

Analyzing German Parliamentary Speeches: A Machine Learning Approach for Topic and Sentiment Classification

Research Paper

Lukas Pätz¹, Moritz Beyer¹, Jannik Späth¹, Lasse Bohlen^{1,2}, Patrick Zschech^{1,2},
Mathias Kraus³, and Julian Rosenberger³

¹ Leipzig University, Leipzig, Germany
{lukas.paetz, zu66evit, j.spaeth}@studserv.uni-leipzig.de

² TU Dresden, Dresden, Germany
{lasse.bohlen, patrick.zschech}@tu-dresden.de

³ University of Regensburg, Regensburg, Germany
{mathias.kraus, julian.rosenberger}@ur.de

Abstract. This study investigates political discourse in the German parliament, the Bundestag, by analyzing approximately 28,000 parliamentary speeches from the last five years. Two machine learning models for topic and sentiment classification were developed and trained on a manually labeled dataset. The models showed strong classification performance, achieving an area under the receiver operating characteristic curve (AUROC) of 0.94 for topic classification (average across topics) and 0.89 for sentiment classification. Both models were applied to assess topic trends and sentiment distributions across political parties and over time. The analysis reveals remarkable relationships between parties and their role in parliament. In particular, a change in style can be observed for parties moving from government to opposition. While ideological positions matter, governing responsibilities also shape discourse. The analysis directly addresses key questions about the evolution of topics, sentiment dynamics, and party-specific discourse strategies in the Bundestag.

Keywords: Natural Language Processing, German Parliamentary, Discourse Analysis, Bundestag

1 Introduction

In today’s digital age, political discussions primarily take place on social media and in televised debates, where public discourse is often shaped by rapid exchanges of opinions (Emmer 2016). However, the parliament remains the primary institution where formal political decision-making occurs and legislative processes unfold. While digital platforms amplify public opinion, it is in the parliament where policies are debated, negotiated and enacted – making it a crucial forum for understanding political discourse. Germany, with its population of over 80 million, is not only one of the largest economies in the world (Sahoo et al. 2024), but also a country where political decisions have

far-reaching national and international consequences. The German parliament, known as the Bundestag, plays an important role in shaping the nation's political discourse and serves as a center for critical debates on various societal issues, and it facilitates rigorous discussions on these topics (Deutscher Bundestag 2025).

Elected representatives utilize this platform to articulate public concerns, commend policy achievements, and critically evaluate governmental actions. During the legislative period between 2021 and 2025, the 20th German Bundestag was the largest freely elected national parliament in the world with 733 members. Each year, approximately 80 plenary sessions are held, during which several thousand speeches on diverse topics are delivered (Deutscher Bundestag 2025). These public debates on key political issues provide a rich foundation of textual data for analysis.

In recent years, a perceived divide between the political elite and the general population in Germany has grown. The German phrase "die da oben" (those up there), often used metaphorically to refer to political decision-makers such as members of parliament, has regained prominence as a symbol of growing public distrust (Hamann 2016). Studies and surveys highlight a growing skepticism among the population towards political institutions and their actors. According to the Körber-Stiftung (2024), 54% of German respondents reported having little or no trust in democracy. In 2019, this figure stood at just 30%, highlighting a dramatic and accelerating loss of confidence in Germany's democratic system over the past years (Körber-Stiftung 2023). This development is confirmed by the European Commission (2024), which shows that by early 2024, only 33% of EU citizens expressed trust in their national governments. Germany, which had traditionally ranked above the EU average, is now among the countries with the sharpest declines in trust. These developments reflect a broader demand for increased transparency, accountability, and comprehensibility in political processes.

Many studies have examined the influence of the media and the public's changing tone on this development (Boulianne 2020, Lorenz-Spreen et al. 2023). While media coverage and social media trends often highlight selective aspects of political discourse, Bundestag debates offer unfiltered insights into policy discussions. Analyzing parliamentary discourse on a large scale reveals the priorities, ideologies, and rhetorical strategies of political actors. From an academic perspective, this enhances the understanding of institutional communication, and from a societal perspective, it supports transparency and accountability by showing how parties address key issues over time.

Modern Natural Language Processing (NLP) techniques offer the possibility of analyzing these debates on a large scale, identifying dominant topics and assessing their tone (Abercrombie & Batista-Navarro 2020).

Understanding which topics dominate Bundestag debates, how they are framed by different political parties, and whether discussions are characterized by support or criticism is essential for assessing the dynamics of political communication. Our study investigates how topics and sentiments in Bundestag speeches evolved between 2019 and 2024. The overarching research question guiding this work is:

Research Question: *What patterns in the evolution of topics and sentiments can be identified in Bundestag speeches using NLP techniques?*

To explore this question, our analysis focuses on three key aspects:

1. Which topics were most prominent, and how did their relevance shift over time?
2. How did the sentiment of speeches change, particularly with shifts between government and opposition roles?
3. Did parties use different sentiments when addressing the topics they emphasized?

To address our research questions, we apply machine learning-based classification techniques to Bundestag speeches from two complementary perspectives. First, topic classification allows us to analyze which issues dominate parliamentary debates, how their prominence shifts over time, and how parties differ in their thematic focus and ideological framing. Second, sentiment classification enables us to explore how parties express positive or negative attitudes toward these topics, revealing patterns in political tone, framing, and changes in sentiment across time and party lines.

It is important to emphasize that this study does not aim to endorse any political position. Instead, the goal is to provide tools for objectively analyzing the political discourse in Germany using quantitative data. By doing so, this study seeks to contribute to a deeper understanding of parliamentary discussions. To support further research and ensure transparency, the full codebase and all associated data are also made available.¹

The remainder of this paper is structured as follows. Section 2 provides a review of relevant literature, thereby establishing the theoretical foundation for the methodology employed and situating the work within existing research. Section 3 outlines the steps involved in generating, annotating and developing the model, as well as external validation. This is followed by a brief description of how the models were applied to the full speech dataset for trend analysis. Section 4 presents the empirical findings, which are subsequently discussed in Section 5, where we examine their broader implications, acknowledge the limitations of the study, and propose directions for future research.

2 Related Work

Transcripts of legislative debates provide rich material for analyzing the positions and rhetorical strategies of political actors and their parties on key societal issues. In the context of the German Bundestag, such transcripts offer unique insights into how political discourse is structured within a formal legislative setting. While interest in parliamentary speech has grown across disciplines, the Bundestag serves as a particularly valuable case due to its institutional transparency, large volume of data, and diverse party landscape (Laver et al. 2003). Political and social scientists, who have traditionally relied on expert coding for analyzing such texts, are increasingly adopting computational approaches, enabling a more systematic analysis of political positions (Abercrombie & Batista-Navarro 2020). These advancements allow for a more data-driven examination of parliamentary discourse and ideological patterns (Rauh 2018).

Building on these computational advances, text classification, which is a foundational task in both manual and automated text analysis (Debortoli et al. 2016, Hofmann et al. 2024), emerges as a key method for analyzing political discourse. In the context of political discourse, topic and sentiment classification are particularly relevant.

¹ <https://github.com/ptzlukas/german-parliamentary-speech-classifier>

Topic classification seeks to assign textual content to specific thematic categories or topics (Nasution & Onan 2024). Topics represent the primary subject matter or prevailing themes (Koltsova & Koltcov 2013) within a given text. Often those texts are derived from sources such as social media platforms, news articles, or consumer behavior data (Akteer et al. 2016). These topics can either be predefined or derived from a pool of texts. This process of topic discovery can be carried out using a range of computational approaches, including clustering-based methods (Ogunleye et al. 2023) and machine learning techniques (Shah et al. 2021), which facilitate the identification and monitoring of emerging trends (Kim et al. 2016).

Sentiment classification is a specific form of text classification that categorizes texts based on the emotional or attitudinal orientation of expressed opinions (Devika et al. 2016). It can be performed through various approaches, including the use of lexicon-based methods that count emotionally charged words, the application of conventional machine learning models such as logistic regression trained on word frequency features or the deployment of advanced neural networks (Stine 2019). The rapid expansion of social media has significantly amplified the accessibility of public opinions and sentiments, thereby establishing sentiment classification as an essential tool for understanding public attitudes across diverse fields including business, politics, and beyond (Tan et al. 2023).

While topic classification aims to uncover latent thematic structures within document collections (Eickhoff & Neuss 2017), sentiment classification is used to extract opinions and attitudes from texts, including those from social media, blogs, or parliamentary debates (Storey & Park 2022). Recent advances focus on finer-grained approaches that account for the multi-aspect nature of sentiment – where a single text can express different sentiments towards various entities or subjects – which is especially pertinent for analyzing complex political speech (Lu et al. 2011).

Over the past decade, interest in parliamentary text corpora has grown steadily, resulting in a series of national projects. One of the first was "DutchParl", introduced by Marx et al. (2010), covering debates from the Netherlands, Flanders, and Belgium. Shortly thereafter, Jakubíček & Kovár (2010) released "CzechParl", based on plenary session transcripts from the Czech parliament. These initiatives were followed by corpora for Finland (Mansikkaniemi et al. 2017), Greece (Dritsa et al. 2022), Poland (Ogrodniczuk 2018), Norway (Fiva et al. 2025), and Turkey (Güngör et al. 2018), illustrating the increasing international interest in the systematic analysis of legislative discourse.

In recent years, various projects have aimed to develop a comprehensive corpus encompassing the parliamentary debates of the German Bundestag. Barbaresi (2018) compiles speeches delivered by the German presidents, the Bundestag presidents, the German chancellors, and the foreign ministers from the years 1982 to 2018. Additionally, several initiatives have focused on integrating parliamentary records from multiple countries. "ParlSpeech V2" provides parliamentary protocols from the national parliaments of Austria, Germany, Denmark, and other countries over various legislative periods (Rauh & Schwalbach 2020), while "GerParCor" focuses on protocols and speeches from the national parliaments of Austria, Switzerland, Liechtenstein, and Germany (Abrami et al. 2022). Moreover, Richter et al. (2023) present a machine-readable dataset of speeches from German Bundestag sessions, covering nearly 900,000 speeches from 1949 to 2022. The project provides a structured, maintainable data format enriched with

political context, enabling advanced text analysis and supporting diverse political science research.

Building on this existing body of work, the following section details the specific methodological approach we developed to systematically analyze discourse patterns in the Bundestag and address our research questions.

3 Methods

3.1 Analytical Approach

Our study uses a two-step supervised machine learning approach to analyze a large corpus of parliamentary speeches from the German Bundestag. First, topic and sentiment classifiers were trained and evaluated on a manually labeled subset, with domain-specific annotations ensuring high-quality training data. This allowed the models to capture nuances of political communication and party behavior. In the second step, the trained models were applied to the full dataset of approximately 28,000 speeches by 963 speakers across seven parties. The resulting thematic and sentiment predictions were integrated with metadata (e.g., speaker, party, date), creating a comprehensive annotated dataset.

Unlike studies using generic labels, this approach ensures contextual relevance through annotations tailored to the parliamentary domain. Combining manual labeling with scalable machine learning enables systematic, large-scale analysis of discourse patterns. The resulting dataset offers an empirical basis for examining trends in political language and behavior across parties and over time.

3.2 Dataset Generation and Annotation

Supervised machine learning in NLP involves training algorithms on labeled text data so they can learn patterns such as topics or sentiment and accurately classify or predict similar patterns in new, unseen texts (Singh et al. 2016). To create such a labeled dataset we used the Documentation and Information System for Parliamentary Materials (DIP)² interface to download Bundestag plenary meeting minutes from October 2019 to October 2024, resulting in 393 downloaded plenary protocols. Subsequently, individual speeches and relevant fields were extracted from those protocols using regular expressions to distinguish speeches by speaker name, party affiliation, and position because multiple speeches are collected within one plenary protocol (Chapman & Stolee 2016).

Data preprocessing involved noise removal, tokenization, lowercase conversion, lemmatization, and partly speech-based stopword removal (via SpaCy) (Smelyakov et al. 2020). A custom stopword list for German parliamentary debates was used³ for topic classification but omitted for sentiment classification to preserve model performance.

Manual labeling was used to create a ground truth, which is essential for effective sentiment and topic classification. The authors conducted the labeling manually for both topic and sentiment based on a randomly selected subset of speeches. Due to the thematic

² <https://search.dip.bundestag.de/api/v1/swagger-ui/>

³ <https://github.com/ptzluks/german-parliamentary-speech-classifier/blob/main/resources/stopwords.txt>

complexity of many of the speeches, the selected samples were cross-checked by other group members to ensure consistency and reduce subjectivity. The number and scope of topic labels were guided by insights from linear discriminant analysis based topic classification (Blei et al. 2003), which revealed coherent thematic clusters. These clusters were then interpreted and consolidated in a way that reflects the committee structure of the German Bundestag,⁴ ensuring both statistical grounding and political relevance. The procedure results in six categories: Environment, Social Affairs and Education, Economy and Finance, Foreign and Security Policy, Infrastructure and Transport, and Health.

The sentiment classification used binary labels: positive and negative. Positive sentiment refers to supportive, solution-oriented language. Negative sentiment denotes critical or evaluative statements, such as disapproval of policies or actors, and highlights conflicts or failures. Although we use the term negative, it reflects critical tone rather than emotional negativity, aligning with common terminology in machine learning (Hofmann et al. 2024). As truly neutral speeches are rare, they were excluded. Sentiment and topic analyses used the same sample.

Following these procedures, we curated a labeled dataset of 557 speeches for topic classification and 445 speeches for sentiment classification. The difference in sample size is due to the greater complexity of the multi-class topic task, which required additional labeled examples to ensure robust model training. Table 1 presents an exemplary data entry.

Table 1. Sample data entry with metadata, assigned topic and sentiment, and speech excerpt (abridged)

Metadata		Speech
U-id	5284	Frau Präsidentin! [...] Mit dem Klimaschutzgesetz und dem Emissionshandelsgesetz wird ein weiteres Bürokratiemonster geschaffen, das Bürgerinnen und Bürger finanziell belastet. Die Kosten werden auf die Verbraucher abgewälzt, ohne dass die Wirkung dieser Maßnahmen belegt wäre. [...] Dieses Gesetz basiert auf einer ideologisch motivierten Hypothese und dient letztlich der Transformation unserer Gesellschaft in eine ökosozialistische Diktatur.
Session-id	19/125	
Date	2019-11-08	
Speaker	Karsten Hilse	
Party	AfD	
Position	Abgeordneter	
Topic	1 (Environment)	
Sentiment	-1 (negative)	

3.3 Model Development and Validation

Separate models were developed for sentiment and topic classification, each trained and evaluated on the manually labeled dataset. Given the distinct nature of these two classification tasks, we applied different strategies tailored to their specific requirements.

Topic classification was approached as a multi-class task, where speeches could be assigned to a single category. This required a more semantically rich representation. A custom Word2Vec embedding model was trained on the whole corpus of parliamentary

⁴ <https://www.bundestag.de/en/committees>

speeches, enabling the capture of domain-specific language patterns (Lilleberg et al. 2015). These embeddings outperformed generic pre-trained alternatives. Several classifiers were evaluated, including logistic regression, XGBoost, and a Bagging ensemble with a Support Vector Machine (SVM) base estimator. To obtain an overall score, the area under the receiver operating characteristic curve (AUROC) values were averaged over all topic classes.

Sentiment classification was performed using Term Frequency-Inverse Document Frequency (TF-IDF) vectorization based on unigrams to transform the speeches into numerical feature representations. This frequency-based method was selected due to its efficiency and strong performance in text classification tasks with limited label complexity (Rosenberger et al. 2025). Several classifiers were evaluated, including Random Forest (Breiman 2001), XGBoost (Chen & Guestrin 2016), and logistic regression.

Model Evaluation. All models were trained and tested using an 80:20 train-test split to ensure robustness and generalizability. Performance was evaluated on the test set using three metrics: accuracy, F1-score, and the AUROC. AUROC quantifies a classifier’s ability to distinguish between classes across all threshold levels, with a value of 1.0 indicating perfect separation and 0.5 representing random guessing (Fawcett 2006). AUROC served as the primary evaluation metric, as it provides a threshold-independent measure of model performance – especially relevant when comparing classifiers with different decision thresholds, where accuracy and F1-score can vary substantially (Huang & Ling 2005). Accuracy and F1-score complemented this by capturing overall correctness and the balance between precision and recall, respectively.

To assess the *generalizability* of our machine learning models beyond their original training domain, we conducted a small validation study using political speech transcripts from the Austrian National Council (Nationalrat). A manually labeled sample of 80 Austrian speeches served as ground truth to test both topic and sentiment classification.

3.4 Application to Field Data

The best-performing models, evaluated on the test set, were applied to the full text corpus. Using the predicted labels, we analyzed relative topic distribution over time, party-specific sentiment trends, and sentiment variation by topic within parties. These models provide the analytical foundation for answering the research question through structured, large-scale classification of thematic and sentimental dimensions in parliamentary discourse.

4 Results

4.1 Model Results

Building on the model development outlined in Section 3.3, both the sentiment and topic classification models achieved strong performance. Table 2 presents an overview of the evaluated classifiers for both tasks. The results highlight the relative strengths of each approach in terms of accuracy, F1-score, and AUROC, providing the basis for model selection in each case.

Table 2. Performance of evaluated classifiers for topic and sentiment classification

Topic Classification				Sentiment Classification			
Model	Accuracy	F1	AUROC	Model	Accuracy	F1	AUROC
Logistic Regression	0.62	0.62	0.90	Logistic Regression	0.76	0.76	0.85
XGBoost	0.63	0.63	0.90	XGBoost	0.71	0.70	0.78
Bagging + SVM	0.68	0.68	0.94	Random Forest	0.85	0.85	0.89

For *topic classification*, the best-performing model was the Bagging classifier combined with custom Word2Vec embeddings. It achieved an accuracy and F1-score of 0.68. AUROC analysis further supports the model’s effectiveness, with an average AUROC of 0.94 across the six topic classes, indicating strong separability even in the presence of thematic overlap.

For *sentiment classification*, the Random Forest classifier reached an accuracy and F1-score of 0.85. The AUROC value of 0.89 further confirms the model’s ability to robustly distinguish between positive and negative sentiments.

Additionally, both models showed reasonable *generalizability* to Austrian parliamentary data. The topic classifier achieved an AUROC of 0.84, down from 0.94 in the original domain, while the sentiment classifier reached 0.78, compared to 0.89.

Despite maintaining decent discriminative power, both models – especially the sentiment classifier – showed lower accuracy and F1-scores. Sentiment accuracy dropped to 0.49, indicating difficulties with threshold calibration in the new context. A closer look revealed a clear asymmetry in misclassifications: genuinely positive speeches were often labeled as negative, while no negative speeches were predicted as positive. This may reflect a distinct tone in Austrian parliamentary discourse, where different expressions are used.

4.2 Field Data Trends

This section presents the empirical results derived from applying the classification models to the full Bundestag dataset. The analysis directly addresses the research question introduced earlier, with a particular focus on the three subquestions: (1) identifying dominant topics and their variation across parties and time, (2) examining shifts in sentiment in response to political role changes, and (3) analyzing how specific topics are framed emotionally by different parties. The results are presented in the form of visual analyses based on the enriched parliamentary speech dataset described in Section 3.4, providing an in-depth view of discourse patterns in the German Bundestag.

Topic Distribution Over Time. Figure 1 illustrates the distribution of political discussions in the Bundestag across different topics over time. From this analysis, it becomes evident that three main topics, Social Affairs and Education, Economy and Finance, and Foreign and Security Policy, together make up around 70% of the discussions. These three topics dominate the political debates in the German Parliament, reflecting their central importance in the political landscape. In contrast, topics such as Environment, Health, and Infrastructure and Transportation each account for roughly 10% of the total discussions, suggesting that they occupy a smaller but still relevant space in the

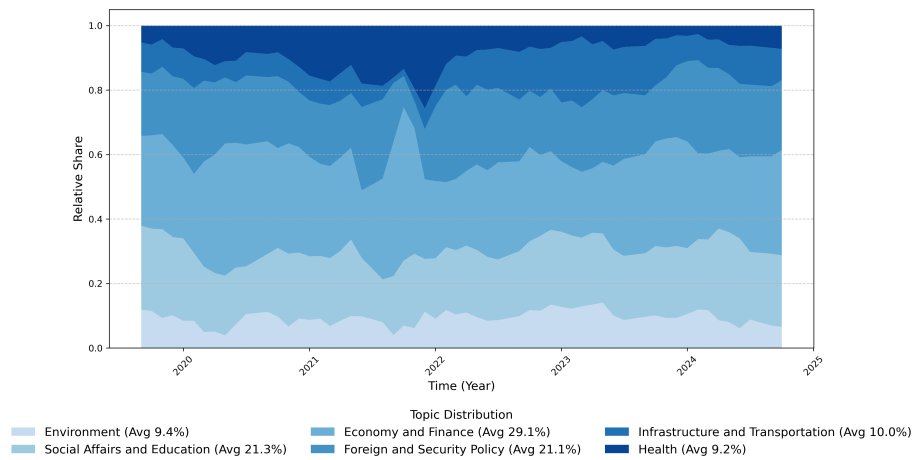


Figure 1. Topic distribution over time

political dialogue. The discussion of Health, in particular, showed a notable fluctuation. During the COVID-19 pandemic, Health-related topics were discussed much more frequently than after the end of the pandemic in early 2022. This shift highlights the substantial impact that the pandemic had on the political agenda. Furthermore, towards the end of 2021, Economy and Finance saw a marked increase in discussions. This surge can be attributed to the formation of the new coalition government (coalition between Sozialdemokratische Partei Deutschlands (SPD), Bündnis 90/Die Grünen and Freie Demokratische Partei (FDP)), which faced substantial challenges related to budget planning. The government's budget had to be passed under the constraints of the debt brake, which had been temporarily suspended during the COVID years (Connolly 2023). This budget was later declared unconstitutional by the Federal Constitutional Court (Bundesverfassungsgericht 2023), which may have further contributed to the heightened discussions around Economic and Financial matters at that time.

Sentiment Distribution Over Time. Figure 2 presents the relative share of negatively labeled speeches per party. This analysis breaks down the distribution of negative sentiments, with each party being represented by its own curve. This visualization provides insight into how the sentiment in speeches has evolved over time for each political party, allowing for a comparative assessment of how parties have shifted in terms of negative speech over the period of observation. One of the most notable observations is the shift in sentiment following the September 2021 elections. Prior to the election, CDU/CSU, as part of the governing Coalition, exhibited a relatively low share of negative speeches. However, after moving into opposition, this share increased sharply. Conversely, when FDP and Bündnis 90/Die Grünen joined the governing coalition, their proportion of critical statements fell markedly. This pattern suggests that parties in government generally express fewer negative sentiments, likely reflecting their perception of their own policies in a more favorable light. In contrast, opposition parties tend to take a more critical stance, resulting in a higher proportion of negative speeches. A comparison between the two governing coalitions further supports this observation.

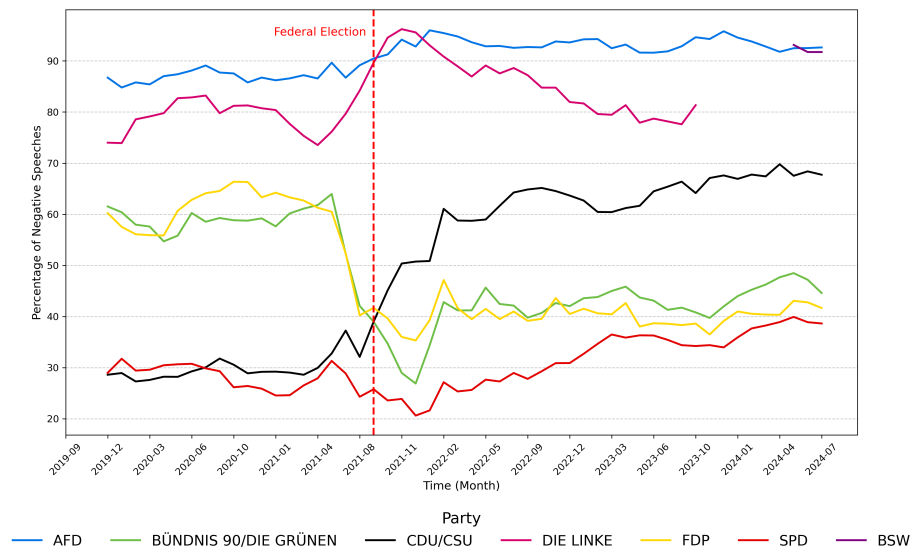


Figure 2. Percentage of negative speeches over time by party

During the 19th legislative period (2017–2021), the governing coalition consisted of CDU/CSU and SPD. During this time, approximately 30% of speeches from coalition parties carried a negative sentiment. In contrast, the subsequent coalition – comprising SPD, Bündnis 90/Die Grünen, and FDP – governing during the 20th legislative period (from 2021), exhibited a notably higher share of negative speeches of around 40%. This may indicate that this coalition was more critical of its own policies than its predecessor, reflecting a greater degree of internal tension or dissatisfaction within the alliance. These inner conflicts eventually culminated in the premature dissolution of the coalition in November 2024 (Meakem 2024). Another key observation is that parties on the political extremes tend to express more negative sentiments. The AfD, widely regarded as the most right-wing party in the German Bundestag, consistently delivers over 90% of its speeches with a negative tone. Similarly, Die Linke, considered the most left-wing party, uses a negative tone about 80% of the time. This high level of negativity not only reflects the parties' positions at the extreme end of the political spectrum, but also underlines their role as sharp critics of the ruling government.

Sentiment Distribution per Topic and Party. Figure 3 combines both models to conduct a joint analysis. The top part displays the average share of negatively labeled speeches for each party (with total counts in parentheses). Below that, deviations in sentiment per topic are visualized: a positive deviation (blue) indicates that the share of negative speeches for a given topic is higher than the party's own average across all topics, while a negative deviation (orange) indicates a lower share. Importantly, these deviations are always calculated relative to each party's individual baseline—not the overall Bundestag average. This enables a relative within-party comparison, showing how critically a party discusses specific policy areas compared to its typical tone.

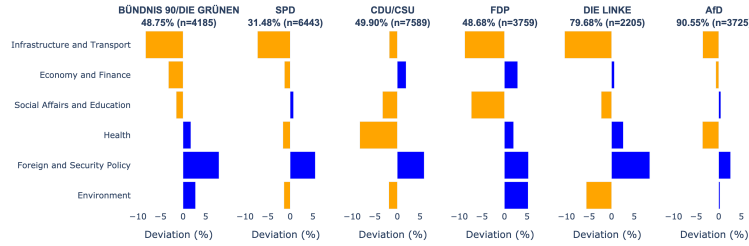


Figure 3. Deviation sentiment distribution per topic and party

One key observation is that Infrastructure and Transportation is evaluated more positively across all parties relative to their respective baselines, suggesting lower political contention in this area. In contrast, Foreign and Security Policy tends to be discussed more negatively, especially by opposition parties. This heightened negativity may reflect dissatisfaction with the current political direction in this domain. However, it is important to note that these sentiment deviations do not necessarily reflect topic-specific critiques, but may instead indicate broader discontent with the political status quo—particularly among opposition actors.

Another notable finding is the thematic focus of individual parties. CDU/CSU and FDP, traditionally associated with market-oriented policies, show higher negativity in the area of Economy and Finance. SPD places emphasis on Social and Education topics, while Bündnis 90/Die Grünen dedicate more attention to Environment. Although the frequency of topics is not depicted here, these patterns correspond to well-known party priorities reflected in broader analyses. Interestingly, the intensity of criticism tends to increase within parties’ core issue areas. This suggests that parties use sentiment not just to evaluate, but to strategically frame their priorities and underline the need for political action. Overall, these findings highlight the strategic nature of political discourse in the Bundestag: parties selectively focus on topics and adjust the tone of their rhetoric to align with their broader political objectives.

5 Discussion

Our study presents an approach for automated classification of parliamentary speeches, contributing both methodologically and empirically to the computational analysis of political discourse. First, a dataset of speeches from the German Bundestag was collected, manually annotated, and preprocessed with respect to domain-specific challenges.

Second, the study demonstrates that machine learning methods, if adapted to the linguistic and structural characteristics of political speech, can yield meaningful insights. The successful application of different classification strategies tailored to sentiment and topic detection shows the practical value of combining established NLP techniques with contextual knowledge. The approach of developing domain-specific classifiers offers a potentially transferable framework for analyzing other structured text domains, such as political communication, media analysis, or institutional reporting.

Third, the application of the trained models to real-world Bundestag data enables large-scale empirical analysis and reveals systematic patterns in parliamentary discourse. Debates are dominated by a few key policy areas such as Economy and Finance, Social and Education, and Foreign and Security Policy. Sentiment patterns differ substantially between government and opposition parties, with opposition parties consistently expressing more negative sentiment overall. Additionally, parties tend to adopt a particularly critical stance within their own core issue areas, suggesting a strategic use of negativity to emphasize political priorities. The resulting dataset, enriched based on model predictions, offers new possibilities for tracking discourse dynamics over time and across party, thereby expanding the analytical scope for political science and discourse studies.

From a theoretical perspective, our study shows that machine learning can be effectively applied to structured political text when domain-specific nuances are considered. Unlike general NLP tasks, classifying parliamentary speech demands sensitivity to institutional language, formal rhetoric, and the hybrid nature of political expression, which shifts between critique and consensus. Custom embeddings and task-specific feature engineering were essential for capturing these aspects.

In contrast to previous work that has relied on unsupervised methods for the analysis of parliamentary texts, such as the study by Rauh (2018), our study adopts a supervised approach grounded in domain-specific manual annotations. This enables a precise classification of sentiment and topics, particularly in complex or institutionally embedded speech segments. By combining the advantages of manual labeling with scalable machine learning techniques, the approach presented extends and refines existing methodologies.

Despite the promising results, several limitations must be acknowledged. First, model performance, though high in relative terms, still leaves room for improvement. Particularly in topic classification, overlapping themes and ambiguous phrasing contributed to misclassifications, suggesting potential gains from multi-label classification or more advanced machine learning models. Second, the manual labeling process, despite internal validation and cross-checking, remains subjective, especially where speeches touch upon multiple thematic areas or where sentiment is expressed in subtle rhetorical forms.

Third, the size of the labeled training data remains limited. While the overall corpus comprises over 28,000 speeches, only a small subset of entries was used for supervised model training. This could constrain model robustness, particularly for underrepresented topics. Future work should focus on expanding the labeled dataset – either through crowdsourcing, active learning, or semi-supervised approaches – to enhance both coverage and performance. To facilitate further research, all datasets and code used in this study are publicly available.⁵

In conclusion, our study demonstrates the feasibility and value of applying machine learning to large-scale political speech data. Our work contributes to a deeper understanding of parliamentary discourse and offers a framework for future machine learning-based political science research. Beyond methodological contributions, our findings answer our research questions by revealing empirical patterns in the evolution of political topics, sentiment dynamics over time, and party roles. They also reveal the emotional framing of different areas. These results highlight the value of computational techniques in generating evidence-based insights into parliamentary communication.

⁵ <https://github.com/ptzlukas/german-parliamentary-speech-classifier>

References

- Abercrombie, G. & Batista-Navarro, R. (2020), 'Sentiment and position-taking analysis of parliamentary debates: a systematic literature review', *Journal of Computational Social Science* **3**(1), 245–270.
- Abrami, G., Bagci, M., Hammerla, L. & Mehler, A. (2022), German Parliamentary Corpus (GerParCor), in 'Proceedings of the Thirteenth Language Resources and Evaluation Conference', Marseille, France, pp. 1900–1906.
- Akter, S., Bhattacharyya, M., Wamba, S. F. & Aditya, S. (2016), 'How does social media analytics create value?', *Journal of Organizational and End User Computing (JOEUC)* **28**(3), 1–9.
- Barbaresi, A. (2018), A corpus of german political speeches from the 21st century, in '11th Language Resources and Evaluation Conference (LREC 2018)', European Language Resources Association (ELRA), pp. 792–797.
- Blei, D. M., Ng, A. Y. & Jordan, M. I. (2003), 'Latent dirichlet allocation', *Journal of machine Learning research* **3**(Jan), 993–1022.
- Boulianne, S. (2020), 'Twenty years of digital media effects on civic and political participation', *Communication research* **47**(7), 947–966.
- Breiman, L. (2001), 'Random forests', *Machine learning* **45**, 5–32.
- Bundesverfassungsgericht (2023), 'Beschluss des Zweiten Senats vom 15. November 2023 - 2 BvF 1/22'. BVerfG, Beschluss vom 15. November 2023 - 2 BvF 1/22. ECLI:DE:BVerfG:2023:fs20231115.2bvf000122. Veröffentlicht u.a. in: NJW 2024, 501.
- Chapman, C. & Stolee, K. T. (2016), Exploring regular expression usage and context in python, in 'Proceedings of the 25th International Symposium on Software Testing and Analysis', ISSTA 2016, Association for Computing Machinery, pp. 282–293.
- Chen, T. & Guestrin, C. (2016), Xgboost: A scalable tree boosting system, in 'Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining', pp. 785–794.
- Connolly, K. (2023), 'Fear and anger rise in germany over handling of budget crisis'.
URL: <https://www.theguardian.com/world/2023/dec/01/germany-scholz-government-handling-of-budget-crisis>
- Debortoli, S., Müller, O., Junglas, I. & Vom Brocke, J. (2016), 'Text mining for information systems researchers: An annotated topic modeling tutorial', *Communications of the Association for Information Systems (CAIS)* **39**(1), 7.
- Deutscher Bundestag (2025), 'Deutscher Bundestag - Arbeit des Deutschen Bundestages'.
URL: <https://www.bundestag.de/services/faq/arbeit-244948>
- Devika, M. D., Sunitha, C. & Ganesh, A. (2016), 'Sentiment analysis: a comparative study on different approaches', *Procedia Computer Science* **87**, 44–49.
- Dritsa, K., Thoma, A., Pavlopoulos, I. & Louridas, P. (2022), 'A greek parliament proceedings dataset for computational linguistics and political analysis', *Advances in Neural Information Processing Systems* **35**, 28874–28888.

- Eickhoff, M. & Neuss, N. (2017), Topic modelling methodology: Its use in information systems and other managerial disciplines, in 'Proceedings of the 25th European Conference on Information Systems (ECIS)', Research Papers, pp. 1327–1347.
- Emmer, M. (2016), 'Soziale medien in der politischen kommunikation', *Handbuch Soziale Medien* pp. 81–99.
- European Commission (2024), 'Standard eurobarometer 101 – spring 2024: Public opinion in the european union'.
- URL:** <https://europa.eu/eurobarometer/surveys/detail/3216>
- Fawcett, T. (2006), 'An introduction to roc analysis', *Pattern recognition letters* **27**(8), 861–874.
- Fiva, J. H., Nedregård, O. & Øien, H. (2025), 'The norwegian parliamentary debates dataset', *Scientific Data* **12**(1), 4.
- Güngör, O., Tiftikçi, M. & Sönmez, Ç. (2018), A corpus of grand national assembly of turkish parliament's transcripts, in 'Proceedings of the Eleventh International Conference on Language Resources and Evaluation'.
- Hamann, J. (2016), *"Das Volk" und "Die da oben": Misstrauen und Entfremdung in Europa*, Vol. 17 of *DGAP kompakt*, Forschungsinstitut der Deutschen Gesellschaft für Auswärtige Politik e.V., Berlin.
- Hofmann, S., Sommermann, C., Kraus, M., Zschech, P. & Rosenberger, J. (2024), Hate Speech and Sentiment of YouTube Video Comments From Public and Private Sources Covering the Israel-Palestine Conflict, in 'Wirtschaftsinformatik 2024 Proceedings'.
- Huang, J. & Ling, C. (2005), 'Using AUC and accuracy in evaluating learning algorithms', *IEEE Transactions on Knowledge and Data Engineering* **17**(3), 299–310.
- Jakubíček, M. & Kovár, V. (2010), Czechparl: Corpus of stenographic protocols from czech parliament., in 'RASLAN', pp. 41–46.
- Kim, E. H.-J., Jeong, Y. K., Kim, Y., Kang, K. Y. & Song, M. (2016), 'Topic-based content and sentiment analysis of ebola virus on twitter and in the news', *Journal of Information Science* **42**(6), 763–781.
- Koltsova, O. & Koltcov, S. (2013), 'Mapping the public agenda with topic modeling: The case of the russian livejournal', *Policy & Internet* **5**(2), 207–227.
- Körber-Stiftung (2024), 'Demokratie in der Krise Umfrage 2024'.
- URL:** <https://koerber-stiftung.de/projekte/demokratie-in-der-krise-umfrage-2024/>
- Körber-Stiftung (2023), 'Deutsche verlieren vertrauen in ihre demokratie'.
- URL:** <https://koerber-stiftung.de/presse/mitteilungen/deutsche-verlieren-vertrauen-in-ihre-demokratie/>
- Laver, M., Benoit, K. & Garry, J. (2003), 'Extracting policy positions from political texts using words as data', *American Political Science Review* **97**(2), 311–331.
- Lilleberg, J., Zhu, Y. & Zhang, Y. (2015), Support vector machines and word2vec for text classification with semantic features, in '2015 IEEE 14th International Conference on Cognitive Informatics & Cognitive Computing (ICCI*CC)', pp. 136–140.
- URL:** <https://ieeexplore.ieee.org/document/7259377/?arnumber=7259377>
- Lorenz-Spreen, P., Oswald, L., Lewandowsky, S. & Hertwig, R. (2023), 'A systematic review of worldwide causal and correlational evidence on digital media and democracy', *Nature human behaviour* **7**(1), 74–101.
- Lu, B., Ott, M., Cardie, C. & Tsou, B. K. (2011), Multi-aspect sentiment analysis with topic models, in '2011 IEEE 11th international conference on data mining workshops', IEEE, pp. 81–88.

- Mansikkaniemi, A., Smit, P. & Kurimo, M. (2017), Automatic construction of the finnish parliament speech corpus, in 'INTER_SPEECH', International Speech Communication Association (ISCA), pp. 3762–3766.
- Marx, M., Schuth, A. et al. (2010), 'Dutchparl. a corpus of parliamentary documents in dutch', *Proceedings Language Resources and Evaluation (LREC)* pp pp. 3670–3677.
- Meakem, A. (2024), 'German government collapses amid uncertainties over u.s. election'.
URL: <https://foreignpolicy.com/2024/11/06/germany-coalition-government-collapse-scholz-lindner-budget-spending-trump/>
- Nasution, A. H. & Onan, A. (2024), 'Chatgpt label: Comparing the quality of human-generated and llm-generated annotations in low-resource language nlp tasks', *IEEE Access* **12**, 71876–71900.
- Ogrodniczuk, M. (2018), Polish parliamentary corpus, in 'Proceedings of the LREC 2018 workshop ParlaCLARIN: creating and using parliamentary corpora', pp. 15–19.
- Ogunleye, B., Maswera, T., Hirsch, L., Gaudoin, J. & Brunsdon, T. (2023), 'Comparison of topic modelling approaches in the banking context', *Applied Sciences* **13**(2), 797.
- Rauh, C. (2018), 'Validating a sentiment dictionary for german political language—a workbench note', *Journal of Information Technology & Politics* **15**(4), 319–343.
- Rauh, C. & Schwalbach, J. (2020), 'The parlspeech v2 data set: Full-text corpora of 6.3 million parliamentary speeches in the key legislative chambers of nine representative democracies', *Harvard Dataverse* **1**(1).
- Richter, F., Koch, P., Franke, O., Kraus, J., Warode, L., Kuruc, F., Heine, S. & Schöps, K. (2023), 'Open discourse: towards the first fully comprehensive and annotated corpus of the parliamentary protocols of the german bundestag', *OSF Preprints* .
URL: <https://doi.org/10.31235/osf.io/dx87u>
- Rosenberger, J., Wolfrum, L., Weinzierl, S., Kraus, M. & Zschech, P. (2025), 'CareerBERT: Matching resumes to ESCO jobs in a shared embedding space for generic job recommendations', *Expert Systems with Applications* **275**, 127043.
- Sahoo, B. P., Ashri, D. & Gulati, A. (2024), 'Drivers of growth: Empirical analysis of ten largest economies', *The Indian Economic Journal* p. 00194662241260815.
- Shah, A. M., Yan, X., Qayyum, A., Naqvi, R. A. & Shah, S. J. (2021), 'Mining topic and sentiment dynamics in physician rating websites during the early wave of the covid-19 pandemic: Machine learning approach', *International Journal of Medical Informatics* **149**, 104434.
- Singh, A., Thakur, N. & Sharma, A. (2016), A review of supervised machine learning algorithms, in '2016 3rd international conference on computing for sustainable global development (INDIACom)', Ieee, pp. 1310–1315.
- Smelyakov, K., Karachevtsev, D., Kulemza, D., Samoilenko, Y., Patlan, O. & Chupryna, A. (2020), Effectiveness of preprocessing algorithms for natural language processing applications, in '2020 IEEE International Conference on Problems of Infocommunications. Science and Technology (PIC S&T)', pp. 187–191.
- Stine, R. A. (2019), 'Sentiment analysis', *Annual review of statistics and its application* **6**(1), 287–308.
- Storey, V. C. & Park, E. H. (2022), 'An ontology of emotion process to support sentiment analysis', *Journal of the Association for Information Systems* **23**(4), 999–1036.
- Tan, K. L., Lee, C. P. & Lim, K. M. (2023), 'A survey of sentiment analysis: Approaches, datasets, and future research', *Applied Sciences* **13**(7), 4550.