this decrease in sentiment is most likely due to the coalition crisis, as the most represented topic is again the “spd-cdu-bundestag-democracy”-topic. This hypothesis is confirmed as the Vertrauensfrage (engl. question of confidence) that led to the dissolution of the German Parliament was put forward by chancellor Olaf Scholz on this day. Losing this vote enabled the new election in February [34].

Another notable change in the number of positive and neutral posts, which dropped considerably by roughly 150 posts compared to the day before, is on the 21st of December. In contrast, the number of negative posts reaches a local maximum of almost 150. The topics for this day are mostly about “next of kin-thoughts-magdeburg-injured persons”, indicating that this change in sentiment and post numbers is related to the attack on a Christmas market in Magdeburg in the evening of the prior day, on the 20th of December [32].

Summary. Taken together, our party-level analysis shows systematic differences in both sentiment and thematic focus across political actors, the longitudinal analysis reveals distinct sentiment shifts over the course of the coalition period, and by combining sentiment analysis with topic modeling, we were able to link changes in the members of parliaments’ communication to identifiable real-world developments.

5 Discussion

In the following section, we discuss the overall results and highlight key findings that emerged from our analysis. In doing so, we also address the extent to which the combination of sentiment analysis and topic modeling constitutes an appropriate approach for exploring large-scale social media data (RQ1).

RQ3: Are there detectable differences between the parties in terms of the sentiment and thematic focus of their posts?

We observed differences in sentiment and thematic focus across the political parties represented in the German Parliament. In line with the findings of Schmidt et al. [27], it becomes evident that most opposition parties — including AfD, Die Linke, and BSW — exhibit a predominantly negative average sentiment. In contrast, the governing parties — SPD, Bündnis 90/Die Grünen, and FDP — show a more consistently positive sentiment. An interesting outlier in this context is the CDU/CSU, which, despite being in opposition during the analyzed coalition period, aligns more closely with the governing parties in terms of sentiment, displaying a distinctly positive tone. One possible explanation for this deviation lies in the CDU/CSU’s long-standing role in government prior to the analysed coalition period, having held executive power for 16 consecutive years.

Regarding thematic focus, and consistent with the observations of Achmann and Wolff [1] and Hellwig et al. [17], we identified two broad topics that revolve around politics and parliamentary processes in general. Other dominant topics are voluntary work, foreign affairs (especially concerning the European Union and the conflict in the Ukraine), memorial days, racism, public health, holidays like Christmas, climate protection, financial policy and violence against women. In addition, topic modeling allowed the identification of party-specific thematic emphases: social policy emerged as a key topic for the SPD; both Bündnis 90/Die Grünen and CDU/CSU showed positively connoted, party-specific themes; digitalization was rather prominent for the CDU/CSU and FDP; climate protection for Bündnis 90/Die Grünen; and labor market policy for Die Linke. A particularly noteworthy finding is the high




Figure 4: Development of the mean-sentiment of each party in December 2024.

proportion of AfD posts containing external links to other websites.

RQ2: Can shifts in sentiment be observed over the course of the coalition, and can topic modeling help to identify potential causes for these changes?

We were able to identify events and timeframes of interest, such as the drop in sentiment in November 2024, through sentiment analysis (cf. Figure 2). A qualitative deep dive into this month's data revealed four noticeable changes within the sentiment: two rises in negative sentiment on the 7th and 25th of November, the 9th had a noteworthy neutral sentiment and the 16th was strikingly positive.

By using topic modeling for additional context, we could attribute the spikes on the 7th and 16th properly to the coalition crises, while the sentiment from the 9th and the 25th was identified as being unrelated to that. Rather, these instances could be traced back to the commemoration of the "Reichsprogromnacht" and the international day against violence against women, as well as the case of Gisèle Pelicot.

RQ1: How can sentiment analysis and topic modeling be combined to enhance the distant reading of social media data and improve our understanding of political communication?

Although the sentiment analysis can be utilized to discover crucial events due to unusual shifts in the sentiments, it has also become apparent that this alone is not enough to determine the reason for these changes. By supplementing the sentiment analysis with topic modeling, we were able to retrace the sentiment shifts to certain events, associated with the respective topic, as could be seen in the incidents regarding the coalition breakup as well as in the case of the Christmas market attack in Magdeburg or the enhanced focus on the day against violence against women.

Even when the combined methods are not expressive enough, they can help identify relevant dates for a focused close reading, which can be complemented with research in news archives, as we showcased regarding the sentiments coinciding with the FDP-Paper leak in the *Stern*.

Overall, our findings suggest that the combination of sentiment analysis and topic modeling offers a highly promising approach to distant reading a large-scale Instagram datasets, enabling a more nuanced understanding of political communication. This methodological framework holds considerable potential for various fields of application, including journalism, socio-political discourse, and academic research.

5.1 Limitations & Future Work

Our study offers a detailed insight into the sentiments and thematic structures of Instagram posts published by members of the German Parliament. Nonetheless, certain limitations must be acknowledged. In particular, our analysis is restricted to textual content and does not account for images, videos, or stories. Given the visual nature of Instagram and the communicative importance of multimedia content on this platform, the exclusion of these elements represents a relevant constraint. Future research should therefore aim to incorporate such formats, for example through visual sentiment analysis[21], in order to provide a more comprehensive understanding of political communication on Instagram.

An important aspect to consider is that the collected posts constitute part of the official public communication of members of the German Parliament, and therefore reflect their professional activity. Most posts were published between Monday and Friday, with noticeable declines during weekends and parliamentary recesses. Some holidays, such as Christmas and New Year's Eve, may temporarily increase posting activity. Future research should examine the influence of working schedules, weekends, holidays, and recess periods on posting frequency, sentiment, and topics.

Regarding model performance, the sentiment classifier achieves a weighted F1 score of 0.72, indicating moderate performance and suggesting that a non-trivial proportion of individual posts may receive incorrect sentiment labels. Since the gold-standard evaluation was conducted on a random sample of the full corpus without stratification by party or time period, we cannot rule out that classification accuracy varies across parties or communicative styles. This should be taken into account when interpreting party-level

comparisons. The fine-tuning was deliberately limited to publicly available training data to ensure reproducibility and transferability. Future work could address this limitation by conducting stratified evaluations across parties and by expanding the in-domain training data to improve classification accuracy. Similarly, some of the extracted topics exhibit overlap and lack clear separation, which represents a further area for refinement.

Moreover, while the combination of sentiment analysis and topic modeling proved effective for identifying periods of communicative shifts, attributing these shifts to specific real-world events relied on qualitative interpretation supported by news archives. A systematic close reading approach, for instance, by sampling posts from dates with notable sentiment changes and manually verifying their thematic content and event references, could serve to validate and further substantiate the findings of the distant reading. Such a mixed-methods design would strengthen the methodological rigor of the event–sentiment linkage and help to delineate the boundaries of what the computational approach alone can reliably detect.

A further simplification concerns party affiliations: for the present study, each member's final party membership was assigned to all of their posts across the entire legislative period. This affects 19 out of 708 politicians (2.7 %), most notably the ten former members of *Die Linke* who formed *Bündnis Sahra Wagenknecht* (BSW) in January 2024, as well as nine individual switches involving other parties. Given that these cases represent a small fraction of both the accounts and the overall corpus, the impact on aggregate-level findings is expected to be minimal. Nonetheless, future work could assign party labels at the post level based on the documented date of each switch.

Since all posts in our corpus are annotated with both sentiment and topic labels, this enables further studies that systematically link these two methods, answering, for example, the following questions: What sentiment do individual parties express toward specific topics? How did the discourse surrounding prominent political actors change after the coalition broke up? Can instances of inter-party attacks be identified, and which policy issues do they address?

Acknowledgments

The fourth annotator of the gold standard was Oscar Aquite-Peña.

Ethical Considerations and Data Availability

We collected only publicly available content from manually verified Instagram accounts of elected members of the German parliament in the 20th legislative period. No personal or sensitive user data is included. Following Venturini and Rogers [37], we treat scraping as a “necessary evil”.

To support transparency and reproducibility, we release our annotated gold-standard corpus and the code used for this project. All available data can be found at <https://github.com/jessimaye/BundesPosts>. Due to size constraints, the whole corpus will only be accessible on request.

References

- [1] Michael Achmann and Christian Wolff. 2023. Policy issues vs. documentation: Using bertopic to gain insight in the political communication in instagram stories

- and posts during the 2021 german federal election campaign. *Digital Humanities in the Nordic and Baltic Countries Publications* 5, 1 (2023), 11–28.
- [2] Alienmaster. n.d.. German Politicians Twitter Sentiment [Dataset]. https://huggingface.co/datasets/Alienmaster/german_politicians_twitter_sentiment. Hugging Face. Accessed on March 1, 2025.
- [3] Jennifer Bast. 2021. Politicians, Parties, and Government Representatives on Instagram: A Review of Research Approaches, Usage Patterns, and Effects. *Review of Communication Research* 9 (July 2021). <https://www.rcommunicationr.org/index.php/rcr/article/view/108>
- [4] Axel Bruns. 2019. After the ‘APicalypse’: Social Media Platforms and Their Fight Against Critical Scholarly Research. *Information, Communication & Society* 22, 11 (2019), 1544–1566. doi:10.1080/1369118X.2019.1637447
- [5] Bundeszentrale für politische Bildung (bpb). 2023. Novemberpogrom 1938. <https://www.bpb.de/kurz-knapp/hintergrund-aktuell/542301/novemberpogrom-1938/>. Accessed April 14, 2025.
- [6] Santiago Castro. 2017. Fast Krippendorff: Fast computation of Krippendorff’s alpha agreement measure. *GitHub github.com/pln-fing-udelar/fast-krippendorff* (2017).
- [7] Priyavrat Chauhan, Nonita Sharma, and Geeta Sikka. 2021. The emergence of social media data and sentiment analysis in election prediction. *Journal of Ambient Intelligence and Humanized Computing* 12 (2021), 2601–2627.
- [8] Anna De Fina. 2024. Bonding with followers: Chronotopes and scales in political communication on Instagram. *Discourse studies* (Sept. 2024). doi:10.1177/14614456241276717
- [9] Deepset. n.d.. German BERT base [Machine Learning Model]. <https://huggingface.co/deepset/gbert-base>. Hugging Face. Accessed on March 3, 2025.
- [10] Zulfadzli Drus and Haliyana Khalid. 2019. Sentiment analysis in social media and its application: Systematic literature review. *Procedia Computer Science* 161 (2019), 707–714.
- [11] Catalina Goanta, Thales Bertaglia, and Adriana Iamnitchi. 2022. The Case for a Legal Compliance API for the Enforcement of the EU’s Digital Services Act on Social Media Platforms. *arXiv preprint arXiv:2205.06666*. <https://arxiv.org/abs/2205.06666>
- [12] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).
- [13] Maarten Grootendorst. 2024. BERTopic - FAQ. <https://maartengr.github.io/BERTopic/faq.html>. Accessed April 11, 2025.
- [14] Maarten Grootendorst. n.d.. BERTopic on Large Datasets. <https://colab.research.google.com/drive/1W7aEdDPxC29jP99GGZphUlqjMFFVKtBC?usp=sharing>. Google Colab Notebook.
- [15] Guhr, Oliver. n.d.. German Sentiment Bert [Machine Learning Model]. <https://huggingface.co/oliverguhr/german-sentiment-bert>. Hugging Face. Accessed on March 1, 2025.
- [16] Nils Constantin Hellwig, Markus Bink, Thomas Schmidt, Jakob Fehle, and Christian Wolff. 2023. Transformer-based analysis of sentiment towards german political parties on twitter during the 2021 election year. In *Proceedings of the 6th international conference on natural language and speech processing (ICNLSP 2023)*. 84–98.
- [17] Nils Constantin Hellwig, Jakob Fehle, Markus Bink, Thomas Schmidt, and Christian Wolff. 2024. Exploring Twitter discourse with BERTopic: topic modeling of tweets related to the major German parties during the 2021 German federal election. *International Journal of Speech Technology* 27, 4 (2024), 901–921.
- [18] Hugging Face. n.d.. Fine-tune a pretrained model. <https://huggingface.co/docs/transformers/training>. Hugging Face. Accessed on March 3, 2025.
- [19] Kamran Kowsari, Kiana Jafari Meimandi, Mojtaba Heidarysafa, Sanjana Mendu, Laura Barnes, and Donald Brown. 2019. Text Classification Algorithms: A Survey. *Information* 10, 4 (2019). doi:10.3390/info10040150
- [20] Caitlin Doogan Poet Laureate, Wray Buntine, and Henry Linger. 2023. A systematic review of the use of topic models for short text social media analysis. *Artificial Intelligence Review* 56, 12 (2023), 14223–14255.
- [21] Alessandro Ortis, Giovanni Maria Farinella, and Sebastiano Battiato. 2020. Survey on visual sentiment analysis. *IET Image Processing* 14, 8 (2020), 1440–1456. doi:10.1049/iet-ipr.2019.1270
- [22] Stijn Peeters. 2023. *Zeeschuimer*. doi:10.5281/zenodo.7525702
- [23] Yilang Peng. 2021. What makes politicians’ Instagram posts popular? Analyzing social media strategies of candidates and office holders with computer vision. *The International Journal of Press/Politics* 26, 1 (2021), 143–166.
- [24] Vincent Raynaud and Mireille Lalancette. 2023. Social media, visuals, and politics: a look at politicians’ digital visual habitus on Instagram. In *Research handbook on visual politics*. Edward Elgar Publishing, 167–180.
- [25] Margarita Rodríguez-Ibáñez, Antonio Casánz-Ventura, Félix Castejón-Mateos, and Pedro-Manuel Cuenca-Jiménez. 2023. A review on sentiment analysis from social media platforms. *Expert Systems with Applications* 223 (2023), 119862.
- [26] Joni O. Salminen, Hind A. Al-Merekhi, Partha Dey, and Bernard J. Jansen. 2018. Inter-Rater Agreement for Social Computing Studies. In *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*. 80–87. doi:10.1109/SNAMS.2018.8554744

- [27] Thomas Schmidt, Jakob Fehle, Maximilian Weissenbacher, Jonathan Richter, Philipp Gottschalk, and Christian Wolff. 2022. Sentiment analysis on twitter for the major German parties during the 2021 German federal election. In *Proceedings of the 18th conference on natural language processing (konvens 2022)*. KONVENS 2022 Organizers, 74–87.
- [28] Ssary. n.d.. XLM RoBERTa German sentiment [Machine Learning Model]. <https://huggingface.co/ssary/XLM-RoBERTa-German-sentiment>. Hugging Face. Accessed on March 1, 2025.
- [29] Stern. 2024. Grundsatzpapier: Christian Lindner fordert die Koalition heraus. <https://www.stern.de/politik/deutschland/grundsatzpapier--christian-lindner-fordert-die-koalition-heraus-35192098.html>. Accessed April 13, 2025.
- [30] TabularisAI. n.d.. Multilingual Sentiment Analysis [Machine Learning Model]. <https://huggingface.co/tabularisai/multilingual-sentiment-analysis>. Hugging Face. Accessed on March 1, 2025.
- [31] Tagesschau. 2023. Gewalt gegen Frauen: Ein Problem in Deutschland. <https://www.tagesschau.de/inland/gesellschaft/gewalt-gegen-frauen-116.html>. Accessed April 14, 2025.
- [32] Tagesschau. 2024. Magdeburg: Weihnachtsmarkt in der Stadt. <https://www.tagesschau.de/inland/gesellschaft/magdeburg-weihnachtsmarkt-104.html>. Accessed April 13, 2025.
- [33] Tagesschau. 2024. Pelicot-Prozess: Frankreichs Justiz vor der Herausforderung. <https://www.tagesschau.de/ausland/europa/frankreich-pelicot-prozess-100.html>. Accessed April 14, 2025.
- [34] Tagesschau. 2024. Scholz verliert Vertrauensfrage. <https://www.tagesschau.de/inland/scholz-verliert-vertrauensfrage-100.html>. Accessed April 13, 2025.
- [35] Martino Trevisan, Luca Vassio, Idilio Drago, Marco Mellia, Fabricio Murai, Flavio Figueiredo, Ana Paula Couto da Silva, and Jussara M Almeida. 2019. Towards understanding political interactions on Instagram. In *Proceedings of the 30th ACM conference on hypertext and social media*. 247–251.
- [36] Ike Vayansky and Sathish AP Kumar. 2020. A review of topic modeling methods. *Information Systems* 94 (2020), 101582.
- [37] Tommaso Venturini and Richard Rogers. 2019. “API-Based Research” or How can Digital Sociology and Journalism Studies Learn from the Facebook and Cambridge Analytica Data Breach. *Digital Journalism* 7, 4 (April 2019), 532–540. doi:10.1080/21670811.2019.1591927
- [38] Wikipedia. 2024. Liste der Mitglieder des Deutschen Bundestages (20. Wahlperiode). [https://de.wikipedia.org/wiki/Liste_der_Mitglieder_des_Deutschen_Bundestages_\(20._Wahlperiode\)](https://de.wikipedia.org/wiki/Liste_der_Mitglieder_des_Deutschen_Bundestages_(20._Wahlperiode)). Accessed December 16, 2025.
- [39] Michael Wojatzki, Eugen Ruppert, Sarah Holschneider, Torsten Zesch, and Chris Biemann. 2017. Germeval 2017: Shared task on aspect-based sentiment in social media customer feedback. *Proceedings of the GermEval* (2017), 1–12.
- [40] Chen Yang. 2021. Research in the Instagram context: Approaches and methods. *The Journal of Social Sciences Research* 7, 1 (2021), 15–21.
- [41] ZDFheute. 2024. Ampel-Aus vorbereitet? FDP weist Vorwurf zurück. <https://www.zdf.de/nachrichten/politik/deutschland/fdp-ampel-bruch-vorbereitung-100.html>. Accessed: 2025-04-10.
- [42] Zeit. 2024. Christian Lindner: Die Ampel-Aus aus der FDP-Bundesregierung. <https://www.zeit.de/politik/deutschland/2024-11/christian-lindner-ampel-aus-fdp-bundesregierung/komplettansicht>. Accessed April 13, 2025.