Augmented Political Leaning Detection: Leveraging Parliamentary Speeches for Classifying News Articles

Charlott Jakob
Technical University Berlin,
German Research Center for AI,
Berlin, Germany,
c.jakob@tu-berlin.de

Pia Wenzel
Technical University Berlin,
Berlin, Germany,
p.wenzel.2@campus.tu-berlin.de

Salar Mohtaj
Technical University Berlin,
German Research Center for AI,
Berlin, Germany,
salar.mohtaj@tu-berlin.de

Vera Schmitt
Technical University Berlin,
German Research Center for AI,
Berlin, Germany,
vera.schmitt@tu-berlin.de

Abstract

In an era where political discourse infiltrates online platforms and news media, identifying opinion is increasingly critical, especially in news articles, where objectivity is expected. Readers frequently encounter authors' inherent political viewpoints, challenging them to discern facts from opinions. Classifying text on a spectrum from left to right is a key task for uncovering these viewpoints. Previous approaches rely on outdated datasets to classify current articles, neglecting that political opinions on certain subjects change over time. This paper explores a novel methodology for detecting political leaning in news articles by augmenting them with political speeches specific to the topic and publication time. We evaluated the impact of the augmentation using BERT and Mistral models. The results show that the BERT model's F1 score improved from a baseline of 0.82 to 0.85, while the Mistral model's F1 score increased from 0.30 to 0.31.

1 Introduction

In an era increasingly dominated by digital landscapes, political discussion has largely migrated online. As readers engage with posts and news articles, they encounter not only the facts but also the authors' inherent viewpoints. This phenomenon is further complicated by the prevalence of misinformation and disinformation, making it difficult for individuals to maintain objectivity and discern the embedded viewpoints.

The presence of political viewpoints in news articles poses significant challenges for the average reader, potentially leading to cognitive biases. According to a recent study by French et al. (2023), readers, among others, fall for Confirmation Bias. It reinforces pre-existing beliefs, as users tend to agree with statements that foster their own opinions. Detecting political viewpoints in textual content is thus crucial for ensuring that readers can access a comprehensive view of the information presented.

Political opinion can be categorized differently. A common way is to categorize text on a left-to-right spectrum, known amongst others as political leaning detection (Doan and Gulla, 2022).

Previous research on political leaning has predominantly focused on English-speaking contexts, yet such models are not directly transferable to the German linguistic and political landscape, where left and right viewpoints on specific topics vary significantly. For example, topics considered right-leaning in the United States, such as gun ownership rights, may not be framed similarly in Germany (Doan and Gulla, 2022).

Moreover, the dynamic nature of political landscapes means that older datasets might not accurately reflect current political climates or the evolving positions of political parties. The recent shifts in German politics with the emergence of the party Alternative für Deutschland (AfD) in the German federal parliament (Bundestag) illustrates these changes, necessitating updated and relevant datasets for analysis. Furthermore, there are differences in legislation terms as the opposition parties usually use more emotional language in their speeches when criticizing the government parties (Bissmann et al., 2016).

This paper proposes a novel system for detecting political leaning. It leverages current speeches from the Bundestag to improve the classification of Ger-

Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 126–133

man news articles by augmenting the articles with timely quotes from right-leaning and left-leaning speeches.

2 Related Work

A political viewpoint is defined as one of a limited number of identifiable opinions on a political subject. In the realm of automatic political viewpoint detection in English text, terms like *ideology*, *leaning*, *party*, and *political bias* are frequently used interchangeably, distinguishing classes typically along a left-to-right spectrum (Doan and Gulla, 2022). Although the specific characteristics of left and right may vary in German text, research indicates a consistent association across countries. The left is typically linked to egalitarianism, progress, social freedom, internationalism, and state intervention in the economy, while the right is associated with tradition, authority, nationalism, and market liberalization (Ferreira and Rosas, 2014; Caprara and Vecchione, 2018). For clarity, we will use the term *political leaning* to denote a viewpoint positioned on the spectrum between left and right ideologies.

Numerous studies have employed traditional machine learning models to classify political leaning (Slapin and Proksch, 2008; Barberá, 2015; Temporão et al., 2018; Goet, 2019; Rheault and Cochrane, 2020). Due to the recent advancements in Deep Learning, the utilization of Large Language Models (LLMs) is increasingly coming into focus. RoBERTa (Liu et al., 2019) in combination with a back-translation technique for augmentation proved to be effective in the CLEF *Checkthat!* shared task on political bias detection (Da San Martino et al., 2023). Maab et al. (2023) proposed another augmentation method where text samples are expanded by samples with the same target and bias.

Varies research was conducted based on the idea of comparing left and right viewpoints on the same topic to highlight differences. In 2020, Roy and Goldwasser (2020) examined the appearance of words from politically opposed news outlets regarding specific topics. Later, Liu et al. (2022) improved classification by pre-training an LLM with article pairs that presented the same story but from different ideological perspectives.

A recurring challenge in the classification of news articles is the inherent media bias. Hong et al. (2023) proposed using multi-head attention mechanisms, applying both document-level and sentence-level labels to ensure that the sentences accurately reflect the document's overall political leaning. Moreover, techniques such as Adversarial Media Adaptation and Triplet Loss Pre-training are introduced by Baly et al. (2020) to mitigate the influence of media bias on model outcomes. Chen et al. (2020) suggested modelling article-level bias by analyzing sentence-level bias along with other features like frequency, position, and sequence of biased terms.

Incorporating political speeches has shown promising results in enhancing the performance of political leaning detection models. Krestel et al. (2012) analyzed the cosine similarities between articles and speeches to classify news outlets along a political spectrum, confirming the relevance of political speeches. Bissmann et al. (2016) applied logistic regression models trained on bag-of-words vectors from German political speeches. Hajare et al. (2021) aligned social media posts with congressional speeches using TF-IDF similarity.

In this study, we augment a news article dataset with text from political speeches. Building on the ideas of Liu et al. (2022) and Bissmann et al. (2016) we leverage discussions on the same topic from different parties in the Bundestag. By using sequence pairs from the furthest left and furthest right parties in the parliament, we train the model to recognize the two most divergent political viewpoints on a given topic.

3 Datasets

We use two corpora, a news article dataset from Aksenov et al. (2021) and German Parliament speeches from Open Discourse (Richter et al., 2020) to evaluate our proposed approach.

3.1 Articles

We use the German news article dataset initially presented by Aksenov et al. (2021) as a basis. The labels were assigned based on a survey carried out by Medienkompass.org, in which subjects were asked to rate different German media outlets on a scale of left to right. Crawling the media outlets' articles resulted in a set of 47,362 articles from 34 different publishers (Aksenov et al., 2021). We rerun the open-source news crawler provided by Aksenov et al. (2021), resulting in retrieving a subset of the original dataset as some links were broken. Before cleaning, the dataset consists of

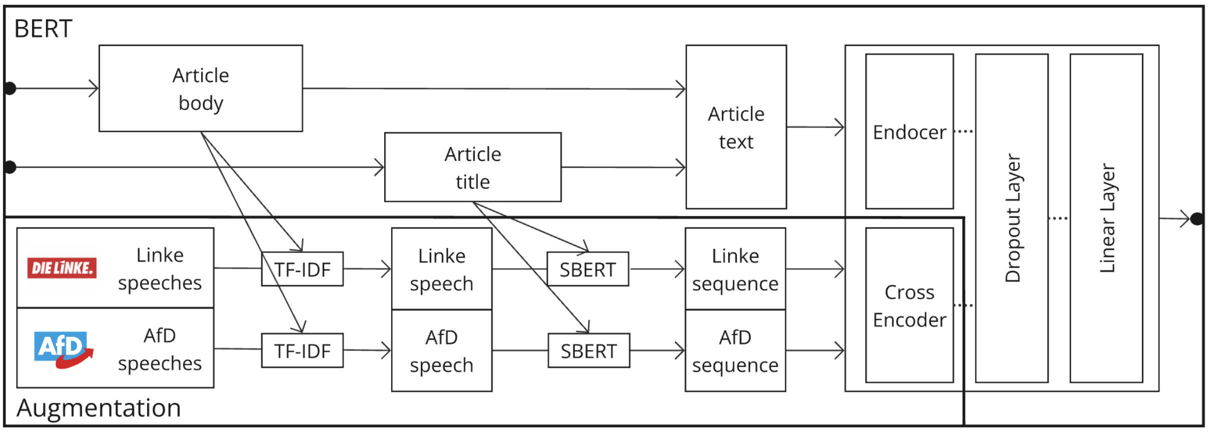

Figure 1: Augmentation pipeline and architecture of models BERT and BERT Augmented

46,191 articles from 2001 to 2021. We then filtered out articles published between March 2018 and December 2021, as AfD entered the Bundestag in October 2017 and we included speeches from six months before to six months after the article's publication date. We also removed articles with less than 100 words, which resulted in 13,831 articles in total. The length of the articles and the left- and right-side speech excerpts measured in word count are presented in Table 1

We decided to run the experiment in two scenarios that differed in the number of classes. We run the original 5 classes distribution and furthermore create new classes by splitting the news-outlet ratings from Medienkompass at quantiles 0.33 and 0.66, resulting in classes "left", "center", "right", often used in the context of political leaning detection (Doan and Gulla, 2022). The distribution of classes in both scenarios is presented in Table 2.

Dataset	max.	avg.	std.
Articles	16.320	865	1318
AfD sequences	185	45	20
Linke sequences	168	45	19

Table 1: Text length statistics within the datasets

3.2 Speeches

The speeches are obtained from a publicly available dataset called Open Discourse (Richter et al., 2020), which is the first fully comprehensive corpus of the plenary proceedings of the Bundestag. The dataset contains speeches ranging from 9.12.1949 to 20.5.2022. We prefilter the speeches from October 2017 to July 2022 to match the article dataset. Only speeches from the factions AfD and Linke, representing the far right and far left of the German political landscape, with more than 100 words

	Class	Total	Train	Eval.	Test
3 Classes Absolute	left	4.949	3.509	723	717
	center	5.972	4.136	921	915
	right	1.058	740	153	165
3 Classes Relative	left	41.31	41.85	40.23	39.90
	center	49.85	49.33	51.25	50.92
	right	8.83	8.83	8.51	9.18
5 Classes Absolute	far-left	210	148	31	31
	center-left	2.875	2.040	419	416
	center	2.450	1.726	373	351
	center-right	5.386	3.731	821	834
	far-right	1.058	740	153	165
5 Classes Relative	far-left	1.75	1.77	1.73	1.73
	center-left	24.00	24.33	23.32	23.15
	center	20.45	20.58	20.76	19.53
	center-right	44.96	44.50	45.69	46.41
	far-right	8.83	8.83	8.51	9.18

Table 2: Label distributions across the two class scenarios

were selected. This choice aims to provide a broad range of opinions from both ends of the political spectrum for model training.

3.3 Augmentation

To assess political leaning in news articles, we augment the content of the articles with parliamentary speeches of two political parties. We consider the articles to be the query in an Information Retrieval sense. For both parties, we follow the following procedure. All speeches of the party are considered as documents. In the first stage, we filter the documents based on the article's publication date, selecting those that fall within a window of six months before and six months after the publication. Then, we identify the most relevant document for our article by calculating the cosine similarity between the query and the documents' TF-IDF vectors. TF-IDF is well-suited for the initial broad document search as it's able to capture the broad context of large documents. On the other side, models like SBERT are designed to embed shorter sentences and capture semantic information which is beneficial for our second stage. There, we tokenize the selected speech, removing the first and last two sentences, as they primarily include the formal opening and closing phrases that follow established protocol. Also, sentences with fewer than 10 words are excluded to remove less informative content. We then break down the speech into sentences, treating each sentence as a separate document. Each sentence is embedded as well as the article title, and we identify the most relevant sentence based on cosine similarity with the title, using SBERT embeddings (Reimers and Gurevych, 2019). The most similar sentence and its following sentence are used as the final augmentation sequence. The final augmented dataset for classification includes five columns: the article text, the 3-point scale label, the 5-point-scale label, the most similar AfD sequence, and the most similar Linke sequence. For the dataset without augmentation, the only difference is the absence of the faction sentences. A visualization of this process is provided in Figure 1. Furthermore, an example of an article with the respective augmentation sequences can be found in Figure 2.

4 Models

We utilize a small encoder model and a small decoder model to demonstrate the effectiveness of speech augmentation. For the encoder, we fine-tune a BERT Model (Devlin et al., 2019) pre-trained on a German corpus. To incorporate larger generative models and the concept of Retrieval Augmented Generation, we conduct inference tests with an LLM equipped with a decoder. We use Mistral-7B-Instruct (Jiang et al., 2023) as it is a small European LLM and is fine-tuned on instructions, which helps to retrieve valid JSON outputs. The prompt details for Mistral are provided in the appendix. Both models are designed as classifiers. BERT is modified by appending a dropout and a linear layer, while Mistral utilizes few-shot learning to output a JSON containing the label prediction. We then compare the baseline models of BERT and Mistral against the augmented versions.

For BERT, we enhance its capabilities by expanding it with a cross-encoder model, allowing it to incorporate the augmentation sequences. Comparing left-leaning and right-leaning sequences may help to interpret the opinions expressed in the news article. Cross-encoders are typically used for similarity calculation in retrieval processes. We exploit this feature to encode contrasting sequences. It takes two sequences as inputs, and we take its last hidden states as input for the dropout layer. The BERT model baseline and expanded BERT model are illustrated in Figure 1.

In the case of the Mistral model, the augmentation process is straightforward. Left and right-leaning sequences are included in the instructions, as detailed in the appendix 5.

5 Experimental Setup

We use the German BERT base model[1] and a cross-encoder for multilingual support[2]. For Mistral, we use Mistral-7B-Instruct[3]. Due to input size limitations in both models, we use a maximum token size of 512 for the articles. For evaluation, we train our BERT models with a train-eval-test split of 70%, 15% and 15%. For Mistral, we perform few-shot learning by adding a left-leaning and right-leaning news article to the prompt and running it on the test set only. For evaluation metrics, we use marco precision, recall and F1 score.

We fine-tune the BERT models using grid search and run 20 random parameter variations. We always train 4 epochs and use batch sizes [8,16,32], learning rates [1e-5, 2e-5, 3e-5, 4e-5, 5e-5], weight decays [0.01, 0.1], warm-up ratios [0.06, 0.08, 0.1] and dropout rates of [0.1, 0.2, 0.3, 0.4].

6 Results

The evaluation metrics for four different models over two scenarios are summarized in Table 3. For the three-class scenario, the BERT Augmented model demonstrated the highest performance with precision, recall, and F1-score all at 0.85, surpassing the standard BERT model's scores of 0.83, 0.82, and 0.82, respectively. Both versions of the Mistral

[1]https://huggingface.co/dbmdz/bert-base-german-cased
[2]https://huggingface.co/cross-encoder/msmarco-MiniLM-L6-en-de-v1
[3]https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.1

Article text: „Babys an die Urne: Politiker wollen Wahlrecht ab Geburt [SEP] Babys an die Urne: Politiker wollen Wahlrecht ab Geburt Eine fraktionsübergreifende Gruppe von Bundestagsabgeordneten will das Wahlalter in Deutschland drastisch herabsetzen - auf null Jahre. Das berichtet der ""Spiegel"" in seiner neuen Ausgabe. Einen entsprechenden Antrag für ein Wahlrecht ab Geburt bereitet der FDP-Politiker und ehemalige Bundestagsvizepräsident Hermann Otto Solms vor. Die SPD-Politiker Ingrid Arndt-Brauer und Swen Schulz sowie Thomas Silberhorn von der CSU unterstützen den Vorstoß. ""Das Ziel ist mehr Generationengerechtigkeit"", sagte Solms, der bereits 2005 und 2009 versucht hat, das Wahlrecht für Kinder durchzusetzen. Der Antrag soll nach der Sommerpause ins Parlament gehen. Foto: Wähler in einem Wahllokal, über dts Nachrichtenagentur"

AfD-sequence: „Aber wir sollten dieser Ausdehnung des Wahlrechts oder, besser gesagt, der Abschaffung der Einschränkung des Wahlrechts hier nicht nur zustimmen, weil es uns Gerichte oder die UN vorgeben. Es ist eine Anerkennung der Menschenwürde, ein notwendiger Schritt des jahrzehntelangen Abschüttelns unserer alten Vorurteile gegenüber behinderten Menschen."

Linke-sequence: „Wir sollten uns zu einem Wahlrecht entschließen können – darüber sollten wir verhandeln –, das allen Parteien etwas abverlangt. Damit komme ich zu den Ideen, die aufseiten der Union vorherrschen."

Figure 2: Example article with respective augmentation sequences

model showed significantly lower metrics. The basic version results in a precision and recall of 0.32 and an F1-score of 0.30, and a small improvement was observed in the augmented version with a precision of 0.32, a recall of 0.33, and an F1-score of 0.31.

In the five-class classification, only slight improvements in precision or recall were observed. The BERT model achieved a precision of 0.67, recall of 0.69, and F1-score of 0.68, and the augmented BERT's scores of 0.67, 0.69, and 0.68. The Mistral models remained less effective, with the augmented version marginally bettering the precision and maintaining the same recall and F1-score as the non-augmented version (0.19 precision, 0.19 recall, and 0.16 F1).

	Model	Precision	Recall	F1
3 classes	BERT	0.83	0.82	0.82
	BERT Augm.	**0.85**	**0.85**	**0.85**
	Mistral	0.33	0.32	0.30
	Mistral Augm.	0.32	0.33	0.31
5 classes	BERT	**0.67**	0.69	**0.68**
	BERT Augm.	**0.67**	**0.70**	**0.68**
	Mistral	0.18	0.19	0.16
	Mistral Augm.	0.19	0.19	0.16

Table 3: Evaluation metrics for all models

Further analysis of the class-wise performance of the BERT Augmented model for the three-class problem reveals high precision and recall across 'left' (0.84 and 0.91), 'center' (0.95 and 0.88), and 'right' categories (0.77 and 0.76), culminating in

	Class	Precision	Recall	F1	Supp.
3 classes	left	0.84	0.91	0.87	723
	center	0.95	0.88	0.92	921
	right	0.77	0.76	0.77	153
	total	0.85	0.85	0.85	
5 classes	far-left	0.00	0.00	0.00	31
	center-left	1.00	0.97	0.99	416
	center	0.72	0.87	0.79	351
	center-right	0.97	0.91	0.94	834
	far-right	0.66	0.76	0.70	165
	total	0.67	0.70	0.68	

Table 4: Class-wise evaluation metrics for BERT Augmented

an average F1-score of 0.85 (Table 4).

In contrast, the five-class classification shows varied performance, with the 'far-left' class failing to identify any true positives (precision and recall at 0.00), while 'center-left' and 'center-right' classes have strong performances with F1-scores of 0.99 and 0.94, respectively. The 'center' and 'far-right' classes show moderate results, contributing to an overall average F1-score of 0.68 for the model.

7 Discussion and Limitations

The results from Mistral appear nearly random, indicating significant room for improvement. Currently, it operates primarily through inference, which might be inadequate. Fine-tuning Mistral on a specific task using the training dataset could enhance its performance. Data augmentation didn't significantly improve Mistral's performance, potentially because Mistral was trained on bigger and more current datasets than models like BERT.

The choice of using a Cross-encoder to enhance an encoder model warrants evaluation. Testing different configurations could reveal more effective alternatives. For instance, employing an encoder with a larger input size and directly appending speech augmentation to article tokens might optimize performance.

Our approach may reduce media bias by utilizing non-media datasets for comparison. However, the challenge of mitigating media bias persists due to our reliance on datasets labelled based on news outlets. Experimenting with a manually labelled German news dataset could provide a clearer indication of the effectiveness of our methodology.

The process of augmentation and similarity search has limitations. To improve, methods such as isolating subjective sentences could be investigated. Furthermore, considering speeches from a narrower time frame, such as within a specific two-month period, might yield more precise insights.

Our analysis currently limits itself to speeches from Linke and AfD. This restriction might oversimplify the complex political spectrum. Incorporating speeches from a broader range of political parties would add complexity and could provide a more comprehensive understanding. However, categorizing parties on a simple left-right scale is challenging as multiple dimensions influence political parties. Shared viewpoints among opposing parties might undermine the utility of the left-right scale. It raises the question of whether detecting political leaning or ideology is more suitable for analyzing the German political landscape.

8 Conclusion

The paper represents an initial attempt to use speeches to augment articles for the detection of political leaning. We explore this approach with two classic LLM methods: fine-tuning BERT and few-shot learning with Mistral. The study incorporated sequences from the most left-leaning and most right-leaning parties as a thematically and temporally relevant input for the articles. For both models, there were slight improvements in F1 scores, with BERT increasing from 0.82 to 0.85 and Mistral from 0.30 to 0.31. Further research is necessary to determine the effectiveness of augmentation.

References

Dmitrii Aksenov, Peter Bourgonje, Karolina Zaczynska, Malte Ostendorff, Julian Moreno-Schneider, and Georg Rehm. 2021. Fine-grained Classification of Political Bias in German News: A Data Set and Initial Experiments. In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 121–131, Online. Association for Computational Linguistics.

Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. We Can Detect Your Bias: Predicting the Political Ideology of News Articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4982–4991, Online. Association for Computational Linguistics.

Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. 23(1):76 – 91.

Felix Bissmann, Pola Lehmann, Daniel Kirsch, and Sebastian Schelter. 2016. Predicting political party affiliation from text. *PolText*, 14(14).

Gian Vittorio Caprara and Michele Vecchione. 2018. On the left and right ideological divide: Historical accounts and contemporary perspectives. *Political Psychology*, 39:49–83.

Wei-Fan Chen, Khalid Al-Khatib, Benno Stein, and Henning Wachsmuth. 2020. Detecting media bias in news articles using gaussian bias distributions. *arXiv preprint arXiv:2010.10649*.

Giovanni Da San Martino, Firoj Alam, Maram Hasanain, Rabindra Nath Nandi, Dilshod Azizov, and Preslav Nakov. 2023. Overview of the clef-2023 checkthat! lab task 3 on political bias of news articles and news media. *Working Notes of CLEF*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Tu My Doan and Jon Atle Gulla. 2022. A Survey on Political Viewpoints Identification. *Online Social Networks and Media*, 30:100208.

Ana Rita Ferreira and João Cardoso Rosas. 2014. *Left and right: the great dichotomy revisited*. Cambridge Scholars Publishing.

Aaron M. French, Veda C. Storey, and Linda Wallace. 2023. The impact of cognitive biases on the believability of fake news. *European Journal of Information Systems*, pages 1–22.

Niels D. Goet. 2019. Measuring polarization with text analysis: Evidence from the UK house of commons, 1811–2015. 27(4):518 – 539.

Prasad Hajare, Sadia Kamal, Siddharth Krishnan, and Arunkumar Bagavathi. 2021. A Machine Learning Pipeline to Examine Political Bias with Congressional Speeches. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 239–243, Pasadena, CA, USA. IEEE.

Jiwoo Hong, Yejin Cho, Jiyoung Han, Jaemin Jung, and James Thorne. 2023. Disentangling Structure and Style: Political Bias Detection in News by Inducing Document Hierarchy. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5664–5686, Singapore. Association for Computational Linguistics.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *CoRR*, abs/2310.06825.

Ralf Krestel, Alex Wall, and Wolfgang Nejdl. 2012. Treehugger or petrolhead?: identifying bias by comparing online news articles with political speeches. In *Proceedings of the 21st International Conference on World Wide Web*, pages 547–548, Lyon France. ACM.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.

Yujian Liu, Xinliang Frederick Zhang, David Wegsman, Nick Beauchamp, and Lu Wang. 2022. POLITICS: Pretraining with Same-story Article Comparison for Ideology Prediction and Stance Detection. *arXiv preprint*. ArXiv:2205.00619 [cs].

Iffat Maab, Edison Marrese-Taylor, and Yutaka Matsuo. 2023. Target-Aware Contextual Political Bias Detection in News. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 782–792, Nusa Dua, Bali. Association for Computational Linguistics.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.

Ludovic Rheault and Christopher Cochrane. 2020. Word embeddings for the analysis of ideological placement in parliamentary corpora. 28(1):112 – 133.

Florian Richter, Philipp Koch, Oliver Franke, Jakob Kraus, Fabrizio Kuruc, Anja Thiem, Judith Högerl, Stella Heine, and Konstantin Schöps. 2020. Open Discourse.

Shamik Roy and Dan Goldwasser. 2020. Weakly Supervised Learning of Nuanced Frames for Analyzing Polarization in News Media. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7698–7716, Online. Association for Computational Linguistics.

Jonathan B. Slapin and Sven-Oliver Proksch. 2008. A scaling model for estimating time-series PartyPositions from texts. 52(3):457–722.

Mickael Temporão, Corentin Vande Kerckhove, Clifton van der Linden, Yannick Dufresne, and Julien M. Hendrickx. 2018. Ideological scaling of social media users: A dynamic lexicon approach. 26(4):457 – 473.

A Appendix

<s>[INST]
TASK:
You are a language model tasked with analyzing the political leaning of an article. Given an article **and opinions on the topic from left and right politicians**, your goal is to:
1. Provide a brief 2-sentence elaboration on the author's perceived political viewpoint based on the article's content **and the left and right-leaning opinions**.
2. Categorize the article's political leaning as "left", "center", or "right".

OUTPUT FORMAT:
Your response must be provided as a JSON object with the following keys and values:
“‘json
{{
"elaboration": "<Two sentences elaborating on the articles political leaning>",
"political leaning": "<left, center or right>"
}}
“‘

EXAMPLE 1:
Article: "Die Welt bereitet sich auf einen Krieg vor [SEP] Michail Gorbatschow, hier bei seinem 85. Geburtstag im März 2016, sieht die Welt auf einen Krieg zusteuern Michail Gorbatschow ist ein Mann des Friedens. Sein Name steht für atomare Abrüstung, Perestroika und Glasnost, den Fall der Berliner Mauer und eine Annäherung der damaligen Supermächte USA und Russland. Der letzte Staatschef der Sowjetunion glaubt sein Erbe im Gefahr. „Es sieht aus, als würde die Welt sich auf einen Krieg vorbereiten“, schreibt der 85-Jährige in einem Beitrag für das „Time“-Magazin. Seine Beobachtung: Es gebe ein neues Wettrüsten, das dringend gestoppt werden müsse. Doch das Gegenteil sei der Fall. Soldaten und schwere Waffen von Nato und Russland würden in Europa immer näher zusammenkommen – jeweils in Schlagdistanz. „Während Staatshaushalte darum kämpfen, die grundlegenden sozialen Bedürfnisse der Menschen zu finanzieren, wachsen die Militärausgaben.“ Das Geld fließe in schwere Waffen und Raketenabwehrsysteme, die die „strategische Stabilität untergraben“. "
Left leaning opinion: "Das Fazit: Der Gründungsgeist der UNO „Frieden durch Diplomatie“ muss endlich wieder gestärkt werden. Die Bundesregierung hat leider mit ihrer Politik in den Vereinten Nationen mehr geostrategische NATO-Politik denn eine aktive UNO-Friedenspolitik betrieben."
Right leaning opinion: "Ich darf noch einmal in Erinnerung rufen: Michail Gorbatschow sprach beim Ende des Kalten Krieges vom „gemeinsamen Haus Europa“. François Mitterrand sah mit dem Ende des Kalten Krieges für Europa die Möglichkeit, zu seiner eigenen Geschichte und seiner eigenen Geografie zurückzukehren, so wie man zu sich nach Hause zurückkehrt."
Output:
“‘json
{{
"elaboration": "The article emphasizes the need for disarmament and criticizes the increase in military spending and aggressive political rhetoric, aligning closely with traditional left-wing values that prioritize social welfare over militarization. ",
"political leaning": "left"
}}
“‘

EXAMPLE 2:
Article: "Corona: Wie gefährlich ist das Virus? Wie kann ich mich besser schützen? [SEP] Fast kein Tag vergeht mehr, wo nicht neue Verbote der Bewegungsfreiheit in Deutschland beschlossen werden. Das düstere Beispiel von Italien und Spanien vor Augen jagt eine Einschränkung die nächste. Was kommt als nächstes? Zahlen wir mit dem Kollaps unserer Wirtschaft nicht einen viel zu hohen Preis? Führt das Kontaktverbot nicht auch zu einem „Bewegungsverbot“, was uns anfälliger für Corona werden lässt? Wie gefährlich ist dieser Virus, der uns allen den Atem nimmt? Wie kann ich mich vor Corona schützen? Spätestens seitdem Virologen die Macht übernommen haben, scheinen wirtschaftliche Grundregeln nicht mehr zu gelten. Wer den gigantischen Wirtschaftskreislauf runterfährt und ihn in Bereichen gar stoppt, muss triftige Gründe haben, dies zu tun. "
Left leaning opinion: "Deshalb unterstützen wir nachdrücklich noch einmal die Forderungen des DGB und des BDI, die sagen: Auch wegen Corona müssen wir jetzt ein Zeichen setzen, dass wir aus dieser Krise herauskommen werden. – Wir fordern nochmals in den nächsten zehn Jahren zusätzliche Investitionen in Höhe von jährlich 45 Milliarden Euro."
Right leaning opinion: "Wir wissen jedoch schon, dass Coronaviren zu den schnell mutierenden Virusstämmen gehören. In bester sozialistischer Manier werden also Impfstoffentwickler mit Steuergeldern unterstützt, um dann Gewinne – allein die erste Impfwelle verspricht 20 Milliarden Dollar Umsatz – privat einzustreichen, und das für einen Virus, dessen Gefährlichkeit nach heutigem Wissensstand weit geringer ist, als zunächst angenommen."
Output:
“‘json
{{
"elaboration": "The article conveys a right-leaning viewpoint, emphasizing skepticism towards the severity of the coronavirus and the economic repercussions of restrictive measures. It criticizes the disruption of economic norms and suggests that the virus might not justify the extensive government interventions and limitations on movement. ",
"political leaning": "right"
}}
“‘

INSTRUCTIONS:
You will be provided with an article text. Your task is to analyze the content and generate a JSON object following the specified OUTPUT FORMAT. Provide your elaboration on the article's political leaning in two sentences under the "elaboration" key, and categorize it as "left", "center", or "right" under the "political leaning" key.
Strictly adhere to the JSON OUTPUT FORMAT. Do not include any other text in your response.

Article: {article}
Left leaning opinion: {linke sequence}
Right leaning opinion: {afd sequence}

[/INST]
Output: {{

Table 5: Input Sequence for Mistral and Mistral Augmented, where bold text is additionally for Mistral Augmented