Analyzing the Online Communication of Environmental Movement Organizations: NLP Approaches to Topics, Sentiment, and Emotions

Christina S. Barz
Darmstadt University
of Applied Sciences
Haardtring 100
64295 Darmstadt
Germany
christina.barz@h-da.de

Melanie Siegel
Darmstadt University
of Applied Sciences
Schöfferstraße 3
64295 Darmstadt
Germany
melanie.siegel@h-da.de

Daniel Hanss
Darmstadt University
of Applied Sciences
Haardtring 100
64295 Darmstadt
Germany
daniel.hanss@h-da.de

Abstract

This project employs state-of-the-art Natural Language Processing (NLP) techniques to analyze the online communication of international Environmental Movement Organizations (EMOs). First, we introduce our overall EMO dataset and describe it through topic modeling. Second, we evaluate current sentiment and emotion classification models for our specific dataset. Third, as we are currently in our annotation process, we evaluate our current progress and issues to determine the most effective approach for creating a high-quality annotated dataset that captures the nuances of EMO communication. Finally, we emphasize the need for domain-specific datasets and tailored NLP tools and suggest refinements for our annotation process moving forward.

1 Introduction

In order to address the escalating environmental crises of our time, it is imperative that individuals and groups worldwide act in a collective manner. Investigating current environmental movements is crucial as they play a significant role in motivating collective environmental action, shaping public opinion, and influencing policy decisions. The online communication of Environmental Movement Organizations (EMOs) provides valuable insights into their strategic approaches, thematic content and emotional appeals (Gulliver et al., 2021; Ackland and O'Neil, 2011). Concurrently, the success of these organizations can be ascertained through the examination of various reactions of other users, such as likes and comments. Therefore, the objective of our project is to analyze four years of X (former Twitter) data, which we refer to as tweets, from a range of international EMOs, including *Greenpeace*, *Friends of the Earth*, *Fridays for Future*, *Extinction Rebellion*, and *Climate Action Network (CAN)* in order to gain a deeper understanding of their communication. We intend to assess the sentiment and emotions conveyed in the EMOs' language, utilizing and evaluating state-of-the-art Natural Language Processing (NLP) models. As part of this research, our ultimate goal is to create a comprehensive and annotated dataset tailored to climate- and environment-specific content, in the future. In this paper, we present an interim stage of our project, focusing on a critical assessment of our current annotation process to refine and enhance its robustness, with the goal of creating a high-quality annotated dataset. The paper is structured as follows: First, we review related work and outline our research questions. Next, we describe our methodology, including the annotation process and the evaluation of existing sentiment and emotion models. We then present our results, highlighting key findings and insights gained from this evaluation. Finally, we conclude by discussing the limitations of our approach and proposing avenues for future research, including further dataset development and model fine-tuning for climate- and environment-related communication.

1.1 Related Work & Research Questions

Sentiment and emotion analysis is rarely applied to climate and environmental contexts. Instead, most research in this field focuses on more general applications, such as product analysis or the determination of stock market trends (Wankhade et al., 2022). Moreover, there is a paucity of datasets and models that have been specifically designed for the purpose of understanding climate-related text in social media. Available datasets include the *ClimaConvo* dataset, which comprises 15,309 tweets from the year 2022 labeled as *relevance*, *stance*, *hate speech*, *direction of hate*, and *humor* (Shiwakoti et al., 2024) and the *Twitter Climate Change Sentiment Dataset* (Qian, 2021) with a to-

Proceedings of the 1st Workshop on Ecology, Environment, and Natural Language Processing (NLP4Ecology2025), pages 68–76
March 2, 2025

tal of 43,943 tweets spanning from 2015 to 2018. Here each tweet was classified into one of four categories: *news*, *pro*, *neutral*, or *anti*.
A significant challenge in the research field is posed by the ambiguity of sentiment and emotions, particularly in social media, where context and tone vary greatly (Pozzi et al., 2016). In contrast to studies that analyze individual users' posts (Dahal et al., 2019; El Barachi et al., 2021), our research exclusively examines content created by groups. It seems probable that the content of posts from EMOs reflects strategic communication approaches pursued by an organization rather than an expression of individual members' sentiments or emotional states. This raises the question of whether state-of-the-art NLP models are capable of adequately capturing the nuances of environmental and climate communication from EMOs. Our study addresses these gaps in the literature by investigating the following research questions: (1) *How effectively do state-of-the-art NLP models perform in analyzing online communication of EMOs?* and (2) *How effective and reliable is our current annotation process in capturing sentiment and emotion in climate- and environment-related tweets and what refinements are necessary to ensure the creation of a high-quality, domain-specific annotated dataset?*
To this end, we first employ topic modeling to describe our dataset and then test the performance of the ClimateBERT sentiment model for sentiment analysis (Webersinke et al., 2021) and the emotion model *bhadresh-savani/bert-base-uncased-emotion* for emotion classification (Savani, 2020). Our analyses regarding the second research question should provide insights for addressing challenges that may arise during our annotation process, such as inter-annotator agreement (IAA), and enhance the reliability of future annotations.

EMO	Number of Documents
Greenpeace	14420
Extinction Rebellion	12004
CAN	5152
Fridays for Future	2353
Friends of the Earth	2230

Table 1: Document Distribution in Dataset

2 Methodology

2.1 Data

The dataset extracted in September 2024 comprises 36,159 tweets from five prominent international EMOs, namely *Greenpeace*, *Extinction Rebellion*, *Friends of the Earth*, *Fridays for Future*, and *CAN*, see Table 1. The tweets were published between 2019 and 2024. The dataset comprises the following information for each document: group name, time, retweet count, reply count, like count and tweet text. All analyses were performed in Python (version 3.11.11) using the *bertopic*, *pandas*, and *Scikit-learn* packages (pandas development team, 2020; Pedregosa et al., 2011; Grootendorst, 2020).

2.2 Annotation Process

A first sub-dataset of 1399 tweets was independently annotated by three annotators. To facilitate this process, an annotation guideline was developed, which provided clear definitions for all constructs and illustrative examples, see osf.io for our guidelines. Annotators were instructed to label sentiment and expressions of the emotions joy, anger, fear, and sadness. The annotation process commenced with a preliminary phase, during which the annotators labelled an initial set of 10 tweets. This was followed by a feedback session, during which ambiguities were addressed and alignment on the labeling criteria was ensured. Subsequently, feedback sessions were conducted at regular intervals, with each session focusing on a specific subset of 500 tweets. As discrepancies are typical in annotation tasks of this nature (Uma et al., 2021), we established a gold standard dataset through majority voting. In instances where all three annotators reached a differing conclusion (this could only occur with sentiment annotations), these cases were subjected to further analysis and resolution through group discussions in order to achieve a consensus. This approach ensured a balance between individual judgments and collaborative decision-making, thereby enhancing the reliability of the annotations.

2.3 Topic Modeling

To describe and analyze the themes present within the whole dataset (36,159), we employed BERTopic, a recently developed topic modeling approach that utilizes embedding-based techniques in contrast to the traditional bag-of-words methods, such as LDA (Jelodar et al., 2019; Grootendorst, 2022). BERTopic uses semantic embeddings to cluster documents and utilizes parameters such as *n_neighbors* and *min_cluster_size* to refine the granularity of topics. In contrast to LDA, BERTopic does not necessitate the pre-definition of the number of topics. Unlike other approaches, the model is designed to determine the number of clusters based on the data itself. A more detailed representation of themes was given priority, which informed our selection of parameters. The complete parameter settings are available for consultation at osf.io.

2.4 Sentiment Classification

In order to conduct a sentiment analysis, we utilized the ClimateBERT sentiment model, which has been specifically trained on texts pertaining to climate-related issues (Webersinke et al., 2021). In their model, the researchers conceptualized sentiment as a framing, categorizing climate-related text as either positive (opportunity), neutral, or negative (risk). It is noteworthy that ClimateBERT was trained on longer documents, such as news articles or financial reports, and its performance on shorter social media texts remains untested. The present study assesses the applicability of this approach to short-form content, such as tweets.

2.5 Emotion Classification

Emotion analysis was conducted using the emotion model *bhadresh-savani/bert-base-uncased-emotion* that had been trained on general tweets (Savani, 2020). The model identifies a number of emotions, from which we selected four for our analysis that overlap with our annotations, i.e. *joy*, *anger*, *sadness*, and *fear*. Despite having been trained on a generic social media dataset, such models typically necessitate fine-tuning for specific domains. Nonetheless, prior to the fine-tuning of a model, a preliminary evaluation is conducted to assess the functionality of existing models. Therefore, in this analysis, the model's capacity to categorize emotions within the context of climate change and environmental issues is assessed without additional fine-tuning.

3 Results

3.1 IAA & Class Distributions

In order to assess the quality of our gold standard, we have calculated the Fleiss' Kappa coefficient, see Table 2, for our sub-data set (1,399 tweets) and examined the class distributions (Fleiss, 1971). We had slight to moderate Kappa depending on sentiment or the specific emotion. Despite the provision of guidelines and feedback sessions, there was a notable discrepancy in the interpretation of sentiment and emotions by the annotators. The feedback conversations revealed that the annotations contained a bias toward personal emotional reactions to the text. This means that annotators tended to label tweets with emotions if the tweets evoked certain emotions in them. For example, neutral texts reporting on extreme weather events were often rated with sadness, fear or anger, even though the texts were written without emotional tone. We found the least agreement for the emotions anger (0.2472) and fear (0.0379).

Construct	Fleiss' Kappa
Sentiment	0.4574
Joy	0.4708
Anger	0.2472
Fear	0.0379
Sadness	0.1825

Table 2: Inter-Annotator Agreement (IAA) measured by Fleiss' Kappa

Class	Label	Count
Sentiment		
Neutral	0	1054
Risk	-1	325
Opportunity	1	20
Joy		
No Joy	0	1378
Joy	1	21
Anger		
No Anger	0	1298
Anger	1	101
Fear		
No Fear	0	1389
Fear	1	10
Sadness		
No Sadness	0	1391
Sadness	1	8

Table 3: Class Distribution for Sentiment and Emotion Labels

Nevertheless, we have created a majority voting gold standard to evaluate the current models. In 14 cases, a lack of consensus was observed among the three annotators, necessitating a collective dis-

Topic	Topic Name	Number of Documents
0	Climate Change, Fossil Fuels & Finance	22711
1	Indigenous People & Biodiversity	1387
2	Ocean	1231
3	Plastic	1049
4	Noise	1039
5	Indigenous People, Brazil & Amazon	751
6	Noise: Posts in other Languages	600
7	Women & Gender	623
8	Forest & Deforestation	620
9	Support XR Groups & Activists	460
10	Food & Agriculture	607
11	Australia & Wildfires	530
12	Stop Shell	480
13	Air Pollution	385
14	Meat & Dairy	329
15	Deep Sea Mining	295
16	Ban Private Jets	272
17	Nuclear Energy & War	255
18	Transport & Mobility	318
19	Indonesia & Palm Oil	430
20	Policing Bill	321
21	Palestine	269
22	Noise: Apply for Climate Jobs	186
23	Fossil of the Day Award	200
24	Vaccine & Covid19	142
25	Black Friday & Buying	168
26	Cars & Vehicles	152
27	Countries	232
28	Human Rights Act	117

Table 4: Identified Topics, Labels, and Frequencies from Initial Topic Modeling

cussion to resolve the discrepancy. The class distribution in our gold standard is very unbalanced, see Table 3. For example, in our annotations we have more labels for climate change as risk (23.23 %; 325 posts) compared to opportunity (1.43 %; 20 posts), which suggests that EMOs view climate change as a high risk. In terms of emotions, we only had 1.50 % joy (21 posts), 0.71 % fear (10 posts) and 0.57 % sadness (8 posts), compared to a higher incidence of anger with 7.22% (101 posts). The distribution of emotional language used by EMOs indicates a lack of emotional expression, with anger being the most prevalent emotion.

3.2 Topic Modeling

Our initial topic analysis yielded 29 topics, which were then subjected to a manual review by one researcher. Three topics consisting solely of documents labeled as *Noise* due to their lack of meaningful content or content not in English (e.g., 'Clearly.', 'Hmm.' or 'Starting in about 1 hour! Make sure to tune in!'), were excluded from further analysis. For each remaining topic, an in-depth analysis of the representative documents and word representations was conducted, which resulted in the ten most frequently discussed topics: *'Climate Change, Fossil Fuels & Finance'*, *'Indigenous People & Biodiversity'*, *'Ocean'*, *'Plastic'*, *'Indigenous People, Brazil & Amazon'*, *'Women & Gender'*, *'Forest & Deforestation'*, *'Support XR Groups & Activists'*, *'Food & Agriculture'* and *Australia & Wildfires*. For a comprehensive list of all 29 topics and their corresponding labels, please refer to Table 4.

A thorough examination of the most frequently occurring topics across EMOs reveals distinct pat-

Author	Topic	Description	Frequency
CAN	0	Climate Change, Fossil Fuels & Finance	4376
CAN	23	Fossil of the Day Award	142
CAN	1	Indigenous People & Biodiversity	114
CAN	21	Palestine	64
CAN	7	Women & Gender	50
Extinction Rebellion	0	Climate Change, Fossil Fuels & Finance	8323
Extinction Rebellion	9	Support XR Groups & Activists	450
Extinction Rebellion	1	Indigenous People & Biodiversity	293
Extinction Rebellion	11	Australia & Wildfires	272
Extinction Rebellion	20	Policing Bill	222
Friends of the Earth	0	Climate Change, Fossil Fuels & Finance	977
Friends of the Earth	1	Indigenous People & Biodiversity	352
Friends of the Earth	10	Food & Agriculture	179
Friends of the Earth	7	Women & Gender	126
Friends of the Earth	21	Palestine	98
Fridays for Future	0	Climate Change, Fossil Fuels & Finance	1693
Fridays for Future	27	Activism in diverse Countries	107
Fridays for Future	1	Indigenous People & Biodiversity	102
Fridays for Future	21	Palestine	59
Fridays for Future	20	Policing Bill	42
Greenpeace	0	Climate Change, Fossil Fuels & Finance	7342
Greenpeace	2	Ocean	1067
Greenpeace	3	Plastic	829
Greenpeace	5	Indigenous People, Brazil & Amazon	532
Greenpeace	1	Indigenous People & Biodiversity	526

Table 5: Distribution of Topics by Author After Initial Topic Modeling Analysis

terns that reflect the issues these groups prioritize and the strategies they employ. For instance, both Extinction Rebellion and Fridays for Future have most frequent topics which are activism related, such as *'Support XR Groups & Activists'* for Extinction Rebellion and *'Activism in diverse Countries'* for Fridays for Future. In addition, both groups have the topic of *'Policing Bill'* in their most common themes, which includes restrictions on unacceptable protest behavior. These subjects, which have been derived from the topic modeling, reflect the identity and strategies of the groups, as Extinction Rebellion and Fridays For Future are more akin to a protest movement in comparison to larger EMOs such as Greenpeace and Friends of the Earth. Furthermore, Greenpeace appears to prioritize subjects such as *'Ocean'* and *'Plastic'*, in contrast to other groups. It should also be noted that the topic of *'Women & Gender'* only appeared frequently at the CAN and Friends of the Earth. For all topic frequencies and representative documents refer to Table 5.

Since the most frequent topic *'Climate Change, Fossil Fuels & Finance'* encompassed the majority of the documents, we conducted another round of topic modeling using only the documents from this topic (22,711) to explore its content in more detail. This analysis revealed several specific subtopics, as shown in Table 6. Further breakdown of these topics by organization provided valuable insights, see Table 7. For example, CAN primarily posts about COP (Conference of the Parties) and financial issues, while Greenpeace frequently communicates about fossil fuels. Extinction Rebellion focuses heavily on peaceful protest and rights, emphasizing advocacy and activism in its messaging. Fridays for Future, on the other hand, focuses almost exclusively on activism-related issues. Their communication strategy is particularly inviting and action-oriented, as reflected in common themes such as *'Join Fridays for Future Strike'* and *'Friendly Reminder to Act Now'*. These findings underscore the different thematic focuses and strategic communication approaches of each organization, shedding

Topic	Topic Name	Number of Documents
0	COP, Loss and Damage & Finance	3009
1	Fossil Fuels	2631
2	Noise: Article, Link, Source, Join & Share	1420
3	Climate Emergency & Denial	1029
4	Peaceful Protest & Protest Rights	953
5	Carbon Emissions & Net Zero	840
6	Fight for Freedom, Peaceful & Just World	812
7	Flood	760
8	Nature & Sustainable Future	760
9	Join Fridays for Future Strike	663
10	Climate Justice & Court	622
11	Climate Crisis Solutions	613
12	Heat	561
13	Covid19	539
14	Friendly Reminder to Act Now	495
15	Environmental Crisis	494
16	Economic Growth	479
17	Climate, Gender & Racial Justice	472
18	Global Warming, Climate Breakdown & Extreme Weather	443
19	Activism Works	433
20	ISDS	403
21	Greenpeace	379
22	Africa & Energy	371
23	Renewable Energy	369
24	(Youth) Climate Activists	342
25	Coal Mine	324
26	Hope & Love	303
27	Extreme Weather Events	277
28	Rebellion & Resistance	257
29	2021 Session of the UNFCCC Subsidiary Bodies	221
30	IPCC	200
31	Ice & Glacier Melting	194
32	Anxiety, Grief & Hope	193
33	Extinction Rebellion	191
34	Philippines & Typhoons	180
35	Vanuatu & Pacific Islands	166
36	Norway, Denmark, Oil & Coal	160
37	Citizens Assemblies	153

Table 6: Identified Topics, Labels, and Frequencies from Second Topic Modeling

light on their priorities and methods of engagement.

3.3 Sentiment Classification

We tested the application of the model on our gold standard. Using the ClimateBERT sentiment model, we achieved an F1 score of 0.4333 (Precision = 0.6504, Recall = 0.3283). This result can be explained by the training data set of the ClimateBERT sentiment model, which consists of longer documents such as financial reports (Webersinke et al., 2021). We conclude that the application to social media posts is not possible without limitations. According to the results, the model should be fine tuned with social media data before it is applied.

Author	Topic	Description	Frequency
CAN	0	COP, Loss and Damage & Finance	2034
CAN	1	Fossil Fuels	437
CAN	29	2021 Session of the UNFCCC Subsidiary Bodies	158
CAN	10	Climate Justice & Court	149
CAN	5	Carbon Emissions & Net Zero	148
Extinction Rebellion	1	Fossil Fuels	883
Extinction Rebellion	4	Peaceful Protest & Protest Rights	707
Extinction Rebellion	6	Fight for Freedom, Peaceful & Just World	447
Extinction Rebellion	7	Flood	417
Extinction Rebellion	0	COP, Loss and Damage & Finance	394
Friends of the Earth	0	COP, Loss and Damage & Finance	199
Friends of the Earth	5	Carbon Emissions & Net Zero	119
Friends of the Earth	1	Fossil Fuels	79
Friends of the Earth	10	Climate Justice & Court	64
Friends of the Earth	22	Africa & Energy	53
Fridays for Future	9	Join Fridays for Future Strike	471
Fridays for Future	1	Fossil Fuels	100
Fridays for Future	24	(Youth) Climate Activists	80
Fridays for Future	14	Friendly Reminder to Act Now	74
Fridays for Future	4	Peaceful Protest & Protest Rights	65
Greenpeace	1	Fossil Fuels	1132
Greenpeace	3	Climate Emergency & Denial	477
Greenpeace	8	Nature & Sustainable Future	320
Greenpeace	0	COP, Loss and Damage & Finance	318
Greenpeace	21	Greenpeace	308

Table 7: Distribution of Topics by Author After Second Topic Modeling Analysis

3.4 Emotion Classification

We tested the application of the emotion model *bhadresh-savani/bert-base-uncased-emotion* on our gold standard. Based on the model's prediction, continuous emotion outputs were generated for each document, such as 0.45959. In two separate analyses, we applied thresholds of 0.2 and 0.5 to these outputs to compare them to our gold standard. Since we categorized emotions as either present (1) or absent (0), regardless of their intensity, values between 0.2 and 1, or 0.5 and 1, were considered indicative of the presence of an emotion. These two thresholds were used to examine whether the choice of threshold influenced the model's performance. The performance metrics for both thresholds are presented in the corresponding Tables 8 and 9. The choice of cutoff only had a minimal effect on performance, with the 0.5 cutoff showing a slight improvement. F1 scores ranged from 0.0270 to 0.1847 for the 0.2 cutoff and from 0.0496 to 0.2302 for the 0.5 cutoff. However, we conclude that the overall performance remained inadequate and unsuitable for practical use in analyzing environmental and climate-related texts. Given the obtained F1 scores, fine-tuning the model for climate and environmental contexts may prove challenging. Therefore, the use of alternative or more advanced models, such as Large Language Models, may be necessary to improve performance.

Emotion	Precision	Recall	F1 Score
Joy	0.0278	0.9048	0.0540
Anger	0.1081	0.6337	0.1847
Fear	0.0142	0.3000	0.0270
Sadness	0.0174	0.6250	0.0339

Table 8: Model Performance with 0.2 Cutoff

4 Limitations

This study has several limitations that should be considered when interpreting the results. First, only 1,399 tweets were used as the gold standard for model evaluation, which may limit the generaliz-

Emotion	Precision	Recall	F1 Score
Joy	0.0335	0.9048	0.0645
Anger	0.1422	0.6040	0.2302
Fear	0.0270	0.3000	0.0496
Sadness	0.0296	0.6250	0.0565

Table 9: Model Performance with 0.5 Cutoff

ability of our findings. These tweets were annotated by three annotators, with the final dataset created using majority voting. While this approach is standard, the IAA was only slight to moderate (ranging from 0.1825 to 0.4708), which complicates the evaluation of the models. Disagreements among annotators, especially for emotions like anger (IAA = 0.2472) and fear (IAA = 0.0379), are not unusual but highlight the subjective nature of the task. Annotators may interpret climate- and environment-related texts in different ways, given their complexity and the challenge of reading such texts neutrally. We question whether an unbiased annotation of such texts is possible, since the various climate and environmental issues addressed in the documents are difficult to read neutrally. We attribute some of the disagreement in sentiment to the possibility of multiple sentiment framings within a single post. For example, a tweet may present both a risk and an opportunity framing, requiring annotators to choose a single sentiment, which can lead to varied interpretations. This additional room for interpretation may explain some of the discrepancies in sentiment. We would like to emphasize that our previous annotations have primarily shown that texts related to climate and environmental issues seem to be difficult to interpret and evaluate, which crystallizes them as a very challenging area in NLP where there still seems to be a need for research, annotation, and training.

Second, the highly imbalanced class distributions in the dataset pose a significant challenge for evaluating model performance. The presence of floor effects further complicates the accuracy of contemporary sentiment and emotion models, making it difficult to assess their full potential.

Third, our analysis was limited to tweets, which may not fully capture the broader communication patterns of EMOs across different social media platforms. Additionally, this study did not account for the impact of Elon Musk's acquisition of Twitter in October 2022, which led to significant changes to the platform's structure and policies. These changes could have influenced the communication strategies of EMOs in ways that our dataset does not reflect, thus limiting the scope of our findings. Lastly, the use of topic modeling tools, such as BERTopic, also has limitations. While helpful in organizing large datasets, such tools are not infallible. They may fail to identify certain topics or assign topics inaccurately, which could impact the interpretation of the results.

5 Conclusion and Future Work

Despite these limitations, our findings provide valuable insights into the strategic communication of EMOs and the challenges associated with annotating data as well as applying current NLP models to climate- and environment-related group discourse. Our study underscores the need for handling disagreement in data annotation, domain-specific datasets, and models to address the unique challenges posed by analyzing climate- and environment-related content. Our preliminary evaluation of the annotation process serves as a crucial step towards refining and enhancing its robustness. It is imperative that future efforts dedicate greater attention to the resolution of disagreement in climate- and environment-related text annotations. Overall, future research should prioritize the development of robust domain-specific datasets and the fine-tuning of models to improve accuracy and interpretability. Additionally, exploring other psychological constructs, such as efficacy beliefs, alongside traditional sentiment and emotion analysis, should provide a more comprehensive understanding of online climate and environment communication. Expanding beyond the current focus on sentiment, emotion, and hate speech to include such constructs can yield a richer and more nuanced perspective on EMO strategies and their impact on public discourse.

References

Robert Ackland and Mathieu O'Neil. 2011. Online collective identity: The case of the environmental movement. *Social Networks*, 33(3):177–190.

Biraj Dahal, Sathish AP Kumar, and Zhenlong Li. 2019. Topic modeling and sentiment analysis of global climate change tweets. *Social network analysis and mining*, 9:1–20.

May El Barachi, Manar AlKhatib, Sujith Mathew, and Farhad Oroumchian. 2021. A novel sentiment analysis framework for monitoring the evolving public opinion in real-time: Case study on climate change. *Journal of Cleaner Production*, 312:127820.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Maarten Grootendorst. 2020. Bertopic: A topic modeling technique using bert embeddings. Accessed: 2024-12-30.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

Robyn Gulliver, Kelly S Fielding, and Winnifred R Louis. 2021. Assessing the mobilization potential of environmental advocacy communication. *Journal of Environmental Psychology*, 74:101563.

Hamed Jelodar, Yongli Wang, Chi Yuan, Xia Feng, Xiahui Jiang, Yanchao Li, and Liang Zhao. 2019. Latent dirichlet allocation (lda) and topic modeling: models, applications, a survey. *Multimedia tools and applications*, 78:15169–15211.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.

Federico Alberto Pozzi, Elisabetta Fersini, Enza Messina, and Bing Liu. 2016. *Sentiment analysis in social networks*. Morgan Kaufmann.

Ed Qian. 2021. Twitter climate change sentiment dataset. Accessed: 2024-12-30.

Bhadresh Savani. 2020. Bert-base-uncased emotion. Accessed: 2024-12-30.

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.

The pandas development team. 2020. pandas-dev/pandas: Pandas. *Zenodo*.

Alexandra N Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.

Mayur Wankhade, Annavarapu Chandra Sekhara Rao, and Chaitanya Kulkarni. 2022. A survey on sentiment analysis methods, applications, and challenges. *Artificial Intelligence Review*, 55(7):5731–5780.

Nicolas Webersinke, Mathias Kraus, Julia Anna Bingler, and Markus Leippold. 2021. Climatebert: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010*.