AfD-CCC: Analyzing the Climate Change Discourse of a German Right-wing Political Party

Manfred Stede, Ronja Memminger
Applied Computational Linguistics
University of Potsdam
{stede|memminger}@uni-potsdam.de

Abstract

While the scientific consensus on anthropogenic climate change (CC) has long been undisputed, public discourse is still divided. Considering the case of Europe, in the majority of countries, an influential right-wing party propagates climate skepticism or outright denial. Our work addresses the German party *Alternative für Deutschland*, which represents the second-largest faction in the federal parliament. In order to make the party's discourse on CC accessible to NLP-based analyses, we are compiling the *AfD climate change corpus*, a collection of parliamentary speeches and other material from various sources. We report on first analyses of this new dataset using sentiment and emotion analysis as well as classification of populist language, which demonstrate clear differences to the language use of the two largest competing parties (social democrats and conservatives). We make the corpus available to enable further studies of the party's rhetoric on CC topics.

1 Introduction

In 2019, a study by a political consultation company analyzed the climate change (CC) policy agendas of 21 right-wing populist parties in European countries (Schaller and Carius, 2019). Interestingly, they found that the positions are not as homogeneous as one might think. The parties were categorized into these three types (p. 10 ff.):

1. *Denialist/skeptical* parties cast doubt on the scientific consensus on human-induced climate change or explicitly reject evidence beyond reasonable doubt.

2. *Disengaged/cautious* parties either have no position on climate change or attribute little importance to the problem.

3. *Affirmative* parties support the scientific mainstream and recognize the danger that climate change poses to the world and their own countries.

In group (3) there are three parties that acknowledge the problem and see a need for action, though this does not necessarily translate into ambitious goals for their national policies.[1] The biggest group is (2) with eleven parties, while group (1) consists of seven parties.

In this paper, we address the case of the German *Alternative für Deutschland* ('Alternative for Germany'), for short AfD. They belong to group (1) above, and have made their position rather explicit on many occasions. Well-known are, for example, quotes from the AfD MP Steffen Kotré, who said in November 2018 in the parliament that there is no scientifically-proven correlation between CO_2 in the atmosphere and the temperature on planet earth.[2]

In the 2025 elections, the AfD captured 20.8% of the vote (exactly twice as much as in the 2021 elections). While climate change was not a prominent topic in the election campaigns at all, nonetheless the result indicates that a sizable proportion of the German public is sympathetic toward the denialist/skeptical position. For comparison, in a poll conducted by the major public TV station *ARD* in December 2023[3], 62% of participants agreed with the statement that "Germany is already doing a lot for climate protection; now it is time for other countries to move forward"; and 60% agreed with "In the climate debate, there is too much propagation of fear".

The climate movement is well aware that new *communication strategies* are needed in order to

[1]The three parties are the (governing) Hungarian *Fidesz*, the Latvian *National Alliance* and the Finnish *Finns Party*.

[2]https://skepticalscience.com/Politiker-und-Falschinformationen-SKotre.shtml

[3]https://de.statista.com/statistik/daten/studie/1427817/umfrage/umfrage-zu-einstellungen-zum-klimaschutz/

Proceedings of the Fourth Workshop on NLP for Positive Impact (NLP4PI), pages 163–174
July 31, 2025

bring the topic back on the political agenda these days. As a prerequisite, we believe that tracking the discourse of highly-influential right-wing parties is an (if small) contribution that NLP can make. For the case of Germany, in this paper we present the *AfD climate change corpus (AfD-CCC)*, a set of different types of documents that were in recent years issued by party officials, comprising political speech, social media posts, and miscellaneous documents. Section 2 discusses the background and related work. Then, in Section 3, we describe the composition of the corpus. As a first use case, in Section 4 we use both a lexicon-based and a transformer model to analyze *populist language* and a transformer model for *emotions*, and show that the AfD results are quite different from the politically-central parties. Section 5 draws some conclusions and discusses the relevance of our work for the notion of *positive impact* by means of NLP.

2 Background

2.1 Corpora on climate change discourse

Several corpora with materials from the discourse on CC have been assembled in recent years (cf. (Stede and Patz, 2021)). This includes, for instance, the Climate-Fever dataset that specifically collects claims made in the domain of CC (Diggelmann et al., 2021); the richly-annotated *ClimaConvo* set of 15,000 tweets (Shiwakoti et al., 2024); a corpus of a 4-year period covering German parliamentary speeches, tweets and press releases by six parties (Schaefer et al., 2023); a multimodal corpus composed of scientific papers, IPCC reports and content from NGO websites (Volkanovska et al., 2025); or a subset of New York Times articles from the NYTACC corpus that deal with the CC topic (Grasso et al., 2024). We are, however, not aware of a corpus that focuses on political texts related to CC and is representative for the portion of the discourse that is run by an influential political entity, such as we are providing here with the AfD-CCC.

2.2 Detecting populist language

Populist language, in the simplest terms, can be considered to be the linguistic expression of populist narratives, such as a contrast of *the people* versus *the elite* (Mudde, 2004) or a rejection of 'established' political parties or institutions. This has been the basis for manual coding and evaluation of populist language (e.g., Küppers 2022; Sturm 2020 for the AfD's CC-related rhetoric). Several studies have utilized manual annotation of parliamentary speech for *people* and *elite* categories to facilitate the training of classification models. Klamm et al. (2023) annotate German and European parliamentary speeches for hierarchical mentions of *people* and *elites*. They apply transformer-based models to detection and classification of mentions, and report that detection of *people*-centric mentions is particularly successful as an identifier of populist language in large text (Klamm et al., 2023).

In a similar approach, Erhard et al. (2025) present PopBERT, a BERT-based model fine-tuned on manually-annotated transcripts of German parliamentary debates. Specifically, sentences are annotated for containing populist elements, resulting in a multi-label classification task. They are annotated for the two main aspects of populism: anti-elitism and people-centrism, as well as left- or right-wing ideology. PopBERT, they report, performs best on anti-elitism labeling (F1=0.84, with F1=0.71 on people-centrism).

Beyond political framing, the lexical and syntactic makeup of populist language may present another dimension for its detection. This extends beyond populist views and their expression and to descriptive analysis of known populist speakers. In a study of the linguistic complexity of populist language, Zanotto et al. (2024) utilize logistic regression and mixed-effect models on IMPAQTS, a corpus of Italian parliamentary speeches. They measure textual, lexical, and syntactic complexity features to find potential predictors of populist language. While no reliable predictor could be identified, they note some features of populist language, such as a tendency toward using proper nouns, absolutist language, and repetitive subjunctive clauses (Zanotto et al., 2024).

On a quantitative level, there have also been dictionary-based approaches to the detection of populist language (Bischof and Senninger, 2018; Rooduijn and Pauwels, 2011; Bonikowski and Gidron, 2015). For the German language specifically, Gründl (2022) presents a dictionary of populist terms and phrases based on analysis of social media posts by political figures in Germany, Austria, and Switzerland. Rather than simply collecting lemmas, the dictionary by Gründl (2022) contains regular expressions covering singular tokens and multi-word expressions. The domain of social media presents a different rhetorical context to that of the studies above; here, the party official is speaking directly to the public rather than to colleagues.

Populist language on social media is found to be especially dense (Gründl, 2022), which lends this channel a unique ability to capture linguistic features of populism.

2.3 Detecting emotions

Following up on the success of sentiment analysis, emotion analysis established itself as method for providing more fine-grained – though more difficult – accounts of subjectivity in text. For English, much work has been based on the NRC emotion lexicons (Mohammad and Turney, 2013); research on German has been done, for example, by Troiano et al. (2019). For political text, sentiment analysis has been applied for a long time, while work on emotions is much more scarce. As an example, Cochrane et al. (2021) undertook a computational analysis of a multimodal corpus of Canadian parliamentary speeches. Turning to the German language, the study by Widmann and Wich (2023) compared methods using lexicons, embeddings and transformers, and made available the tool that we will be using in Section 4.3.

3 Data Collection

We constructed our corpus primarily from sources that are open to the public, and this material is what we are making available with this paper.[4] It consists of:

- Speeches by AfD MPs in the German parliament (Bundestag), 2017-2022. We extracted them from the *OpenDiscourse* corpus (Richter et al., 2020).
- Speeches by AfD members in the European parliament (2014-2024). We extracted them from the dataset ParlLawSpeech[5] that was made available by Schwalbach et al. (2025).
- Press releases by AfD MPs, 2017-2021. We extracted them from press releases made available by Schaefer et al. (2023).
- Tweets by official AfD MPs, 2017-2022. We extracted them from a large set of tweets provided by Lasser et al. (2022).
- Telegram posts from public channels of the AfD and its former youth organization *Junge Alternative* ('Young Alternative', JA for short)[6], 2019-2025. We retrieved them using FROG, a tool for Telegram data extraction by Primig and Fröschl (2024).

In addition, for our analysis to be reported in the next section we drew on two sources of material that at this point cannot be made freely available:

- A set of Tweets issued by official AfD accounts in 2023 and 2024
- *Mitgliedermagazin Kompakt*, an AfD membership magazine of short articles posted online. We analyzed articles from 2018-2020, provided by Küppers (2022).

We filtered all texts for climate change topicality by using a set of keywords that (Schaefer et al., 2023) had employed to build a German climate-text corpus (see Sct. 2.1). Filtering was done on the level of paragraphs, i.e., our final data set is a collection of climate-related paragraphs taken from larger documents. We opted for this approach as in particular the parliamentary speeches usually address many different topics, and we wanted to eliminate the non-topical material as much as possible.

Telegram channel posts required additional preprocessing. They were split into paragraphs and the text cleaned by removing emoticons and special characters, as well as promotional segments frequently found at the end of posts in some channels. Table 1 provides an overview of the individual data sizes after filtering. We tokenized the texts using the spaCy model `de-core-news-lg`[7], with stopwords removed in the process. Phrases addressing the president and colleagues typically used at the beginning of speeches in the European Parliament were also removed. Punctuation is excluded from the token count.

Our distribution of the AfD-CCC provides both the original text of the paragraphs (with only Telegram posts minimally cleaned for noise, as described above), and the tokenized text. Additional metadata provided by the corpus includes the text's author (if applicable), the date of publication (or, in the case of speeches, the date of delivery), and a unique id.

Aside from the AfD-CCC, for our analyses in the next section, we use texts from the two other

[4] https://github.com/rmemminger/afd-ccc

[5] https://parllawspeech.org/, File Version 1.0.0

[6] These channels are: the AfD faction in the German and the Brandenburg state parliament, as well as the AfD Rhineland-Palatinate channel, and 5 state-level JA channels.

[7] https://spacy.io/models/de

Subcorpus	Tokens	Sentences
European Parliament	10,195	916
German Parliament	44,644	4,158
Press releases	30,633	2,800
Telegram	7,438	727
Sum	**92,910**	**8,601**
Twitter	76,197	5,492
Magazine	5,163	460
Sum Total	**174,270**	**14,553**

Table 1: Composition of the AfD-CCC dataset. The parts above the line are the subcorpora that we make publicly available.

	SPD		**CDU/CSU**	
	Sent.	Tokens	Sent.	Tokens
GP	11,069	105,271	13,822	135,280
EP	1,074	10,932	1,984	19,742
Press	819	8,059	460	4,524
Twitter	9,804	78,575	8,239	66,449
Sum	**22,766**	**202,837**	**24,505**	**225,995**

Table 2: Composition of the comparison datasets for SPD and CDU/CSU, given in sentence counts and token counts excluding punctuation, where GP = German Parliament, EP = European Parliament, Sen. = sentences.

largest parties in the German Parliament—as of the 2025 election—the SPD (*Sozialdemokratische Partei Deutschlands*, 'Social Democrat Party of Germany') and the Christian Democratic & Social Unions CDU & CSU, which share a faction in the German Parliament and are thus considered as a pair. We draw from a subset of the same data sources utilized for the AfD-CCC and collect German and European parliament speeches, tweets by official accounts, as well as press releases for the SPD and CDU/CSU respectively. For European parliament speeches, we only consider those delivered between 2013-2024. This is done not only to delimit the dataset, but also to better represent the rhetorical climate in the parliament since the AfD entered it in 2013. Table 2 illustrates the sizes of the comparison datasets' subsets in sentence and token counts.

4 Analyses: Languages of Populism and Emotion

In a first use case of the AfD-CCC, we address the question of the presence of linguistic features of populism and of emotions. For populism, our main instrument is the lexical approach by Gründl (2022), and as a secondary method we check whether the PopBERT model (Erhard et al., 2025) yields comparable results. For emotions, we make use of the BERT-based model by Widmann and Wich (2023). Additionally, we perform sentiment analysis using German-sentiment-BERT by Guhr et al. (2020).

4.1 Populism: Lexicon-based analysis

We utilize the Populism Dictionary by Gründl (2022) to perform a quantitative, lexical analysis of the party's CC-rhetoric, in comparison to that of CDU/CSU and SPD. That is, we examine how frequently markers of populist language, as defined by the dictionary, occur in paragraphs relating to CC. For this, we perform matching operations with the regular expressions in the dictionary with the paragraphs and calculate their relative frequencies. Importantly, we apply this to the cleaned but not the tokenized text, so as to retain the negation structures and other grammatical morphemes contained in the dictionary entries. We count the number of matches identified and normalize the frequency by dividing absolute counts by the size of the text corpus (in number of tokens)[8].

Party	Matches	Frequency
AfD	840	0.48%
SPD	142	0.07%
CDU/CSU	207	0.09%

Table 3: Lexicon-based analysis results the AfD-CCC, as well as the SPD and CDU/CSU dataset for comparison, given in the total count of dictionary matches over the texts, as well as the relative frequency of matches derived from them.

Table 3 lists the absolute counts of dictionary entry matches and relative match frequencies for all three parties. We find that, compared to the other two parties, the AfD-CCC produces more than five times as many matches to the dictionary. While the datasets vary in size, with the AfD-CCC being the smallest (likely due, in part, to the fact that the AfD has been afforded less speaking time in parliaments), they remain comparable. The AfD exhibits a match-frequency of 0.48% over a corpus of 174,270 tokens, while the SPD only achieve 0.07%

[8]Note that this is not a percentage of tokens matching the dictionary, as tokenizing removed stop-words, and dictionary entries are often multi-token expressions. The token count is used as a measure of corpus size that is more reliable than the number of documents, as their length varies.

(over 202,837 tokens), and the Union a frequency of 0.09% (over 225,995 tokens). The comparison parties, then, display similar levels of populist language, as according to the dictionary by Gründl (2022), and remain generally low in frequency compared to the AfD.

	AfD	SPD	CDU/CSU
German P.	0.55%	0.09%	0.1%
European P.	0.5%	0.08%	0.06%
Press	0.52%	0.05%	0.06%
Twitter	0.37%	0.05%	0.02%
Telegram	0.93%	-	-
Magazine	0.58%	-	-

Table 4: Relative frequencies of populism dictionary matches in the texts for each party, over each text domain, where P. = Parliament.

When looking into the different text domains underlying the datasets, this tendency remains. Table 4 displays the match frequencies (calculated as described above) for each party over each text domain. The SPD and CDU/CSU datasets did not include Telegram or magazine data. When considering only texts from domains available for all parties, speeches delivered in parliaments contain the most dictionary matches for all parties. Within each domain, the AfD consistently produces more than five times as many matches as SPD and CDU/CSU.

Turning to the AfD-CCC itself, we find that Telegram messages display the highest match frequency, at 0.93%. This may be due to the fact that not only is it a small subset of the corpus, but the Telegram channels are also not necessarily operated by AfD MP's. Almost half of the Telegram data is made up of paragraphs from the channel of the AfD faction in the Brandenburg state parliament (with 3,868 tokens and a match frequency of 1.14%). The five youth organization *Junge Alternative* channels in total contribute 1,882 tokens and produce 15 matches (match frequency 0.8%). The official channel of the AfD's faction in the German Parliament showed a comparatively low match frequency of 0.5% (5 matches over 1,009 tokens).

Overall, we find that across all domains, be they transcribed speeches or published text, the AfD consistently out-scores the other parties in match frequencies. The parties also consistently display a pattern of higher match frequency for populist language in parliament speeches and press releases, than tweets.

4.2 Populism: BERT-based analysis

The out-of-the-box tool PopBERT (Erhard et al., 2025) was made available just very recently, and we ran a first experiment to check whether this model confirms our findings regarding differences in the language use of AfD, SPD and CDU/CSU. For this, we used the climate-related paragraphs from the Bundestag speeches held by speakers of the three parties between 2017 and 2021. The total numbers of tokens are: AfD - 107,487; SPD - 208,897; CDU/CSU - 266,477.

Following the training strategy of PopBERT (cf. Sct. 2.2) and the underlying codebook, we computed the following values on sentence level and then aggregated the results to averages for the complete texts by a party:

- Anti-Elite: The sentence conveys resentment toward the ruling parties or toward established and influential organizations ("those up there").
- People-Centric: The sentence makes a statement fromt the perspective of "the normal people."

A sentence that coders labeled with one or both tags from above can in addition be labeled with:

- Host-Left: In the sentence, a left-wing "host ideology" (e.g., an argument from a class-based analysis) can be discerned.
- Host-right: Likewise, for a right-wing ideology.

Table 5 shows the results of applying the model to the corpora. While people-centrism is distributed evenly across the parties, for all other dimensions, the values for SPD and CDU/CSU are very similar, but they differ notably from those for the AfD; notice especially their high values for anti-elite and host-right. We thus conclude that the smaller-scale experiment with a different method confirms the results that we found with the lexicon-based approach.

4.3 Sentiment: German-Sentiment-BERT

To accompany the more fine-grained emotional analysis of CC-rhetoric undertaken in Section 4.4, we perform sentiment analysis of the texts using German-sentiment-BERT[9] (Guhr et al., 2020).

[9]https://huggingface.co/oliverguhr/german-sentiment-bert

Party	Anti-Elite	People-Centric	Host-Left	Host-Right
AfD	0.283	0.017	0.024	0.072
SPD	0.037	0.018	0.005	0.001
CDU	0.034	0.013	0.002	0.004

Table 5: PopBERT results for four populism dimensions on the climate-related parts of Bundestag speeches 2017-2021

	AfD	SPD	CDU/CSU
positive	3.00%	5.42%	4.82%
negative	13.89%	9.72%	9.07%
neutral	83.12%	84.87%	86.12%

Table 6: Results of sentiment classification using German-sentiment-BERT, given in percent of sentences classified as the respective sentiment over the total number of sentences.

This sentiment classification model is trained on German language texts, such as social media posts and reviews (Guhr et al., 2020). Classification is done on the sentence level, and sentiment can be one of three: *positive*, *negative*, or *neutral*. Sentence counts for the AfD-CCC can be found in Table 1 and for comparison parties in Table 2. The results are given in Table 6. For all three parties the predominant sentiment is *neutral* (AfD: 83.12%, SPD: 84.87%, CDU/CSU: 86.12%). While for all three parties there is a greater amount of *negative* than *positive* sentences, the distribution differs primarily for the AfD, with 13.89% *negative* and 3% *positive*, whereas the comparison parties display roughly 5% *positive* and roughly 9% *negative* sentiment.

4.4 Emotions: Pol_emo_mDeBERTa2

While emotions are not inherently markers of populism, an analysis of the emotional undercurrents of populist language can nonetheless contribute to the study of its effects on the audience. We therefore perform an analysis of the emotional aspects of the language in the AfD-CCC using a combination of the dictionary- and transformer-based approach. To this end we follow Widmann and Wich (2023), who present the German emotion dictionary *ed8* in an effort to mitigate the shortcomings of using valence-based, bag-of-words emotion dictionaries. Ed8 contains 20,582 terms and "is capable of measuring language associated with eight different emotions: anger, fear, disgust, sadness, joy, enthusiasm, pride, and hope" (Widmann and Wich, 2023, p. 629). It is well-suited for analyzing the AfD-CCC, as it was developed specifically to capture emotion in German political text. It does not only consider "emotional terms", but also "words that hint toward the presence of a specific emotional appeal that might be appraised by humans as such" (Widmann and Wich, 2023, p. 630).

Alongside the dictionary and implementations presented in their paper, Widmann and Wich (2023) have since released pol_emo_mDeBERTa2[10], a fine-tuned multilingual BERT model (mDeBERTa-v3-base)[11] that functions as a multi-label text classifier for the emotions in the ed8 dictionary. We apply this model to the AfD-CCC, as well as the comparison datasets for SPD and CDU/CSU. For the experiment, we followed the suggested implementation as given in the repository, which sets the decision threshold for labeling as 0.65. As in Section 4.3, classification is done on sentence level. Pol_emo_mDeBERTa2 returns binary scores for each of the 8 emotions per instance, whereby 1 signifies the presence of said emotion in the sentence (and 0 its absence). A sentence can contain zero, one, or several emotions. To evaluate the distribution of emotional language across the texts for each party, we calculate the percentage of each emotion's presence as the number of sentences classified as containing said emotion divided by the total number of sentences in the party's dataset.

The resulting distributions for each emotion are shown in Figure 1. For the AfD, the most prevalent emotion by a large margin is *anger* (35.6%), with all other emotions detected in less than 2% of all sentences (*fear*: 1.7%, *enthusiasm*: 1.1%, *joy*: 0.7%, *sadness*: 0.4%, *hope*: 0.3%, *disgust*: 0.2%, *pride*: 0.1%). Conversely, the SPD and CDU/CSU display a broader range of emotional speech. While *anger* remains the most prevalent (SPD: 8.7%, CDU/CSU: 8.2%), other emotions are detected more frequently than for the AfD, such as *enthusiasm* (SPD: 5%, CDU/CSU: 4.8%) and *joy* (SPD: 3.6%, CDU/CSU: 2.9%). The SPD and CDU/CSU display similar distributions between each other, and are generally found to produce more positive emotions (joy, enthusiasm, pride, hope) than the AfD.

We find, therefore, that while all three parties dis-

[10] https://github.com/tweedmann/pol_emo_mDeBERTa2

[11] https://huggingface.co/microsoft/mdeberta-v3-base

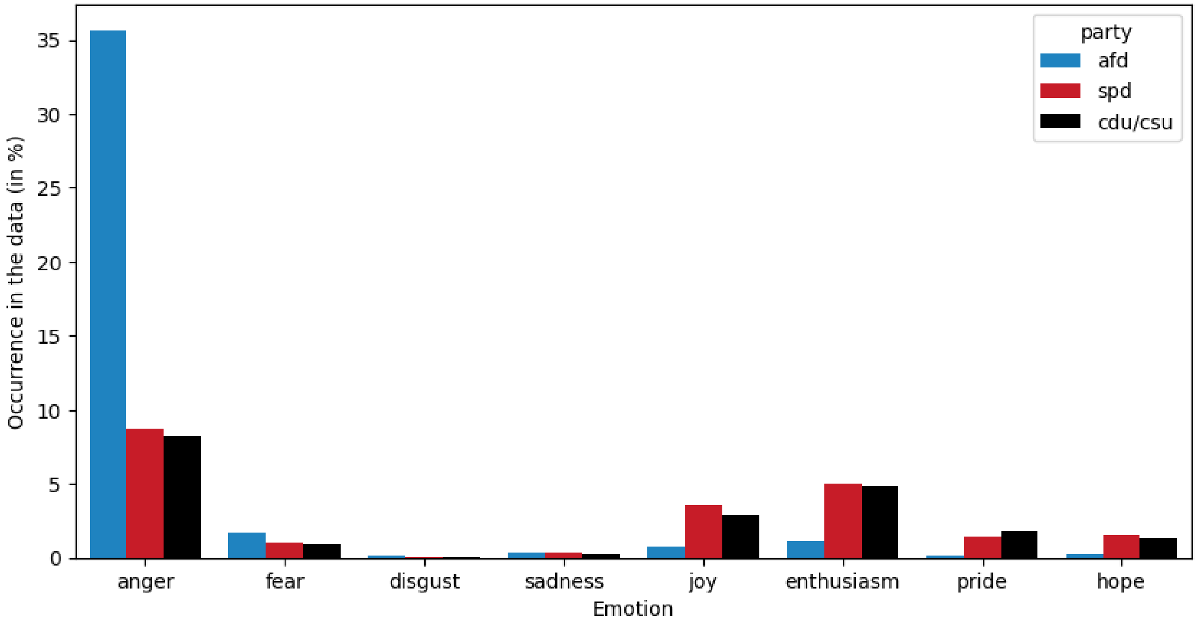

Figure 1: Results of emotion-classification on the AfD-CCC and comparison datasets for the SPD and CDU/CSU using pol_emo_mDeBERTa2. Distributions are given for each emotion and differentiated by party.

cuss the topic of CC with some level of anger, according to pol_emo_mDeBERTa2, the AfD is particularly outraged, with over three-times as many sentences classified as containing angry language as the SPD and CDU/CSU. This results in over a third of the corpus marked for *anger*.

In a follow-up step, since a sentence may be classified to contain several emotions, we calculate the Pearson correlation coefficients and find only weak correlations ($-0.1 < \rho < 0.1$); this can be due to the fact that on sentence level there is not much room for placing multiple emotional words that moreover indicate different emotions. The correlation heatmaps for all three parties are supplied in the appendix.

Table 7 gives translated examples for sentences that were classified for the 4 most frequent emotions in the AfD-CCC: anger, fear, joy, and enthusiasm[12]. The anger, in this example, is directed at the state of Germany as a result of its legislation. Implicitly, current lawmakers are blamed for the shameful state Germany finds itself in, according to the speaker. Criticism of the "old parties" (*Altparteien*), being primarily SPD and CDU/CSU, is a core-aspect of AfD-rhetoric, also around the topic of climate change legislation (Sturm, 2020). This drives its criticism, and may thus be among the causes for its CC-rhetoric being significantly more angry than that of the parties it criticizes, which were part of the government for most (CDU/CSU) or the entirety (SPD) of the time frame that the dataset comprises.

If we consider *anger, fear, disgust*, and *sadness* to be negative, and the remaining four (*joy, enthusiasm, pride, hope*) to be positive emotions, we find the distribution of sentiment not entirely mirrored by that in Section 4.3. The trend, however, remains the same: The AfD displays a greater amount of negative emotion (and sentiment) than the comparison parties.

5 Conclusions

Summary. We have presented the AfD-CCC, a corpus of texts produced by members (for the most part, MP's) of the German right-wing populist party *Alternative für Deutschland*, addressing the topic of climate change. The AfD-CCC expands over five different text domains, a substantial part of which we make publicly available. The public corpus contains transcripts of political speeches in German and European parliaments, press releases, tweet IDs, and Telegram channel messages. The variability of the text sources allows for expansive studies of the party's CC-related rhetoric, both as it is communicated to the political body and to the public.

To showcase the applicability of the corpus, we have further presented three first use cases,

[12] A corresponding table with the original wording is supplied in Appendix A, Table 8.

emotion	subset	id	text
anger	German P.	sp19_703	"The way in which Germany is burying its powerful energy industry and the competitiveness of its companies borders on self-destruction, especially considering the fact that it imports solar panels and battery cells for electric cars from countries that understandably do not care about CO2 emissions."
fear	Telegram	645	"A country that is governed like that must (inevitably) end up in an emergency state."
joy	German P.	sp19_2450	"Boris Johnson achieved a landslide victory and the British have clearly voted against the EU."
enthusiasm	European P.	eu9_28287	"And we need an exit out of the Green Deal for safe energy and for social peace and for prosperity for us all."

Table 7: (Translated) examples for emotion classification with pol_emo_mDeBERTa2, taken from the AfD-CCC's publicly available subset. A table with the original wording is provided in the appendix.

whereby we compared the results on the AfD-CCC to similar datasets for the two other largest parties in the German parliament as of 2025, the SPD and CDU/CSU Union. A lexicon-based analysis using the populism dictionary by Gründl (2022) showed that the AfD's CC-related texts contain over five times as many matches to the dictionary as those of the other parties. This suggests a higher level of populist language in the AfD's rhetoric around CC. We confirmed this in a second small experiment with PopBERT (Erhard et al., 2025), where the AfD scored higher in anti-elite and host-right dimensions that comparison parties. Sentiment analysis using German-Sentiment-BERT (Guhr et al., 2020) and emotion detection using Pol_emo_mDeBERTa2 (Widmann and Wich, 2023) showed that, in cases where non-neutral sentiment or emotions were detected, the language was primarily negative. The AfD especially returned greater levels of *anger* than comparison parties, which, in turn, displayed higher proportions of positive emotions than the AfD.

Future work. As follow-up steps, we plan to study on the one hand the particular subtopics of CC that are being addressed by the AfD over time (also comparing the different communication channels), and on the other hand employ argument mining methods for detecting *claims* and *premises* (see, e.g., Lawrence and Reed (2019)) that shed more light on the argumentation strategies that are being employed.

Positive impact. Times are difficult for the climate movement, because their topic is not among the top of the agendas of societies these days. In many countries, both politicians and the public mood are currently preoccupied with other crises and problems. But on top of that, in many countries, climate-skeptic or -denialist parties have gained significant influence, and even if climate is not one of their top priorities either, they do actively exploit the issue – together with other ecological concerns – by framing it as an elitist project of people who lack connection with the "real problems" of the "real people". In this situation, which is to a large extent being shaped by social media communications but is also reflected in parliamentary debate, the CC movement has become aware that transmitting *facts* about causes and consequences of climate change will not be enough for changing the public mood to the better. Instead, it has been argued, disseminating positive narratives that offer constructive steps toward solutions can be more successful. For building such narratives, it is important to first be aware of the thinking and reasoning of "the other side". Influential right-wing parties, such as the AfD in Germany, are an important player there. Having access to their opinions and arguments, and using NLP to analyze their materials at scale, can support the monitoring of climate-skeptical discourse, and thereby help in reacting to changes in attitudes and in building counter-narratives. AfD-CCC is meant to contribute to this groundwork.

References

Daniel Bischof and Roman Senninger. 2018. Simple politics for the people? complexity in campaign messages and political knowledge. *European Journal of Political Research*, 57(2):473–495.

Bart Bonikowski and Noam Gidron. 2015. The Populist Style in American Politics: Presidential Campaign Discourse, 1952–1996. *Social Forces*, 94(4):1593–1621.

Christopher Cochrane, Ludovic Rheault, Jean-Francois Godbout, Tanya Whyte, Michael Wong, and Sophie Borwein. 2021. The automatic analysis of emotion in political speech based on transcripts. *Political Communication*, 39:1–24.

Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. 2021. Climate-fever: A dataset for verification of real-world climate claims. *arXiv:2012.00614*.

Lukas Erhard, Sara Hanke, Uwe Remer, Agnieszka Falenska, and Raphael Heiko Heiberger. 2025. PopBERT. Detecting Populism and Its Host Ideologies in the German Bundestag. *Political Analysis*, 3(1):1–17.

Francesca Grasso, Ronny Patz, and Manfred Stede. 2024. NYTAC-CC: A Climate Change Subcorpus based on New York Times Articles. In *Proceedings of CLiC-it - 10th Italian Conference on Computational Linguistics*, Pisa, Italy.

Johann Gründl. 2022. Populist ideas on social media: A dictionary-based measurement of populist communication. *New Media & Society*, 24(6):1481–1499.

Oliver Guhr, Anne-Kathrin Schumann, Frank Bahrmann, and Hans Joachim Böhme. 2020. Training a broad-coverage german sentiment classification model for dialog systems. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 1620–1625, Marseille, France. European Language Resources Association.

Christopher Klamm, Ines Rehbein, and Simone Paolo Ponzetto. 2023. Our kind of people? detecting populist references in political debates. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1227–1243, Dubrovnik, Croatia. Association for Computational Linguistics.

Anne Küppers. 2022. 'Climate-Soviets,' 'Alarmism,' and 'Eco-Dictatorship': The Framing of Climate Change Scepticism by the Populist Radical Right Alternative for Germany. *German Politics*, 33(1):1–21.

Jana Lasser, Segun Taofeek Aroyehun, Almog Simchon, Fabio Carrella, David Garcia, and Stephan Lewandowsky. 2022. Social media sharing of low-quality news sources by political elites. *PNAS Nexus*, 1(4).

John Lawrence and Chris Reed. 2019. Argument mining: A survey. *Computational Linguistics*, 45(4):765–818.

Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.

Cas Mudde. 2004. The populist zeitgeist. *Government and Opposition*, 39(4):541–563.

Florian Primig and Fabian Fröschl. 2024. Introducing the FROG tool for gathering Telegram data. *Mobile Media & Communication*, 12(2):449–453.

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

Matthijs Rooduijn and Teun Pauwels. 2011. Measuring populism: Comparing two methods of content analysis. *West European Politics*, 34(6):1272–1283.

Robin Schaefer, Christoph M. Abels, Stephan Lewandowsky, and Manfred Stede. 2023. Communicating Climate Change: A Comparison Between Tweets and Speeches by German Members of Parliament. In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, Toronto, Canada and Online. Association for Computational Linguistics.

Stella Schaller and Alexander Carius. 2019. *Convenient Truths: Mapping climate agendas of right-wing populist parties in Europe*. adelphi, Berlin.

Jan Schwalbach, Lukas Hetzer, Sven-Oliver Proksch, Christian Rauh, and Miklós Sebők. 2025. ParlLawSpeech. GESIS, Cologne.

Shuvam Shiwakoti, Surendrabikram Thapa, Kritesh Rauniyar, Akshyat Shah, Aashish Bhandari, and Usman Naseem. 2024. Analyzing the dynamics of climate change discourse on Twitter: A new annotated corpus and multi-aspect classification. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 984–994, Torino, Italia. ELRA and ICCL.

Manfred Stede and Ronny Patz. 2021. The climate change debate and natural language processing. In *Proceedings of the 1st Workshop on NLP for Positive Impact*, pages 8–18, Online. Association for Computational Linguistics.

Georg Sturm. 2020. Populismus und Klimaschutz. der AfD-Klimadiskurs. *Soziologiemagazin*, 13(2):69–92.

Enrica Troiano, Sebastian Padó, and Roman Klinger. 2019. Crowdsourcing and validating event-focused emotion corpora for German and English. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4005–4011, Florence, Italy. Association for Computational Linguistics.

Elena Volkanovska, Sherry Tan, Changxu Duan, Sabine Bartsch, and Wolfgang Stille. 2025. The InsightsNet Climate Change Corpus (ICCC). *Datenbank Spektrum*, 23:177–188.

Tobias Widmann and Maximilian Wich. 2023. Creating and comparing dictionary, word embedding, and transformer-based models to measure discrete emotions in german political text. *Political Analysis*, 31(4):626–641.

Sergio E. Zanotto, Diego Frassinelli, and Miriam Butt. 2024. Language complexity in populist rhetoric. In *Proceedings of the 4th Workshop on Computational Linguistics for the Political and Social Sciences: Long and short papers*, pages 61–80, Vienna, Austria. Association for Computational Linguistics.

A Appendix

A.1 Correlation Heatmaps

The following Figures 2 and 3, display heatmaps of the Pearson correlation coefficients for the 8 emotion categories for each party, as outlined in Section 4.4.

A.2 Emotion Examples

Table 8 contains the original German wording of the translated examples in Table 7, Section 4.4.

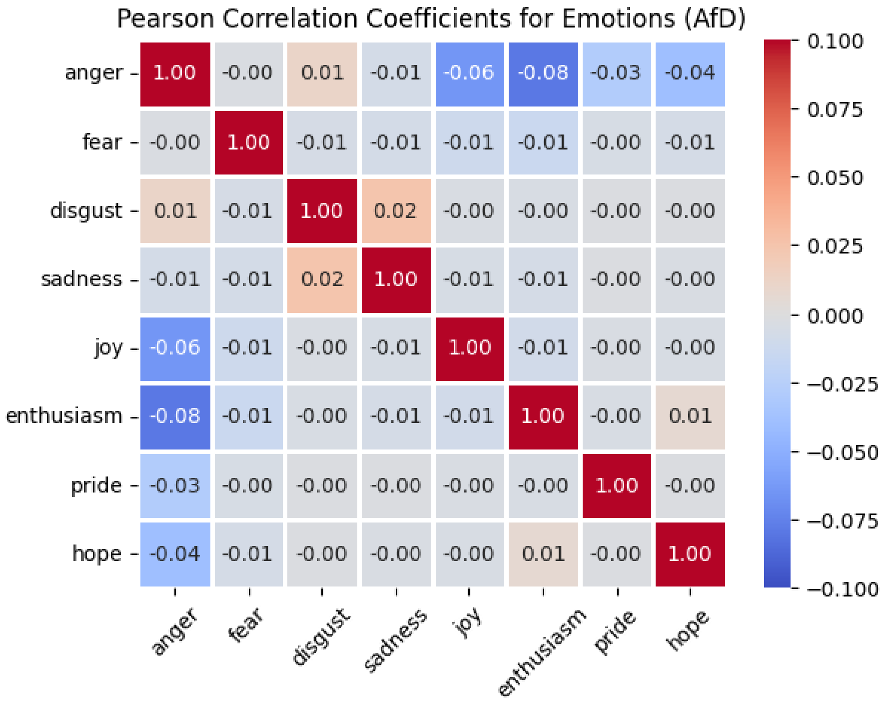

Figure 2: Pearson correlation coefficients for the emotions in the ed8 emotion dictionary, as found in the AfD-CCC.

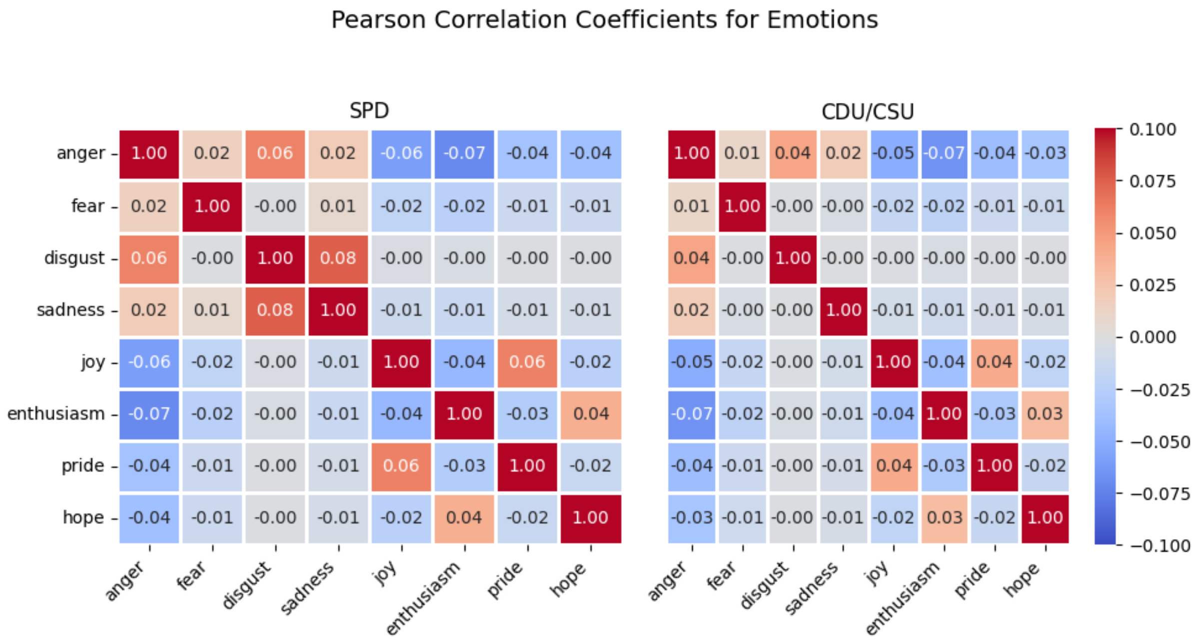

Figure 3: Pearson correlation coefficients for the emotions in the ed8 emotion dictionary, as found in the comparison datasets for SPD (left) and CDU/CSU (right).

emotion	subset	id	text
anger	German P.	sp19_703	"Es grenzt schon an Selbstzerstörung, wie Deutschland seine leistungsfähige Energiewirtschaft und die Wettbewerbsfähigkeit seiner Unternehmen zu Grabe trägt, insbesondere im Hinblick darauf, dass es Sonnenkollektoren und Batteriezellen für Elektroautos genau aus den Ländern importiert, die sich nachvollziehbar nicht um CO2-Emissionen scheren."
fear	Telegram	645	"Ein Land, das so regiert wird, muss in eine Notlage kommen."
joy	German P.	sp19_2450	"Boris Johnson hat einen Erdrutschsieg erzielt, und die Briten haben sich damit ganz klar gegen die EU entschieden."
enthusiasm	European P.	eu9_28287	"Und wir brauchen den Ausstieg aus dem Grünen Deal für sichere Energie und für sozialen Frieden und für Wohlstand für uns alle."

Table 8: Examples for emotion classification with pol_emo_mDeBERTa2, taken from the AfD-CCC's publicly available subset, in their original German wording.