

Article

Measuring Nuclear Democracy: A Large-Scale Analysis of German Nuclear Energy Discourse

Maximilian Teich , Arne Cypionka , Thomas Haider  and Malte Rehbein 

Chair of Computational Humanities, University of Passau, 94032 Passau, Germany;

arne.cypionka@uni-passau.de (A.C.); thomas.haider@uni-passau.de (T.H.); malte.rehbein@uni-passau.de (M.R.)

* Correspondence: maximilian.teich@uni-passau.de

Abstract

Nuclear energy has been a highly controversial issue in post-war (West) German politics. Previous research has identified discourse on nuclear energy as an important factor in the development of German political culture. However, little work has been done on the quantitative analysis of this discourse, despite the availability of large amounts of text data. We use large language models to classify texts in a corpus of Bundestag proceedings and news articles according to speakers' stance and framing. In combination with the texts' metadata, this enables us to draw conclusions about the positions taken by political parties and media publications on nuclear energy over time. We find that while media reporting remained mostly neutral, Bundestag speakers became divided on the issue along party lines by the 1970s. The framing of nuclear energy in both media and politics shifted in response to events and policy needs. More generally, our approach could be applied to other problems in discourse analysis where large amounts of data are available but no high-quality annotations for classifier training exist.

Keywords: digital history; artificial intelligence; LLM; discourse analysis; framing analysis; stance analysis; nuclear energy; Germany

1. Introduction

In the spring of 2023, Germany decommissioned its last nuclear power plants. As energy prices had surged in the preceding months, the decision provoked political opposition, culminating in a parliamentary inquiry into the government's handling of the process. For outside observers, the episode marked a fitting conclusion to decades of controversy. After initial enthusiasm for nuclear technology faded in the 1970s, much of West German political discourse kept returning to nuclear energy: its benefits, its dangers, and its long-term consequences. Plans for a nuclear future were drawn up and rescinded; phase-out dates were set, postponed, brought forward, and extended again. Nuclear energy polarized German society and politics like few other issues: it allowed post-war Germans to debate fundamental questions at a societal scale for the first time, as the controversy engulfed and shaped protest movements, NGOs, churches, political parties, and courtrooms. Uekötter (2022) coined the term “*Atomare Demokratie*” (“nuclear democracy”) to describe the development of contemporary German political culture around the nuclear question.

Researchers have devoted much attention to the history of nuclear energy in Germany, focusing variously on protest movements and their effects on civil society (Opp and Roehl 1990; Uekötter 2022), on international dimensions (Kirchhof 2014; Meyer 2024;



Academic Editor: Doris Gruber

Received: 7 March 2026

Revised: 9 May 2026

Accepted: 13 May 2026

Published: 27 May 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Ogawa 2023), or on discourse analysis (Jordan 2018; Jung 2013; Kliment 1994). The Bundestag—Germany’s legislature responsible for setting and overseeing the country’s nuclear policies—has been studied less frequently, even though a wealth of parliamentary proceedings on nuclear issues is available. These sources have been used to qualitatively assess the discourse (Agyemang 2024) and to quantitatively examine aspects of parliamentary engagement, for instance by analyzing term frequencies (Roose 2010) and collocations (Amri-Henkel 2021).

Building on the existing scholarship on Germany’s nuclear history, we propose to study the surrounding historical discourse through text classifications. Traditional machine learning approaches require a considerable amount of domain specific training data to solve such tasks. To circumvent this requirement, we use large language models (LLMs) to predict text classes for a large corpus of available Bundestag proceedings and an equally large sample of news texts for comparison. In total, over 30,000 parliamentary and journalistic texts are examined in the process. By classifying the stances of speakers and authors on nuclear energy, we can quantitatively assess how different political parties and media outlets have positioned themselves on the issue over time. An additional classification of framing strategies used to contextualize nuclear energy illuminates how the issues raised by participants in the discourse have evolved.

By generating and analyzing automated annotations for this large corpus of historical texts, the development of German nuclear energy discourse can be traced in detail. Following Uekötter (2022)’s notion that nuclear energy served as a test case for German society to practice democratic negotiation, we place particular emphasis on aspects indicating the evolution of post-war political culture in Germany. In particular, we consider the following questions: First, how did stances on nuclear energy change, and what differences can be observed among members of various political parties and authors of different news outlets. An understanding of these changes enables us to scrutinize the assumption that German politics moved from a consensus on nuclear energy to intense controversy along party lines. Second, how did the various participants in the debates frame nuclear energy and how did their framing evolve. This question allows us to contextualize observed changes in stance and to gain insights into the roles of parties and media in negotiating nuclear energy and shaping its public perception. Third, we assess how changes in stance on and framing of nuclear energy in Germany were related to contemporary events, such as nuclear disasters, protests, and election campaigns.

2. Materials and Methods

2.1. Historical Material

This study is based first on the proceedings of the Bundestag, the main legislative body of the Federal Republic of Germany. Under the German “Grundgesetz” (basic law; constitution), the power to regulate the use of nuclear energy is vested in the federal legislature, making the Bundestag a central actor in shaping related policies. The first nuclear energy law (*Atomgesetz*) was passed in late 1959, providing a legal framework that allowed research institutions and energy providers to develop a nuclear infrastructure. In the following decades, this law was amended numerous times, both to account for advances in the technical understanding of nuclear energy and to reflect shifting preferences among lawmakers and the public. The most recent amendments were adopted in November 2022, extending the operating permits of Germany’s last commercial reactors—originally scheduled for decommissioning at the end of that year—by three months in order to mitigate energy shortages. Given the scope of federal legislation on the matter and the Bundestag’s role as a central venue for political debate in the post-war German body politic, its parliamentary proceedings constitute a valuable source for the study of German nuclear

policy. They can also be expected to shed light on broader societal perspectives on nuclear energy over time.

The Bundestag has published its proceedings since its establishment in 1949. More recently, the parliamentary administration has also made the proceedings available online as XML files. Unfortunately, these files are, for the most part, poorly structured and annotated. It appears that the results of OCR processing of the originally printed proceedings were simply dumped into a single tag for each XML file, along with some metadata on the respective session. Only in the most recent years, starting in September 2023, has the Bundestag developed a more comprehensive schema, which now includes annotations for speakers, speeches, and the issues discussed. Still, the glaring lack of annotated proceedings for the years 1949 to 2023 has prompted several teams of German researchers to step into the breach and develop more accessible digital editions.

The DeuParl corpus collects proceedings of both the Bundestag and earlier German parliaments dating back to the 1860s. The texts are preprocessed to normalize words and exclude digressive parts such as name lists (Walter et al. 2021). The ParlSpeech corpus compiles proceedings from multiple European parliaments, including the Bundestag, and uses regular expressions to identify individual speeches, speakers, and their respective party affiliations (Rauh and Schwalbach 2020). In a similar manner, but focused specifically on the Bundestag, the GermaParl corpus segments proceedings into individual speeches, with speaker and party identification; the results are available as TEI files (Blaette et al. 2022; Blätte and Blessing 2018). The OpenDiscourse project provides annotations of comparable quality to GermaParl and makes them accessible through a convenient web interface designed to promote public access to the sources (Richter et al. 2023). Additionally, several smaller corpora containing partial Bundestag proceedings have been published. For this research on nuclear energy debates, we selected the GermaParl corpus because of the high quality of its annotations and its easily parsed TEI-based format. The latest version of GermaParl (v2.0.0) covers all proceedings up to and including the 19th parliamentary term, which ended in October 2021.

The GermaParl corpus can be further enriched by linking it to the master data (“*Stammdaten*”) provided by the Bundestag alongside their XML files. These contain additional information on individual members of parliament, such as dates of birth, party affiliations, and details on the districts they directly represented or, in the case of indirect elections, the party lists from which they ran. Incorporating this information enables a more nuanced understanding of who the individual speakers were and how their speeches should be interpreted.

To analyze how the discourse developed among the public outside the Bundestag, we added newspaper articles as a second basis for our study. This allows for comparisons between the political institution of the Bundestag and the subject’s treatment in the media. *Der Spiegel* and *Die Zeit* were chosen as the publications to investigate. Both have been circulated nationally since the mid-1940s and are regarded as papers of record in Germany. In addition, both are published weekly, which we assume allows for more reflective articles and broader perspectives. Rather than focusing on daily news reporting, weeklies can offer their readers summaries of events over longer periods, contextual information, and more nuanced opinions on controversial political issues such as nuclear policy.

Der Spiegel and *Die Zeit* are available as part of the *German Reference Corpus* (Deutsches Referenzkorpus), a collection of contemporary German texts provided by the *Leibniz Institute for the German Language (IDS)* and accessible via their COSMAS II interface (Kupietz et al. 2018). The corpus includes all texts published in *Der Spiegel* from its inception in 1947 through 2024, as well as all texts from *Die Zeit* from 1953 to 2024.

2.2. Preprocessing

Before examining stances toward nuclear energy and the frames used to contextualize the issue, the material is first preprocessed into manageable chunks. Rather than processing all available proceedings and news texts, only the parts directly relevant to nuclear energy were considered. To do this, we searched for keywords in the material and extracted each sentence containing the keyword along with a context window of two preceding and two following sentences. We found that such chunks are relatively quick to read and comprehend, while still providing enough context to capture the nuances of expressed opinions.

As keywords, the terms commonly found in German to denote nuclear energy were used: “Atomenergie” (atomic energy), “Atomkraft” (atomic power), “Kernenergie” (nuclear energy), and “Kernkraft” (nuclear power). These are mostly used synonymously, although the prefix “Atom-” tends to carry more negative connotations than “Kern-”. This preprocessing yields 13,021 text chunks from the Bundestag’s proceedings, 8347 text chunks from *Der Spiegel*, and 10,535 text chunks from *Die Zeit*. These result in a corpus of 31,903 text chunks.

2.3. Stance Classification

The first issue we address is the position that speakers and authors of the texts in the corpus take on nuclear energy. In linguistics, such questions are often framed as problems of stance detection or stance classification. Du Bois (2007) defined stance “as a public act by a social actor, achieved dialogically through overt communicative means (language, gesture, and other symbolic forms), through which social actors simultaneously evaluate objects, position subjects (themselves and others), and align with other subjects, with respect to any salient dimension of value in the sociocultural field.” In other words, stance indicates the position an actor takes on a specific target in their communication. In our case, we examine how authors and speakers publicly position themselves through their texts in the political discourse on nuclear energy. They presumably do so by expressing approval or disapproval of the technology and its application, as well as by indicating how their opinion relates to that of other actors or groups of actors who participate in the discourse. In the case of newspaper articles, one may also encounter a large proportion of neutral stances, in which an author does not reveal their opinion on the technology but rather describes aspects of the discourse or surrounding events without showing personal judgment.

Scholarship on natural language processing (NLP) has worked on detecting or classifying stances automatically for years, and systematic surveys of the methods employed by researchers are available (Alturayef et al. 2023; Küçük and Can 2020). Early research primarily approached stance classification as a machine learning problem, treating it alongside similar NLP tasks. Researchers trained supervised machine learning algorithms on relatively large datasets of manually annotated stances to produce reasonable predictions. Such datasets most frequently originated from social media, with some built specifically for political content (Sobhani et al. 2017; Taulé et al. 2017) and others derived from parliamentary proceedings (Balahur et al. 2009; Thomas et al. 2006). To train the algorithms, several features were considered, including n-grams, TF-IDF weights, and word embeddings (Hasan and Ng 2013; Küçük and Can 2020).

Recent scholarship has examined the use of LLMs to automatically classify stance through prompting. The primary advantage of this approach is that no training data are required. Tasks for which no specific datasets exist can thus be processed without extensive prior annotation by experts. Cruickshank and Ng (2024) found that LLM-based approaches match conventional machine learning classifiers on many sample tasks, and that fine-tuning the LLMs does little to improve overall performance. Upadhyaya et al. (2025) conclude that

stance classification with prompted LLMs outperforms supervised learning models and can be further enhanced when embedded in a multi-stage framework. Scholarship demonstrating the effectiveness of LLMs in sentiment analysis of historical texts (Cherukuri et al. 2025) suggests that the approach can work well on older documents.

Following this recent work, we used a zero-shot approach based on prompted LLMs to classify stances of the texts in the corpus. The principal reason for this is that no task-specific training data are available for the implementation of a supervised approach. Generating such training data would require considerable human effort. Although zero-shot methods eliminate the need for task-specific training data, a manually annotated sample is still required for evaluation. To this end, a random sample of 200 preprocessed text chunks was drawn, comprising 100 chunks of Bundestag proceedings and 50 each from *Die Zeit* and *Der Spiegel* articles.

2.3.1. Manual Annotation of Stance

In order to derive reliable gold standard data against which to test an automated approach for labeling texts, a manual annotation of this sample of 200 text chunks was performed. Three human annotators, all with academic backgrounds in history, were involved. First, the annotators discussed and labeled example texts to develop a shared understanding of the task and drew up annotation guidelines to handle edge cases. Second, each annotator independently annotated the 200 sampled texts. Annotators judged the speaker's or author's stance on nuclear energy and assigned a corresponding label. Initially, labeling stance on a five-point Likert scale (positive, slightly positive, neutral, slightly negative, negative) was proposed, but we found that the annotators could not establish consistent criteria to reliably distinguish extreme stances from somewhat negative or positive ones. Consequently, a three-point scale (negative, neutral, positive) was employed. Neutral stances were assumed in cases where the speaker's or author's stance could not be determined with certainty, as well as in cases where the speaker or author only reproduced the words of others. As a result, the majority (82%) of the journalistic texts were labeled neutral by all annotators, as these are mostly reports or interviews in which the author withholds their opinion. In contrast, speeches from the Bundestag were more likely to be labeled as conveying a positive or negative stance. Table 1 shows examples of the sampled texts and how they were handled by the annotators. In 148 cases all annotators agreed, in 47 only two agreed, and in five cases all three annotators reached different conclusions, making a discussion necessary to determine the best label.

Table 1. Examples for manually annotated texts with full, partial, or no agreement by the annotators.

Text Chunk	Translation	Annotations	Interpretation
Mit dem Ausstieg aus der Kernenergie würde ein gigantischer Arbeitsplatzverlust einhergehen und auch eine Zunahme des CO ₂ -Ausstoßes wäre zu befürchten. Es ist ein Märchen, zu glauben, dass Kernkraftwerke durch Energieeinsparung oder erneuerbare Energien kompensiert werden können. [...]	With the phase-out of nuclear energy, a gigantic loss of jobs would occur, and an increase in CO ₂ emissions would also have to be feared. It is a myth to believe that nuclear power plants can be compensated for through energy savings or renewable energies. [...]	Positive/ Positive/ Positive	All annotators agreed on the speaker's positive stance, as indicated by his emphasis on the importance of nuclear plants for the energy supply, the job market, and the environment.

Table 1. *Cont.*

Text Chunk	Translation	Annotations	Interpretation
Gewalt gegen Sachen und Menschen ist ein Irrweg. Liebe Kolleginnen und Kollegen, gewalttätige Aktionen schaden den Gegnern der Kernenergie. [...]	Violence against property and people is a misguided path. Dear colleagues, violent actions harm the opponents of nuclear energy. [...]	Neutral/ Negative/ Negative	Annotators 2 and 3 judged that the speaker's stance was negative because of his advice to protesters. Annotator 1 felt that the speaker's own stance on nuclear energy could not be determined with certainty.
Welche Antwort haben Sie darauf, wie der Ausstieg aus der Atomenergie ohne den nötigen Netzausbau gelingen kann? Worum geht es hier eigentlich? Wir haben uns – mit breiter Zustimmung in der Bevölkerung – dazu entschieden, als erstes Industrieland aus der Atomenergie auszusteigen [...]	What answer do you have to the question of how the phase-out of nuclear energy is supposed to succeed without the necessary expansion of the power grid? What is this really about? We decided—with broad support among the population—to be the first industrialized country to phase out nuclear energy. [...]	Neutral/ Negative/ Positive	The questions raised by the speaker make it difficult to discern his stance on nuclear energy. Annotator 2 suspected that he was opposed to nuclear energy, while Annotator 3 felt that he criticized the mishandled phase-out from a positive stance. Given these uncertainties, a neutral label was ultimately assigned.

The labels assigned by the three annotators were analyzed to see how reliably the stance classification task can be solved manually. To allow for pairwise comparison between the individual annotators, Cohen's kappa coefficients were used (Cohen 1960). This has the advantage of estimating inter-annotator agreement beyond the expected agreement by chance. We found that on average, the agreement between annotators was 0.62 for the Bundestag proceedings and 0.66 for the newspaper articles, as shown in Table 2. According to the widely used benchmark of Landis and Koch (1977), Cohen's kappa scores between 0.6 and 0.8 are considered substantial, while scores above 0.8 indicate almost perfect agreement. This indicates that the parliamentary proceedings could be annotated with substantial agreement by the annotators, with slightly higher agreement on newspaper articles. When interpreting the annotation of the journalistic texts, however, one has to consider that Cohen's kappa often yields comparably low scores when a single label dominates the distribution, as the expected chance agreement for a predominant label is high (Feinstein and Cicchetti 1990; Gwet 2014). As this is likely the case with mostly neutral newspaper articles, percentage agreement was also calculated as a check on the Cohen's kappa scores. As annotators on average agree most times on the assigned labels for the proceedings (74%) and the news texts (89%), we conclude that both types of texts could be annotated with reasonable reliability.

The manual annotations were subsequently processed into a gold dataset, giving one label for each of the manually annotated texts. The label assumed to be correct was determined using a majority rule. When at least two of the three annotators agreed on a label, that label was used in the gold standard. In cases where no two annotators assigned the same label, the correct label was decided through discussion.

Table 2. Inter-annotator agreement for the stance annotation task described in Section 2.3.1 and the framing annotation task described in Section 2.4.2. The values shown are averages of pairwise Cohen’s kappa coefficients. Label frequency indicates the average share of texts that annotators labeled as relevant (not false).

	Bundestag Proceedings	Newspaper Articles	Label Frequency
stance annotation	0.62	0.66	1.00
alternatives	0.76	0.70	0.26
Atomausstieg	0.67	0.65	0.36
costs	0.45	0.77	0.11
economic effects	0.59	0.53	0.21
environment	0.96	0.73	0.17
innovation	0.50	0.36	0.14
international relations	0.50	0.60	0.12
protest	0.75	0.78	0.10
reactor safety	0.82	0.76	0.28
reliability	0.34	0.42	0.07
uranium	0.51	0.44	0.04
waste	0.65	0.48	0.14
weapons	0.87	0.79	0.06

2.3.2. Automatic Classification of Stance

To evaluate how well LLMs can match human annotators’ judgments, the same texts we annotated manually were subsequently embedded in prompts describing the stance detection task. Three LLMs (Gemma2 27B, GPT-5, and Llama3.1 70B) were instructed through these prompts to label all the texts. To ensure valid results, the models’ responses were automatically checked in a validation loop. If no valid answer was found in a response, the prompting was repeated multiple times. Across the 200 test cases, this mechanism reliably ensured validity.

The LLMs’ results were compared to those of the annotators by scoring them against the presumably correct labels of the gold dataset. By measuring recall and precision, the LLMs’ performance can be judged. The resulting scores are shown in Table 3 for the Bundestag’s proceedings and Table 4 for the newspaper articles.

Table 3. Recall and precision scores of tested LLMs on stance (see Section 2.3.2) and framing (see Section 2.4.3) tasks of Bundestag proceedings measured against the gold data. The n value indicates the number of relevant text chunks (not labeled as false) in the gold data.

	Gemma2 27B		GPT-5		Llama3.1 70B		n
	Precision	Recall	Precision	Recall	Precision	Recall	
stance annotation	0.61	0.60	0.83	0.81	0.74	0.70	100
alternatives	0.84	0.87	0.94	0.95	0.81	0.74	35
Atomausstieg	0.85	0.86	0.93	0.93	0.79	0.77	42
costs	0.78	0.86	0.74	0.95	0.73	0.80	9
economic effects	0.80	0.80	0.81	0.85	0.74	0.58	22
environment	0.72	0.80	0.79	0.88	0.78	0.78	24
innovation	0.74	0.64	0.79	0.85	0.60	0.53	13
international relations	0.75	0.90	0.73	0.89	0.80	0.82	13
protest	0.55	0.91	1.00	1.00	0.75	0.99	2
reactor safety	0.79	0.88	0.84	0.94	0.72	0.72	20
reliability	0.54	0.61	0.67	0.72	0.53	0.60	6
uranium	1.00	1.00	0.69	0.82	0.99	0.83	3
waste	0.83	0.93	0.93	0.93	0.86	0.74	16
weapons	0.88	0.99	0.93	0.99	0.99	0.92	6

Table 4. Recall and precision scores of tested LLMs on stance (see Section 2.3.2) and framing (see Section 2.4.3) tasks of news texts measured against the gold data. The n value indicates the number of relevant text chunks (not labeled as false) in the gold data.

	Gemma2 27B		GPT-5		Llama3.1 70B		n
	Precision	Recall	Precision	Recall	Precision	Recall	
stance annotation	0.52	0.74	0.59	0.89	0.54	0.65	200
alternatives	0.75	0.84	0.87	0.94	0.75	0.78	17
Atomausstieg	0.70	0.74	0.87	0.86	0.81	0.70	29
costs	0.77	0.83	0.81	0.88	0.80	0.80	11
economic effects	0.77	0.87	0.75	0.80	0.77	0.65	18
environment	0.56	0.66	0.66	0.88	0.63	0.66	10
innovation	0.88	0.72	0.77	0.83	0.79	0.59	11
international relations	0.69	0.88	0.70	0.92	0.73	0.75	9
protest	0.67	0.82	0.83	0.86	0.68	0.70	14
reactor safety	0.87	0.89	0.88	0.91	0.80	0.76	33
reliability	0.54	0.66	0.60	0.74	0.49	0.46	11
uranium	0.74	0.66	0.61	0.65	0.99	0.67	3
waste	0.75	0.87	0.92	0.95	0.96	0.69	13
weapons	0.81	0.90	0.88	0.99	0.86	0.74	6

When labeling text chunks from the Bundestag’s proceedings, the annotators achieved a substantial agreement. Their high consistency could not be fully reproduced by the LLMs. Nevertheless, the best-performing model, GPT-5, reached an average precision of 83% and an average recall of 81%, with Gemma2 and Llama 3.1 scoring a little lower. This suggests that a well-instructed model can provide reasonably good approximations of the labels assigned on a three-point scale by human raters on the task.

For the sampled newspaper articles, the observed Cohen’s kappa agreement between annotators was comparable to that on the proceedings. However, annotator labels on average coincided for 89% of the sampled texts, indicating that the majority of neutral texts were labeled consistently by the three annotators. The high disagreement likely results from differing labels on the relatively few non-neutral texts. Again, GPT-5 performed fairly well on this task, achieving a precision of 59% and a recall of 89% respectively, with Gemma2 and Llama3.1 scoring lower.

2.4. Framing Analysis

The second issue we address is framing of nuclear energy by speakers and authors. The concept of framing is commonly attributed to Goffman (1974) who drew on earlier work by Bateson (1955, 1972) to describe how individuals organize experience and define social situations. In its broadest sense, framing refers to the ‘packaging’ of meanings by structuring concepts and events into a coherent ‘schema of interpretation’ that facilitates understanding (Goffman 1974). Such packaging operates at multiple levels, including semantic, cognitive, and communicative dimensions (Sullivan 2023). As the concept proliferated across disciplines, Entman (1993) sought to provide a more unified definition tailored to communication research. Entman described frames as means by which speakers “select some aspects of a perceived reality and make them more salient in a communicating text in such a way as to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation.” As such, framing functions on multiple levels: the speaker’s wording, the act of communication itself, and the perception of the audience (Otmakhova et al. 2024).

Work in Natural Language Processing on media framing spans a wide set of approaches and conceptualizations, reflecting the interdisciplinary origins of the concept. A growing body of applied work addresses framing in broader settings of media bias and online discourse, often intersecting with stance, persuasion techniques, or ideology

(Hamborg et al. 2019; Hofmann et al. 2022; Piskorski et al. 2023). On an operational level, three main types of framing may be distinguished (Otmakhova et al. 2024):

First, the largest group of studies focuses on *emphasis framing*, also referred to as *issue framing*, by modeling frames as the selection of relevant issues that are foregrounded in a text. These approaches typically operationalize framing as multi-label or topic classification, with a focus on providing a level of abstraction in the labelset based on the Policy Frames Codebook by Boydston et al. (2014) that can be broadly applied to political discourse (Card et al. 2015; Klamm et al. 2022; Sagi et al. 2013) or adapted to specific issues and domains (Card et al. 2022; Hartmann et al. 2019; Jurkschat et al. 2022). Issue frames have been captured through topic models in newspaper articles (Burscher et al. 2014; DiMaggio et al. 2013) and parliamentary records (Klamm et al. 2022; Nguyen et al. 2015), or through zero-shot classification (Mohammed Afzal and Nakov 2023; Reiter-Haas et al. 2023).

A second line examines *framing as lexical choice and labeling*, capturing how entities are depicted in discourse. One strand investigates metaphorical framing, including dehumanizing metaphors (Mendelsohn and Budak 2025; Mendelsohn et al. 2020), other studies examine naming practices (van den Berg et al. 2019), modifiers or adjectives employed in framing (Jing and Ahn 2021; Kwak et al. 2021; Luo et al. 2024), or models connotation frames, to detect asymmetries in agency and power through verb semantics and predicate–argument structure (Antoniak et al. 2023; Field and Tsvetkov 2019; Rashkin et al. 2016; Sap et al. 2017). The third line examines *narrative framing*, modeling frames as higher-level story structures that organize events, characters, and responsibility, for example through assigning narrative roles as to who is the hero, villain, or victim in a particular discourse (Otmakhova and Frermann 2025; Stambach et al. 2022; Steinbach et al. 2025).

In this study, we follow the emphasis framing perspective and approach framing as *issue framing*. This choice is motivated by the characteristics of the German nuclear-energy discourse and by the aims of our analysis: First, the controversy surrounding nuclear energy is multidimensional: speakers and authors routinely move between emphasizing economic viability, reactor safety, waste management, climate and environmental impacts, geopolitical dependencies, and governance questions. Second, our research interest is not in evaluative language at the level of wording, but which *policy issues* are invoked to define the problem and justify positions. This approach also allows us to adapt a labelset that is specific to the discourse domain we are investigating, and then provides a bridge between computational detection and the policy-relevant content of the debate, enabling us to trace how the discourse foregrounds particular dimensions across actors and over time.

2.4.1. Conceptualization of Framing Categories

Accordingly, we examine the framing of nuclear energy in our text corpus as a multi-label topic classification problem, which can also be addressed using an LLM-based approach. Again, such an approach has the advantage of not requiring task-specific training data which would be difficult to generate. By analyzing the topic referenced in each text chunk, we can infer how participants frame the issue in the discourse. To avoid relying on an overly abstract label set for political policy issues, the domain-specific set introduced by Jurkschat et al. (2022) was adapted. In contrast to generic emphasis-framing inventories such as the Policy Frames Codebook (Boydston et al. 2014), which defines cross-issue dimensions such as *Economic*, *Morality*, or *Health & Safety*, Jurkschat et al. (2022) develop a set of categories tailored explicitly to recurring argumentative aspects in the British nuclear energy debate. Their Argument Aspect Corpus–Nuclear Energy (AAC-NE) encodes the following issues:

- Alternatives: Sources of energy other than nuclear fission. In earlier texts mostly coal, later renewables.

- Costs: Costs of building and maintaining nuclear infrastructure and prices for nuclear energy.
- Environment: Effects of nuclear energy on the environment, for example through radiation or the avoidance of CO₂ emissions.
- Innovation: Research and innovation in the realm of nuclear technology.
- Reactor safety: Safety of nuclear reactors for humans and the environment.
- Reliability: Reliability of nuclear fission as an energy source.
- Waste: Handling and depositing of nuclear waste.
- Weapons: Military applications of nuclear technology.

This label set was originally intended to describe issue frames related to nuclear energy in British news texts. We expanded this set by including additional issues which the annotators found to be pertinent to the German debates when initially discussing the texts, such as the German term *Atomausstieg* which relates to the policy trajectory of the phase-out, or more generic frames such as economy or uranium (resource) which are also found in the Policy Frames Codebook. The issue frames added after annotator discussion to capture issues frequently addressed in German texts are:

- *Atomausstieg* (nuclear phase-out): The phase-out of nuclear energy, both in Germany or any other jurisdiction.
- Economic effects: Effects of nuclear energy on the economy, especially on economic growth and unemployment.
- International relations: Diplomacy, international organizations, and international cooperation on questions of nuclear technology.
- Protest: Public protest against nuclear technology or its application.
- Uranium: Mining, trading, and processing of uranium.

2.4.2. Manual Framing Annotation

The three annotators who labeled the sampled 200 texts for stance also annotated the same texts for framing issues using a multi-label approach. If an issue was found to be mentioned in a text, the annotators assigned the label “True”; otherwise, “False” was used. As with the stance annotation, the annotators discussed sample cases to develop shared criteria for determining when an issue is invoked before independently annotating the texts. The gold standard dataset was derived from the manual annotations as before. Because only two values (“True” or “False”) were used for labeling, the majority principle always yielded a definitive solution.

When considering the agreement among human annotators shown in Table 2, vast differences between issue frames become apparent. Expressed in Cohen’s kappa coefficient values, annotators achieve almost perfect (>0.8) or substantial (>0.6) agreement on alternatives, *Atomausstieg* (nuclear phase-out), environment, protest, reactor safety, and weapons. The remaining issues are labeled less consistently. This is unsurprising, as these issues were more difficult for annotators to define during the initial discussion. For example, annotators had difficulty discerning the cost issue. Even in cases where texts explicitly stated amounts of money drawn from the federal budget, it was often a matter of interpretation whether this money was directly dedicated to nuclear infrastructure. In other cases, the labeling heavily depended on the specific annotators’ contextual knowledge. For instance, Bundestag speakers often merely mentioned the town name “Kalkar” in reference to the fast breeder research reactor located there, which may or may not count as an allusion to the innovation issue, depending on the annotator’s awareness of the site.

2.4.3. Automatic Annotation of Issue Frames

Following the analysis of human annotator agreement, three LLMs were prompted with instructions based on the annotators' discussions and the text chunks to replicate the task. Table 3 shows the recall and precision of the LLMs' output for the issue frame labeling of the Bundestag proceedings, while Table 4 presents the corresponding scores for the sampled newspaper articles.

The three tested LLMs vary in performance when evaluated against the gold dataset. GPT-5 demonstrates the strongest overall results, with precision scores often exceeding 80% and recall scores frequently above 90%. Gemma2 and Llama3.1 are less consistent but still produce strong predictions for some labels, albeit often with a low number of true cases. The issues themselves also show notable differences, reflecting the difficulty faced by the human annotators in assigning them. Whereas issues like *Atomausstieg* and reactor safety are predicted with consistently high scores across both parliamentary and journalistic texts, identifying issues such as reliability is more challenging.

2.5. Large Scale Prediction of Issue Frames

In a final step, all 31,903 preprocessed text chunks of the corpus were annotated automatically using an LLM. For each task, we selected the model that performed best in evaluations for the respective text type and label. As shown in the evaluation tables, this was always GPT-5. For stance annotation, each text chunk was labeled as negative, neutral, or positive. For the framing task, only those issues for which the human annotators showed substantial agreement were included (alternatives, *Atomausstieg*, environment, protest, reactor safety, waste, and weapons).

3. Results

After annotating the full corpus, we assess the resulting data. Before examining the annotations themselves, the distribution of text chunks by year is worth considering, as it serves as a good proxy for both legislative and journalistic interest in nuclear energy over time. Figure 1 shows that the respective distributions for the Bundestag's proceedings and the newspaper articles are quite similar, although the absolute number of the latter is much greater. The high correlation between the two distributions (Spearman correlation $r_s = 0.78$) suggests that parliamentary and journalistic texts followed similar trends in dealing with nuclear energy. However, the variation displayed by the newspaper article distribution is higher than that of the Bundestag's texts. This indicates that media interest concentrates more strongly around certain events, whereas the legislative discourse is more drawn out.

The number of both parliamentary and journalistic texts remains very limited until the 1970s, with somewhat elevated counts in the mid-1950s, when the prospects of nuclear energy for Germany's future energy supply were evaluated and several research reactors were planned and commissioned, the first in 1957 in Garching. In the same year, the first noticeable anti-nuclear protests in Germany took place. These were encouraged by a public warning from prominent nuclear scientists and supported by trade unions and the SPD, which opposed any nuclear armament proposed by the governing CDU/CSU. The ensuing campaign "Kampf dem Atomtod" (roughly translating to "fight against nuclear death") reached peak participation in 1958 (Lorenz 2011). By the time the nuclear energy law was passed by the Bundestag in 1959, most of this limited interest—as indicated by the number of texts—had already subsided. Minor amendments to the law in the 1960s also appear to have drawn little attention.




Figure 1. Distribution of text chunks in the corpus by year. Noticeable events are indicated by dashed vertical lines.

Only from the 1970s onward does a substantial number of texts appear each year. This development begins roughly with the occupation of the planned Wyhl nuclear site by protesters in 1975, the first such protest to receive national attention. Elevated text counts also align with the particularly violent protests against the construction of a nuclear power plant in Grohnde in 1977 and the large rally against a planned nuclear waste repository in Gorleben in 1979. By the early 1980s, however, overall interest seems to have declined already. This is somewhat surprising given that the mass protest at Brokdorf in 1981—held despite a county ban that was later ruled to be unconstitutional—established an important legal precedent for the right to protest in Germany (Lepsius 2025). Notably, when the Green Party first entered the Bundestag in 1983, the number of texts on nuclear energy was lower than during much of the preceding decade.

The highest number of texts, however, occurs slightly later, in 1986, the year of the Chernobyl disaster. Similarly, 2011, the year of the Fukushima accident, shows noticeable peaks in both distributions. This speaks to the large interest German society took in both events. A marked discrepancy between parliamentary and journalistic texts appears around the year 2000, shortly before the nuclear energy law was amended to mandate the phase-out of German nuclear power plants. The negotiations between the government, lawmakers, and energy providers at the time are reflected in the Bundestag proceedings but received comparatively little coverage in newspapers.

3.1. Stances

Figure 2 visualizes the LLM-annotated stance labels by showing the mean stance labels of newspapers and parties within overlapping ten-year windows. This provides an overview of changing attitudes toward nuclear energy across political divides over time. The articles of both *Der Spiegel* and *Die Zeit* average closely to the neutral value throughout the entire sample period. A total of 76.4% of news texts are annotated as neutral by the LLM. In the 1970s, the average for both publications shifts from a slight positive leaning to a slight negative leaning. These observations indicate that both papers maintained objectivity in their reporting on nuclear energy, with commentary playing a secondary role and remaining relatively balanced (16.4% of news texts are labeled as negative, 7.2% as positive). The shift of *Die Zeit* from a political center-right to a center-left orientation beginning in the mid-1950s (Staas 2021) may account for some of the differences between the publications, as its averages still lean more positively than those of *Der Spiegel* in the 1970s, before dropping to a more negative mean.




Figure 2. Mean stance labels for newspapers and parties in overlapping ten-year windows. Negative labels are shown as -1 , neutral labels as 0 , positive labels as 1 .

When examining speakers from different political parties in the Bundestag, a greater variety in average stances can be observed. From the 1950s to the 1970s, slightly positive averages are found for speakers of all major parties. By the late 1970s and early 1980s, the positions of representatives from the center-left SPD begin to trend negative, while the center-right Union of CDU and CSU, as well as the liberal FDP, remain relatively steady. In the 1980s, the Green Party emerges as a political competitor with a noticeably more negative average stance on nuclear energy, which—after 1990—is matched by the left-wing PDS/Linke. In the 2010s, the center-right parties approach the negative stances of the center-left parties, while the new right-wing AfD maintains markedly positive average stances on nuclear energy.

3.2. Frames

Consideration of the automatically annotated issue frames adds further nuance to the analysis. Figure 3 shows the proportion of texts classified as belonging to one of the annotated issues for which substantial agreement could be achieved by the annotators. It becomes apparent that the framing of nuclear energy in German political debates not only changed over time, but also exhibits differences between media coverage and legislative discussions. Whereas the analyzed newspapers show very similar trends for most framing issues, with few exceptions, the use of these issues varies noticeably among members of the Bundestag's different political parties.




Figure 3. Proportions of text chunks belonging to each annotated issue for newspapers and parties in overlapping ten-year windows.

Since the 1940s, alternatives to nuclear energy have been addressed in a relatively constant proportion of news texts, but appear with increasing frequency in parliamentary speeches. The small differences between political parties in their use of this framing issue may indicate a shared understanding of the need to balance nuclear power plants with other energy sources, if only for economic reasons. In the 1970s and 1980s, coal mining was still a relevant source of employment in West Germany, and parties agreed that it should not be forced out of the market by cheap nuclear energy. In the 21st century, renewable energy sources are increasingly debated as alternatives to phased-out nuclear plants.

The issue of *Atomausstieg* has received increasing attention in the media since its introduction into public debate in the 1970s. An analysis of the proportion of political parties' references to the issue reveals that the Green Party was its most active driver in parliament, pushing other center-left parties in the same direction by the 1980s. In contrast, center-right parties largely avoided the issue until the 21st century. In the most recently analyzed years, all parties show similarly high proportions, indicating that the issue has become central to discussions of nuclear energy across party lines.

Similar trends can be observed in the data on environmental and reactor safety issues. Both received increasing media attention over time—certainly driven by events such as the Chernobyl disaster, which caused a peak in the 1980s—and were emphasized by the Green Party after its first election to the Bundestag. However, in recent years, both issues have received less attention than *Atomausstieg*.

Differences between media and parliamentary discourse become very apparent in the case of the protest issue. While a high proportion of news texts has covered nuclear energy in relation to protests since the 1970s, this is not reflected in the parliamentary records. Noticeably, even the Green Party and the left-wing PDS/Linke did not focus strongly on the issue, although they gave it slightly more attention than the remaining parties. This may hint at the practical nature of lawmakers' debates, which tend to revolve more around technical and legal issues than daily news. Moreover, the protest movement lent greater political weight to the Greens' anti-nuclear policy, even without the need to explicitly introduce the protests into parliamentary debates.

A similar effect may be observed in the case of nuclear waste. Here, the Bundestag's proceedings show higher proportions than the media, again with center-left parties addressing the issue more frequently. As the parliament was actually charged with making decisions on the disposal of nuclear waste—undoubtedly a pressing problem that was never permanently resolved—it is unsurprising that it repeatedly returned to the issue, whereas the media may have covered only the most relevant stages of the process.

Notably, the issue of nuclear weapons received very little attention from both newspapers and lawmakers, despite having been an important aspect of early nuclear debates in Germany. The issue was addressed in particular by *Der Spiegel* in the 1950s, which was known for its critical coverage of German rearmament at the time. The slightly later references to the issue by SPD and FDP politicians may have been caused by emerging protests against nuclear armament and NATO nuclear sharing in the 1960s. The consistently low proportion of texts addressing nuclear weapons since is likely indicative of a strong discursive separation between military and civilian nuclear technology that developed at the time. Such a separation might also suggest that some discussions of nuclear weapons were not captured in the data, as debates on armament did not necessarily rely on the matched nuclear energy key words used in the preprocessing of the corpus.

3.3. Interpretation

For the most part, our observations on the Bundestag's proceedings and the sampled news texts align well with established knowledge of German nuclear energy discourse.

Through the 1960s, nuclear fission figured as the most spectacular technological achievement of the 20th century; the awe surrounding nuclear weapons was matched by hopes of equally great civilian benefits: Internationally, these expectations led to the creation of the International Atomic Energy Agency in 1957, which was intended to function as a body for global governance and cooperation in the field of nuclear technology despite Cold War tensions. In 1958, its European counterpart, the European Atomic Energy Community, was founded as a platform for independent European cooperation in the domain of nuclear technology. West Germany took an active part in these organizations, hoping to regain

recognition as a trustworthy partner for international cooperation and to achieve a leading role in emerging nuclear technology markets (Gantner 2023).

Whereas the Union-led governments and the SPD-led opposition disagreed on the question of nuclear armament, there was broad consensus on the technology's civilian potential. Gleitsmann (2011) has characterized this early cross-party enthusiasm as a dogma of indispensability: Without its own nuclear energy production, West Germany would not be able to re-attain prosperity in a nuclear post-war world. The annotations reflect this consensus in similar average stances of the three relevant political parties and slightly positive reporting in the sampled news media. At the same time, framing issues which might indicate opposition to this consensus, like *Atomausstieg* or protest, do hardly appear at all.

Only in the 1970s did the non-military use of nuclear energy spark widespread controversy in West Germany. As the politically envisioned nuclear infrastructure of research reactors, power plants, reprocessing facilities, and disposal sites was planned and built, local opposition formed. Initial challenges to nuclear projects through street protests and lawsuits grew into a networked movement of anti-nuclear activism. By the early 1980s, this movement had not only contributed to the introduction of post-war Germans to the notion of a civil society outside the sphere of governmental oversight, but had also become a veritable political force with the election of the Green party to the Bundestag on an anti-nuclear platform. Uekötter (2022) has thus described West Germany as an "*Atomare Demokratie*" (nuclear democracy), shaped by the experience of societal discourse on nuclear energy.

The breakdown of the early post-war consensus on nuclear energy and the accompanying diversification of the political culture can be traced in the annotations: The decisive political process between 1970 and 2000 is the realignment of the SPD, clearly visible in stance averages. Pressured by left-leaning protesters and their new Green Party competition to abandon the pro-nuclear consensus, the SPD began to adopt a more skeptical attitude. After the Chernobyl disaster, this skepticism quickly turned into opposition. In the framing annotations, the convergence of SPD and Green speakers in the frequency of addressing the issue of *Atomausstieg* is apparent. When a coalition government of SPD and Greens came to power in 1998, this consolidation of anti-nuclear sentiment among the center-left parties made possible a first parliamentary decision to phase out nuclear energy. The success of the Green Party in challenging the old consensus and introducing the notion of *Atomausstieg* into the debates supports Uekötter (2022)'s idea of a nuclear democracy, in which protest movements could reshape party politics. One can also see that whereas the shock of Chernobyl prompted the media to report primarily on the issue of reactor safety and public protests, parliamentary debates also revolved around the practical question of alternatives. These differences highlight the ability of the two sampled media outlets to determine their coverage on an increasingly controversial subject independently from political institutions, as would be expected in an open society.

After a new center-right government had eased the plans to phase out nuclear energy, the Fukushima accident in 2011 transformed attitudes on nuclear energy again. Around this time, the annotations show how even the Union parties' and the FDP's speakers averaged increasingly negative stances. Also, *Atomausstieg* and alternatives became dominant issues in the debates. Although the impact of the events in Japan was not as strong as that of the Chernobyl disaster, it likewise illustrates the influence of current events on political attitudes. The increasing focus on alternatives to nuclear energy after 2011 is likely related to the ongoing debates on the replacement of nuclear and fossil energy with renewable energy sources, broadly debated as the "*Energiewende*" (energy turn) in the German public.

4. Discussion

We have used LLMs to label a sample of nuclear energy–related text chunks from the Bundestag’s proceedings and the newspapers *Der Spiegel* and *Die Zeit*. The LLMs were tasked with assigning class labels for both speaker or author stance on nuclear energy and for one of several topics, which were used as a proxy for the framing of nuclear energy. To evaluate the quality of the models’ labeling, we measured performance against gold data derived from three human annotators’ assessments. We found that LLMs could predict stance and some issues with high reliability. However, some of the tested issues were too difficult to differentiate, even among the human annotators, and thus were not used in the automated labeling.

An analysis and interpretation of the machine-generated text annotations suggest that the Bundestag’s proceedings and the sampled news texts align well with the findings of previous research on German nuclear energy discourse. On the question of stances taken by relevant political parties on nuclear energy, the annotations illustrate long-term developments. They show how the parties went from alignment on the matter to a split in stance between center-right and center-left parties. Journalistic texts, on the other hand, retained a neutral tone throughout the analyzed time period. The frame annotation revealed in greater detail the contexts in which the discourse placed nuclear energy. While some frames faded out, others persisted or gained in prominence. In some cases, such as the issue of a nuclear phase-out, the annotation data reveal how individual parties introduced frames and reshaped the debates’ trajectory. Differences between media and legislative interests also become apparent, for example, the lack of references to protests in parliamentary proceedings despite strong media coverage, or the low frequency of news reports during the negotiations for a nuclear phase-out in the early 2000s. Lastly, on the question of the influence of events, the data show how trends in stance and framing for both the media and political institutions correspond to specific incidents: Disasters such as Chernobyl and Fukushima were driving concerns about reactor safety, and the election of the Green Party introduced the increasingly prominent issue of a nuclear phase-out.

The results provide relevant insights into the German nuclear energy discourse at scale. They show how the use of nuclear energy shifted from a consensus issue to a controversial topic. The positions of political parties on nuclear energy, as well as the political landscape itself, adapted over time in response to nuclear events and sustained popular movements. The diversification of German politics was accompanied by media outlets that maintained journalistic distance and managed to frame the issue independently of agendas set by political institutions. In many ways, then, the annotations can be read as supporting Uekötter (2022)’s idea of a nuclear democracy. Even if the discourse on nuclear energy did not by itself transform post-war German society and institutions, it nevertheless exemplifies the development toward a stable pluralistic democracy.

4.1. Limitations

The limitations of our work should be considered when parsing the results. Clearly, our annotations are limited by selection of source texts: The analyzed parliamentary proceedings only cover the years 1949 (when the Bundestag was constituted) through 2021, missing the potentially important years immediately following the war as well as the most recent iteration of the debate on nuclear plant decommissions. The selection of two newspapers covers a slightly longer time period, but cannot reflect the entirety of the German media landscape. Accordingly, our results may only cover a section of the post-war German debates.

Furthermore, an LLM-based approach requires rigorous evaluation. Neither are the pretraining data used to derive the employed models transparent, nor are the resulting

biases the models may exhibit. The quality of LLM outputs on the described tasks can only be judged in comparison to human assessments on the same material, which makes considerable effort necessary. It should also be acknowledged that the quantification of some 30,000 texts leads to a reduction in resolution, which may obscure aspects and nuances of the debates accessible to qualitative methods.

4.2. Future Work

In general, the approach of using LLMs to automatically assign class labels to texts holds great promise for quantitative discourse analysis. In scenarios where textual sources are abundant but high-quality annotations are insufficient to provide training data for conventional machine learning algorithms, LLMs can serve as a viable alternative, approaching the quality of human annotators in our test case. However, since use cases vary, it will always be necessary to establish some test data in order to properly evaluate the models' results against reliable human annotations. With regard to the corpora of Bundestag proceedings and news texts used in this paper, the demonstrated method could be adapted by modifying the label set and creating a new gold standard dataset through manual annotation to explore other topics in German public discourse.

Author Contributions: Conceptualization, M.T., T.H. and M.R.; methodology, M.T. and T.H.; software, M.T.; validation, M.T.; formal analysis, M.T.; investigation, M.T. and A.C.; resources, M.T. and A.C.; data curation, M.T. and A.C.; writing—original draft preparation, M.T.; writing—review and editing, M.T., A.C., T.H. and M.R.; visualization, M.T.; supervision, M.R.; project administration, M.T. and M.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in the study are openly available in Zenodo at <https://doi.org/10.5281/zenodo.20071984> (accessed on 12 May 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Agyemang, Jasmin. 2024. "Wer gegen Atomkraft ist, ist gegen Klimaschutz": Eine Argumentationsanalyse der Debatte zur Laufzeitverlängerung der Atomkraft im deutschen Bundestag. In *Exploring Eco-Discourses. Perspektiven - Ansätze - Methoden*. Edited by Johanna Freudenberg, Lisa Rhein and Niklas Simon. Darmstadt: Universitäts- und Landesbibliothek Darmstadt, pp. 8–36. [\[CrossRef\]](#)
- Alturayef, Nora, Hamzah Luqman, and Moataz Ahmed. 2023. A systematic review of machine learning techniques for stance detection and its applications. *Neural Computing and Applications* 35: 5113–44. [\[CrossRef\]](#)
- Amri-Henkel, Andrea. 2021. Die Energiewende im Bundestag: Ein politisches Transformationsprojekt?: Eine Diskursanalyse aus feministischer und sozial-ökologischer Perspektive. In *Die Energiewende im Bundestag: Ein politisches Transformationsprojekt?* Bielefeld: Transcript Verlag.
- Antoniak, Maria, Anjalie Field, Jimin Mun, Melanie Walsh, Lauren Klein, and Maarten Sap. 2023. Riveter: Measuring power and social dynamics between entities. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*. Toronto: Association for Computational Linguistics, pp. 377–88.
- Balahur, Alexandra, Zornitsa Kozareva, and Andrés Montoyo. 2009. Determining the Polarity and Source of Opinions Expressed in Political Debates. In *Computational Linguistics and Intelligent Text Processing*. Edited by Alexander Gelbukh. Berlin/Heidelberg: Springer, pp. 468–80. [\[CrossRef\]](#)
- Bateson, Gregory. 1955. A theory of play and fantasy; a report on theoretical aspects of the project of study of the role of the paradoxes of abstraction in communication. *Psychiatric Research Reports* 1: 39–51.
- Bateson, Gregory. 1972. A Theory of Play and Fantasy. In *Steps to an Ecology of the Mind*. New York: Ballantine Books, pp. 177–93.
- Blaette, Andreas, Julia Rakers, and Christoph Leonhardt. 2022. How GermaParl Evolves: Improving Data Quality by Reproducible Corpus Preparation and User Involvement. In *Proceedings of the Workshop ParlaCLARIN III Within the 13th Language Resources and Evaluation Conference*. Edited by Darja Fišer, Maria Eskevich, Jakob Lenardič and Franciska de Jong. Marseille: European Language Resources Association, pp. 7–15.

- Blätte, Andreas, and Andre Blessing. 2018. The GermaParl Corpus of Parliamentary Protocols. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Edited by Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo and et al. Miyazaki: European Language Resources Association (ELRA).
- Boydston, Amber E., Dallas Card, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2014. Tracking the development of media frames within and across policy issues. In *APSA 2014 Annual Meeting Paper*. Washington, DC: APSA, pp. 1–25.
- Burscher, Björn, Daan Odijk, Rens Vliegthart, Maarten de Rijke, and Claes H. de Vreese. 2014. Teaching the Computer to Code Frames in News: Comparing Two Supervised Machine Learning Approaches to Frame Analysis. *Communication Methods and Measures* 8: 190–206. [\[CrossRef\]](#)
- Card, Dallas, Amber Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Beijing: Association for Computational Linguistics, pp. 438–44.
- Card, Dallas, Serina Chang, Chris Becker, Julia Mendelsohn, Rob Voigt, Leah Boustan, Ran Abramitzky, and Dan Jurafsky. 2022. Computational analysis of 140 years of us political speeches reveals more positive but increasingly polarized framing of immigration. *Proceedings of the National Academy of Sciences of the United States of America* 119: e2120510119. [\[CrossRef\]](#)
- Cherukuri, Komala Subramanyam, Pranav Abishai Moses, Aisa Sakata, Jiangping Chen, and Haihua Chen. 2025. Large Language Models for Oral History Understanding with Text Classification and Sentiment Analysis. *arXiv*, arXiv:2508.06729. . [\[CrossRef\]](#)
- Cohen, Jacob. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20: 37–46. [\[CrossRef\]](#)
- Cruickshank, Iain J., and Lynnette Hui Xian Ng. 2024. Prompting and Fine-Tuning Open-Sourced Large Language Models for Stance Classification. *arXiv*, arXiv:2309.13734. [\[CrossRef\]](#)
- DiMaggio, Paul, Manish Nag, and David Blei. 2013. Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. *Poetics* 41: 570–606. [\[CrossRef\]](#)
- Du Bois, John W. 2007. The stance triangle. In *Stancetaking in Discourse: Subjectivity, Evaluation, Interaction*. Pragmatics & Beyond New Series. Amsterdam: John Benjamins Publishing Company, pp. 139–82. [\[CrossRef\]](#)
- Entman, Robert M. 1993. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication* 43: 51–58. [\[CrossRef\]](#)
- Feinstein, Alvan R., and Domenic V. Cicchetti. 1990. High agreement but low Kappa: I. the problems of two paradoxes. *Journal of Clinical Epidemiology* 43: 543–49. [\[CrossRef\]](#)
- Field, Anjalie, and Yulia Tsvetkov. 2019. Entity-centric contextual affective analysis. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence: Association for Computational Linguistics, pp. 2550–60.
- Gantner, Manuela. 2023. *Das “friedliche Atom” als Narrativ in Politik, Wissenschaft und Design. Ikonografie eines Diskurses zwischen 1953 und 1978*. Karlsruhe: KIT Scientific Publishing. [\[CrossRef\]](#)
- Gleitsmann, Rolf-Jürgen. 2011. Der Vision atomtechnischer Verheißungen gefolgt: Von der Euphorie zu ersten Protesten—Die zivile Nutzung der Kernkraft in Deutschland seit den 1950er Jahren. *Journal of New Frontiers in Spatial Concepts* 3: 17–26. [\[CrossRef\]](#)
- Goffman, Erving. 1974. *Frame Analysis: An Essay on the Organization of Experience*. Cambridge: Harvard University Press, pp. ix, 586.
- Gwet, Kilem L. 2014. *Handbook of Inter-Rater Reliability, 4th Edition: The Definitive Guide to Measuring the Extent of Agreement Among Raters*. Columbus: Advanced Analytics, LLC.
- Hamborg, Felix, Karsten Donnay, and Bela Gipp. 2019. Automated identification of media bias in news articles: An interdisciplinary literature review. *International Journal on Digital Libraries* 20: 391–415. [\[CrossRef\]](#)
- Hartmann, Mareike, Tallulah Jansen, Isabelle Augenstein, and Anders Søgaard. 2019. Issue framing in online discussion fora. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Edited by Jill Burstein, Christy Doran and Thamar Solorio. Minneapolis: Association for Computational Linguistics, pp. 1401–7. [\[CrossRef\]](#)
- Hasan, Kazi Saidul, and Vincent Ng. 2013. Stance Classification of Ideological Debates: Data, Models, Features, and Constraints. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*. Edited by Ruslan Mitkov and Jong C. Park. Nagoya: Asian Federation of Natural Language Processing, pp. 1348–56.
- Hofmann, Valentin, Xiaowen Dong, Janet Pierrehumbert, and Hinrich Schuetze. 2022. Modeling ideological salience and framing in polarized online groups with graph neural networks and structured sparsity. In *Findings of the Association for Computational Linguistics: NAACL 2022*. Seattle: Association for Computational Linguistics, pp. 536–50.
- Jing, Elise, and Yong-Yeol Ahn. 2021. Characterizing partisan political narrative frameworks about COVID-19 on twitter. *EPJ Data Science* 10: 53. [\[CrossRef\]](#)
- Jordan, Katrin. 2018. *Ausgestrahlt: Die mediale Debatte um “Tschernobyl” in der Bundesrepublik und in Frankreich 1986/87*. Medien und Gesellschaftswandel im 20. Jahrhundert. Göttingen: Wallstein Verlag.
- Jung, Matthias. 2013. *Öffentlichkeit und Sprachwandel: Zur Geschichte des Diskurses über die Atomenergie*. Berlin/Heidelberg: Springer.

- Jurkschat, Lena, Gregor Wiedemann, Maximilian Heinrich, Mattes Ruckdeschel, and Sunna Torge. 2022. Few-Shot Learning for Argument Aspects of the Nuclear Energy Debate. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Edited by Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani and et al. Marseille: European Language Resources Association, pp. 663–72.
- Kirchhof, Astrid Mignon. 2014. Spanning the Globe: West-German Support for the Australian Anti-Nuclear Movement. *Historical Social Research/Historische Sozialforschung* 39: 254–73.
- Klamm, Christopher, Ines Rehbein, and Simone Paolo Ponzetto. 2022. Frameast: A framework for second-level agenda setting in parliamentary debates through the lense of comparative agenda topics. In *Proceedings of the Workshop ParlaCLARIN III Within the 13th Language Resources and Evaluation Conference*. Marseille: European Language Resources Association, pp. 92–100.
- Kliment, Tibor. 1994. *Kernkraftprotest und Medienreaktionen*. Wiesbaden: Deutscher Universitätsverlag. [\[CrossRef\]](#)
- Kupietz, Marc, Harald Lungen, Paweł Kamocki, and Andreas Witt. 2018. The German Reference Corpus DeReKo: New Developments—New Opportunities. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Edited by Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo and et al. Miyazaki: European Language Resources Association (ELRA).
- Küçük, Dilek, and Fazli Can. 2020. Stance Detection: A Survey. *ACM Computing Surveys* 53: 12:1–12:37. [\[CrossRef\]](#)
- Kwak, Haewoon, Jisun An, Elise Jing, and Yong-Yeol Ahn. 2021. Frameaxis: Characterizing microframe bias and intensity with word embedding. *PeerJ Computer Science* 7: e644. [\[CrossRef\]](#)
- Landis, J. Richard, and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics* 33: 159–74. [\[CrossRef\]](#)
- Lepsius, Oliver. 2025. Wandel der Protest-Kultur und “Brokdorf-Beschluss”. In *Handbuch Bundesverfassungsgericht im Politischen System*. Edited by Robert Chr. van Ooyen and Martin H. W. Möllers. Wiesbaden: Springer Fachmedien, pp. 539–61. [\[CrossRef\]](#)
- Lorenz, Robert. 2011. *Protest der Physiker: Die »Göttinger Erklärung« von 1957 (1 ed.)*, Volume 3 of *Studien des Göttinger Instituts für Demokratieforschung zur Geschichte politischer und gesellschaftlicher Kontroversen*. Bielefeld: Transcript Verlag. [\[CrossRef\]](#)
- Luo, Yiwei, Kristina Gligorić, and Dan Jurafsky. 2024. Othering and low status framing of immigrant cuisines in us restaurant reviews and large language models. In *Proceedings of the International AAAI Conference on Web and Social Media*. Washington, DC: AAAI Press, vol. 18, pp. 985–98.
- Mendelsohn, Julia, and Ceren Budak. 2025. When people are floods: Analyzing dehumanizing metaphors in immigration discourse with large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vienna: Association for Computational Linguistics, pp. 8079–103.
- Mendelsohn, Julia, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. *Frontiers in Artificial Intelligence* 3: 55. [\[CrossRef\]](#)
- Meyer, Teva. 2024. Bordering Nuclearity: Very Low-level Radioactive Wastes’ Clearance and the Production of Spatial Nuclearities in Germany. *Historical Social Reserach* 49: 227–45.
- Mohammed Afzal, Osama, and Preslav Nakov. 2023. Team TheSyllogist at SemEval-2023 Task 3: Language-Agnostic Framing Detection in Multi-Lingual Online News: A Zero-Shot Transfer Approach. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Edited by Atul Kr. Ojha, A. Seza Doğruöz, Giovanni Da San Martino, Harish Tayyar Madabushi, Ritesh Kumar and Elisa Sartori. Toronto: Association for Computational Linguistics, pp. 2058–61. [\[CrossRef\]](#)
- Nguyen, Viet-An, Jordan Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. Tea Party in the House: A Hierarchical Ideal Point Topic Model and Its Application to Republican Legislators in the 112th Congress. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Edited by Chengqing Zong and Michael Strube. Beijing: Association for Computational Linguistics, pp. 1438–48. [\[CrossRef\]](#)
- Ogawa, Akihiro. 2023. *Antinuclear Citizens: Sustainability Policy and Grassroots Activism in Post-Fukushima Japan*. Anthropology of Policy. Stanford: Stanford University Press.
- Opp, Karl-Dieter, and Wolfgang Roehl. 1990. *Der Tschernobyl-Effekt*. Wiesbaden: VS Verlag für Sozialwissenschaften. [\[CrossRef\]](#)
- Otmakhova, Yulia, and Lea Frermann. 2025. Narrative Media Framing in Political Discourse. In *Findings of the Association for Computational Linguistics: ACL 2025*. Edited by Wanxiang Che, Joyce Nabende, Ekaterina Shutova and Mohammad Taher Pilehvar. Vienna: Association for Computational Linguistics, pp. 9167–96. [\[CrossRef\]](#)
- Otmakhova, Yulia, Shima Khanehzar, and Lea Frermann. 2024. Media Framing: A typology and Survey of Computational Approaches Across Disciplines. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Bangkok: Association for Computational Linguistics, pp. 15407–28. [\[CrossRef\]](#)
- Piskorski, Jakub, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. Semeval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Toronto: Association for Computational Linguistics, pp. 2343–61.

- Rashkin, Hannah, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin: Association for Computational Linguistics, pp. 311–21.
- Rauh, Christian, and Jan Schwalbach. 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*. [\[CrossRef\]](#)
- Reiter-Haas, Markus, Alexander Ertl, Kevin Innerhofer, and Elisabeth Lex. 2023. mCPT at SemEval-2023 Task 3: Multilingual Label-Aware Contrastive Pre-Training of Transformers for Few- and Zero-shot Framing Detection. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*. Edited by Atul Kr. Ojha, A. Seza Doğruöz, Giovanni Da San Martino, Harish Tayyar Madabushi, Ritesh Kumar and Elisa Sartori. Toronto: Association for Computational Linguistics, pp. 941–49. [\[CrossRef\]](#)
- Richter, Florian, Philipp Koch, Oliver Franke, Jakob Kraus, Lukas Warode, Fabrizio Kuruc, Stella Heine, and Konstantin Schöps. 2023. Open Discourse: Towards the first fully Comprehensive and Annotated Corpus of the Parliamentary Protocols of the German Bundestag. *Harvard Dataverse*. [\[CrossRef\]](#)
- Roose, Jochen. 2010. Der endlose Streit um die Atomenergie. Konfliktsoziologische Untersuchung einer dauerhaften Auseinandersetzung. In *Umwelt- und Technikkonflikte*. Edited by Peter H. Feindt and Thomas Saretzki. Wiesbaden: VS Verlag für Sozialwissenschaften, pp. 79–103. [\[CrossRef\]](#)
- Sagi, Eyal, Daniel Diermeier, and Stefan Kaufmann. 2013. Identifying issue frames in text. *PLoS ONE* 8: e69185. [\[CrossRef\]](#)
- Sap, Maarten, Marcella Cindy Prasettio, Ari Holtzman, Hannah Rashkin, and Yejin Choi. 2017. Connotation frames of power and agency in modern films. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen: Association for Computational Linguistics, pp. 2329–34.
- Sobhani, Parinaz, Diana Inkpen, and Xiaodan Zhu. 2017. A Dataset for Multi-Target Stance Detection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Edited by Mirella Lapata, Phil Blunsom and Alexander Koller. Valencia: Association for Computational Linguistics, pp. 551–57.
- Staas, Christian. 2021. Die ZEIT und die NS-Zeit. *Die Zeit*, May 6, pp. 16–17.
- Stammbach, Dominik, Maria Antoniak, and Elliott Ash. 2022. Heroes, villains, and victims, and gpt-3: Automated extraction of character roles without training data. In *Proceedings of the 4th Workshop of Narrative Understanding (WNU2022)*. Seattle: Association for Computational Linguistics, pp. 47–56.
- Steinbach, Hannah Mathilde, Imge Yüzüncüoğlu, Raluca Rilla, and Manfred Stede. 2025. Heroes, villains, and victims: Character narratives in the WPS agenda of the UNSC. Paper presented at the 21st Conference on Natural Language Processing (KONVENS 2025): Workshops, Hildesheim, Germany, September 10–13; pp. 54–73.
- Sullivan, Karen. 2023. Three levels of framing. *Wiley Interdisciplinary Reviews: Cognitive Science* 14: e1651. [\[CrossRef\]](#) [\[PubMed\]](#)
- Taulé, Mariona, Maria Antònia Martí, Francisco Manuel Rangel Pardo, Paolo Rosso, Cristina Bosco, and Viviana Patti. 2017. Overview of the Task on Stance and Gender Detection in Tweets on Catalan Independence. Paper presented at the Second Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2017), Berlin, Germany, September 12.
- Thomas, Matt, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*. Edited by Dan Jurafsky and Eric Gaussier. Sydney: Association for Computational Linguistics, pp. 327–35.
- Uekötter, Frank. 2022. *Atomare Demokratie: Eine Geschichte der Kernenergie in Deutschland*. Stuttgart: Franz Steiner Verlag.
- Upadhyaya, Apoorva, Wolfgang Nejdl, and Marco Fisichella. 2025. Interpretable zero-shot stance detection with proactive content intervention. *Information Processing & Management* 62: 104223. [\[CrossRef\]](#)
- van den Berg, Esther, Katharina Korfhage, Josef Ruppenhofer, Michael Wiegand, and Katja Markert. 2019. Not my president: How names and titles frame political figures. In *Proceedings of the Third Workshop on Natural Language Processing and Computational Social Science*. Minneapolis: Association for Computational Linguistics, pp. 1–6.
- Walter, Tobias, Celina Kirschner, Steffen Eger, Goran Glavaš, Anne Lauscher, and Simone Paolo Ponzetto. 2021. Diachronic Analysis of German Parliamentary Proceedings: Ideological Shifts through the Lens of Political Biases. Paper presented at the 2021 ACM/IEEE Joint Conference on Digital Libraries (JCDL), Champaign, IL, USA, September 27–30; pp. 51–60. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.