Climate Change Discourse Over Time: A Topic-Sentiment Perspective

Chaya Liebeskind[a,*,1] **and Barbara Lewandowska-Tomaszczyk**[b,1]

[a]Department of Computer Science, Jerusalem College of Technology, Israel
[b]Department of Language and Communication, University of Applied Sciences in Konin, Poland
ORCID (Chaya Liebeskind): https://orcid.org/0000-0003-0476-3796, ORCID (
Barbara Lewandowska-Tomaszczyk): https://orcid.org/0000-0002-6836-3321

Abstract. The present paper focuses on the study of opinion dynamics and opinion shifts in social media in the context of climate change discourse in terms of the quantitative NLP analysis, supported by a linguistic outlook. The research draws on two comparable collections of climate-related social media data from different time periods, each based on trending climate-related hashtags and annotated for relevant sentiment values. The quantitative computer-based research methodology has been supported by a language-based perspective in the pragma-linguistic form. The research shows that the latter data source, for the majority of identified topics, exhibits a significant reduction in negative sentiment and a dominance of positive sentiment, i.e., a potential temporal evolution in public sentiment toward climate change. To achieve this, we used a BERT-based clustering approach to identify dominant themes within a combined dataset of tweets from both periods. Subsequently, a unified sentiment classification framework using a Large Language Model (LLM) was applied to reclassify all tweets, ensuring consistent and climate-specific sentiment analysis across both datasets. This methodology allowed for a coherent comparison of public attitudes and their evolution in different time periods and thematic structures.

1 Introduction

Understanding human communication in its full richness requires delving beyond the mere factual content of utterances. It requires an exploration of the underlying layers of opinion, sentiment, and attitude [10] that imbue discourse with meaning, intention, and interpersonal resonance. These three interconnected yet distinct concepts form the background of subjective expression, shaping how individuals perceive, interpret, and react to the world around them. The manifestation and interplay of opinions, sentiments, and attitudes are ubiquitous, fundamentally influencing the dynamics of human interaction and the propagation of ideas.

The study of opinion, sentiment, and attitude in discourse has emerged as an increasingly critical area within diverse academic disciplines, including linguistics, computational linguistics, and others. This interdisciplinary focus reflects the impact these objective and subjective elements have on information processing, decision-making, and social cohesion.

Opinion can be broadly defined as a belief or judgment held by an individual about a particular subject. It represents a cognitive stance, intermixed with existential, epistemological and moral stances [11], perceived in terms of propositional structures, or else as proposed in Lewandowska-Tomaszczyk et al. [12] perceived in terms of a whole Opinion Event.

The topic of this study undertakes research in the dynamicity of opinions. Opinions are dynamic and can evolve over time as individuals and communities acquire new information, engage in critical reflection, or are exposed to diverse perspectives. In discourse, opinions are frequently expressed through explicit statements, but can also be inferred from argumentative structures, rhetorical devices, patterns of reasoning in terms of figurative, indirect or implicit structures [13].

Sentiment refers to the emotional tone or feeling conveyed by a piece of text or speech. It encapsulates the affective dimension of discourse, reflecting positive, negative, or neutral emotional or, more generally, affective, states [5]. For example, saying "I love this new policy" expresses a positive sentiment, while "I'm frustrated with the current situation" conveys a negative one. Sentiment is often, though not always, expressed through emotionally charged vocabulary (e.g., "amazing," "terrible," "joy," "anger"), but can also be conveyed through intonation, facial expressions, or linguistic cues like intensifiers or hedges. The granularity of sentiment analysis can vary, from broad polarity (positive/negative/neutral) to more fine-grained emotions (e.g., joy, sadness, anger, fear, surprise, disgust) and affective states such as curosity, reserve, boredom [5]. The prevalence of sentiment in online reviews, social media posts, and customer feedback has made sentiment analysis a cornerstone of natural language processing (NLP), enabling businesses to gauge public perception, track brand reputation, and identify emerging trends.

Attitude represents a more stable and enduring predisposition towards a person, object, idea, or issue [3]. It is a psychological construct that reflects an individual's overall evaluative stance, encompassing cognitive, affective, and behavioral components. Attitudes are often seen as the underlying drivers of opinions and sentiments. For example, a pro-environmental attitude might lead to the opinion that renewable energy is essential and positive sentiments toward green initiatives. Attitudes are more deeply ingrained than sentiments or even situation-specific opinions; they shape an individual's worldview and influence their long-term behaviors and decision-making. In discourse, attitudes are often less explicitly stated than opinions

* Corresponding Author. Email: liebchaya@gmail.com
[1] Equal contribution.

or sentiments, but can be inferred from consistent patterns of expression, recurring themes, and the overall evaluative orientation of an individual's communication. Analyzing attitudes in discourse involves a more holistic and interpretive approach, often drawing on psychological theories and qualitative methodologies.

The relationship between these three concepts is hierarchical and interdependent. Attitudes form the broadest and most stable foundation, influencing the opinions an individual holds, which in turn are often expressed with a particular sentiment. For example, a deeply ingrained attitude of skepticism toward government intervention might lead to the opinion that a new social welfare program is flawed, which could then be articulated with sentiment of anger or frustration. However, this relationship is not always linear or one-directional. A particularly strong sentiment or a newly formed opinion, especially when reinforced by social interaction, can sometimes contribute to the formation or modification of an attitude. The dynamic interplay between these layers makes their individual and collective analysis crucial for a comprehensive understanding of human communication.

To sum up, opinions tend to be more rooted in cognitive processes - they are often the result of consideration, evaluation, and assessment. They can be can be logical, informed, or even biased and can be strongly held or loosely formed. Sentiment, in this context, refers to the underlying emotional tone or feeling associated with an opinion, statement, or topic. It is the affective component of an opinion. Thus, it is most often a complex combination of affect (emotions, feelings or moods) and a cognitive process of thinking.

2 Influence of external events

The dynamism between opinions, emotions, and attitudes with respect to climate change, could be influenced externally both by external weather conditions and by political events. The period between 2015 and 2019 from which we draw our data stands as a pivotal time in the evolution of global climate consciousness, as it was during these five years that the tangible impacts of a warming planet became increasingly undeniable and visible to the broader public due to a series of major climate-related events, coupled with significant policy milestones, served as stark reminders of the escalating crisis, effectively shifting public perception from abstract threat to immediate reality. One of the most significant overarching trends during this period was the consistent shattering of temperature records and the increasingly evident link between these extreme events and human-induced climate change became a central theme in public discourse. The political landscape also played a crucial role in shaping climate consciousness during this period. The Paris Agreement on climate change[2], adopted in December 2015, marked a landmark moment. Furthermore, the Intergovernmental Panel on Climate Change (IPCC) Special Report on Global Warming of 1.5°C, released in October 2018, delivered a stark warning. On the other hand, Donald Trump's presidency (which began in January 2017) was marked by a significant departure from previous U.S. climate policy, actively seeking to dismantle environmental regulations and withdraw from international climate agreements[3]. Paradoxically though, Trump's actions and controversy also exerted impact on climate change awareness. And it was then, in August 2018, that Greta Thunberg[4] emerged as an unparalleled global icon for climate action during the latter part of this period, dramatically amplifying climate consciousness, especially among young people.

3 Computational and linguistic approaches

The rise of computational approaches, particularly in natural language processing (NLP) and machine learning, has revolutionized the study of opinion, sentiment, and attitude in discourse. The sheer volume of digitally available text data – from social media feeds and online forums to news articles and customer reviews – has created an unprecedented opportunity to analyze subjective language at scale. Early efforts in sentiment analysis focused primarily on the lexicons of positive and negative words, but more sophisticated techniques now employ machine learning algorithms, deep learning models, and contextual embeddings to capture the nuances of human emotion and subjective expression. These computational tools enable researchers and practitioners to automatically detect sentiment [4], identify opinion holders, track opinion evolution, and even infer underlying attitudes from large corpora of text.

However, the computational analysis of opinion, sentiment, and attitude is not without its challenges. The inherent subjectivity and context-dependency of human language pose significant hurdles. Sarcasm, irony, negation, and implicit expressions of sentiment can easily mislead automated systems. For example, "This movie was so good I fell asleep" expresses negative sentiment despite the positive word "good." Furthermore, differentiating between factual statements, expressions of opinion, and expressions of sentiment can be complex, as these often intertwine in natural discourse. The development of robust and accurate computational models requires sophisticated linguistic knowledge, large annotated datasets, and continuous refinement to account for the complexities of human communication.

Beyond computational approaches, traditional linguistic and discourse analytic methodologies remain indispensable for a deeper qualitative understanding. Discourse analysis and rhetorical strategies, for example, examines how opinions, sentiments, and attitudes are constructed, negotiated, and contested within specific communicative contexts. Semantics and pragmatics provide frameworks for understanding the meaning and intended effect of subjective language, considering factors such as the intention of the speaker, the shared knowledge, and the social norms. By integrating computational and qualitative approaches, researchers can achieve a more comprehensive and nuanced understanding of these multifaceted phenomena and uncover stances [1, 8, 9] towards climate change.

4 Aims of the analysis

The aims of our analysis are to uncover the social media users' attitudes, embracing sentiment, towards the climate change in two social media datasets from different periods of time and investigate to what extent the public opinions and attitudes towards this issue change within that period of time. As individuals categorize information based on their existing beliefs and attitudes, new information is evaluated in relation to these existing attitudes, leading to either the new information assimilation (acceptance) or contrast (rejection) [7]. A general theoretical approach adopted here refers to the identification of the users' opinions and attitudes via the sentiments expressed with reference to the opinions and argumentation persuasive effects,

[2] https://unfccc.int/process-and-meetings/the-paris-agreement

[3] https://www.nytimes.com/2017/06/01/climate/trump-paris-climate-agreement.html

[4] https://awpc.cattcenter.iastate.edu/2019/12/02/speech-at-cop24-dec-4-2018/

also in terms of Aristotle's persuasive appeals. An important part of this theory is played by three main factors - logic, emotion and trust-based types of persuasion. Logos is the appeal to the logic or reasoning by referring to evidence. Pathos is the appeal to the emotions of the audience. Trust-based ethos appeal invokes the credibility or trustworthiness of individuals considered the source of information. Sentiment marking in datasets can refer to any of the three appeals while opinions have specific content orientation (Mao et al. 2024). That is why some opinionated utterances can be ambiguous between sentence polarity and opinion polarity or between literal and non-literal (eg ironic) utterances such as e.g., You'd rather drown in these floods than admit global warming.

Previous attempts at opinion sentiment analysis are multiple and varied. In their (2024) review paper Mao et al. [15] critically assess available systems and highlight their challenges and limitations. To minimize the limitation of these methodologies, two types of sentiment analysis are used in the present study, one based on human marking, which can be more reliable but time-consuming and hence inefficient, and the other - on automatic lexicon-based labels - which is fast but can provide erroneous results with regard to the results of the opinion sentiment marking due to the above mentioned possible distinction between appeal polarity and opinion stance sentiments [12, 11]. As the results obtained were not fully consistent, we used LLMs to assist in the clarification of the distinction and to scrutinize a possible statistically significant sentiment flow between the earlier and later public opinion expression.

5 Literature review

The topic of automated sentiment analysis is not new in computational linguistics. The reference to and basing the identification on emotion word dictionaries is typically the first method to think of. However, the emotion sentiment values are often ambiguous outside context. Prediction of sentiment data with the use of machine-learning algorithms required heavy pre-processing preparations but it is machine learning that is regarded as being one of the latest and most prevalent techniques in sentiment analysis [2, 17]. Another related problem refers to an unreflexive use of lexicon sentiment analysis for opinion sentiment which could lead to erroneous and untrusthworthy results due to considering sentiment identification synonymous with opinion polarity marking. The most recent model of sentiment analysis and opinion mining is developed by Maruthupandi et al. [16]. This framework (SemAI) analyzes opinion sentiment by applying AI-based semantic analysis to assess and classify user views from social data. The authors use the Degree of Correlation Network Model (DCNM) to extract the subset of basic features from social network data with the meta-heuristics model - Heap based Optimization (HbO), to reduce the dimensionality of retrieved features. Finally, the unique classification algorithm, Self-Attention based Deep Analyzing Network (SA-DAN), is used to identify and categorize attitudes into positive, neutral, and negative. The SemAI model achieves particularly good results for accuracy, precision, as well as F1-scores in a variety of social data sets. Our approach is an attempt to reconcile the lexicon- and text-based sentiment value approaches via contextual BERT text embeddings, their clustering and topic labeling and interpretation to observe the dynamicity of opinion shift in public opinion sentiment values related to climate change.

6 Experiments

6.1 Datasets

To investigate the evolution of public opinion regarding climate change across different time periods, we employed two datasets composed of climate-related tweets. Each dataset captures sentiment toward climate change but differs in its time range, method of data collection, and sentiment labeling.

6.1.1 Twitter Climate Change Sentiment Dataset (2015–2018)

The first dataset, titled **Twitter Climate Change Sentiment Dataset**[5], was collected between April 27, 2015, and February 21, 2018, with support from the Canada Foundation for Innovation JELF Grant to Chris Bauch at the University of Waterloo. It consists of 43,943 tweets related to climate change, each manually annotated by three independent reviewers. Only tweets for which all three annotators reached unanimous agreement were retained, while others were discarded to ensure label reliability.

Each tweet in this dataset is assigned one of four sentiment categories:

- 2 (News): Links to factual news about climate change
- 1 (Pro): Supports the belief in anthropogenic climate change
- 0 (Neutral): Neither supports nor refutes the belief in anthropogenic climate change
- -1 (Anti): Rejects the belief in anthropogenic climate change

This dataset provides high-quality, human-annotated sentiment labels that are well-suited for training and evaluating classification models over historical climate discourse.

6.1.2 TwitterSocialMediaAnalysis_ClimateChange (2019)

The second dataset, referred to as **TwitterSocialMediaAnalysis_ClimateChange**[6], comprises tweets collected throughout 2019 via the Twitter API. More than 50,000 tweets were gathered from major cities across the United States and various international locations. Tweets were obtained using advanced search strategies, focusing on trending climate-related hashtags, including:

> #climateStrike, #climatestrike, #climatechange, #GreenNewDeal, #climatecrisis, #climateAction, #FridaysForFuture, #environment, #globalwarming, #GlobalWarming, #ActOnClimate, #sustainability, #savetheplanet, #bushfiresAustralia, #bushfires

To analyze public sentiment, two sentiment analysis tools were applied: TextBlob[7] and VADER (Valence Aware Dictionary and sEntiment Reasoner)[8]. In particular, VADER's compound sentiment score was averaged for each hashtag to quantify general sentiment trends. Unlike the first dataset, which contains categorical human labels, this dataset provides continuous sentiment scores based on automatic, lexicon-based sentiment analysis methods.

Together, these two datasets enable a comparative analysis of climate change sentiment across different years and geographic regions, while also contrasting manual and automated sentiment annotation approaches.

---

5 https://www.kaggle.com/datasets/edqian/twitter-climate-change-sentiment-dataset
6 https://github.com/tasvora/TwitterSocialMediaAnalysis_ClimateChange
7 https://textblob.readthedocs.io/en/dev/
8 https://github.com/cjhutto/vaderSentiment

6.2 *Topic Identification via BERT-Based Clustering*

To explore how public sentiment toward specific climate-related topics evolved over time, we first aimed to identify dominant themes within these dataset. For this purpose, we employed a BERT-based clustering approach that allows for the extraction of semantically coherent topics from unstructured text.

Our method consists of the following key steps:

1. **Text Embedding with BERT** Each tweet was encoded into a high-dimensional dense vector using a pre-trained BERT model[6]. These embeddings capture contextual semantic information, allowing similar tweets to be represented by similar vectors in the embedding space. This step enables clustering to be based on meaning rather than surface-level lexical similarity.
2. **Clustering with k-means** The resulting tweet embeddings were clustered using the k-means algorithm. This unsupervised method partitions the data into k clusters by minimizing intra-cluster distances. Each cluster is interpreted as representing a distinct topic, with the assumption that tweets within the same cluster discuss related subject matter.
3. **Topic Labeling and Interpretation** To interpret each cluster, we extracted the most representative terms and manually reviewed example tweets. This allowed us to assign meaningful labels to clusters, such as "climate policy," "natural disasters," or "climate activism." These topics were then tracked across datasets to assess how sentiment shifted around each theme over time.

It is crucial to note that we first identify these topics from the combined dataset of both periods. Subsequently, we analyze the dominance of each identified topic within each individual period. This approach avoids the complex challenge of needing a mapping function to align similar topics across different periods. Such a mapping function could potentially complicate the comparison, as it might not be a one-to-one correspondence between a topic's representation in one period and its equivalent in another. By identifying topics from the outset on the aggregated data, we ensure a consistent thematic framework for cross-period analysis.

This BERT-based clustering framework enables nuanced detection of thematic structures in social media discourse and provides a robust basis for analyzing topic-specific opinion dynamics.

6.3 *Unified Climate-Focused Sentiment Analysis*

To enable a coherent comparison of public attitudes toward climate change across datasets, we applied a unified sentiment classification framework as a method to uncover public opinions. While both datasets contain some form of sentiment annotation, they differ in scope and interpretation. Our goal was to standardize sentiment analysis so that, for every tweet, the sentiment reflects its stance *specifically toward climate change*—independently of any associated political, social, or contextual references.

We adopted a three-class sentiment scheme: *positive*, *negative*, and *neutral*, referring strictly to the tweet's attitude toward climate change. However, a review of the original sentiment annotations revealed two key limitations. First, as noted in Section 6.1.1, the first dataset encodes sentiment in terms of belief in anthropogenic climate change, while the second uses general sentiment labels without topic-specific context. This mismatch complicates direct cross-dataset analysis.

Second, neither dataset consistently isolates sentiment toward climate change itself. Many tweets express sentiment toward peripheral topics such as politicians, movements, or media coverage, which may not reflect the author's view on the climate issue directly.

To resolve these issues, we employed a large language model (LLM) — specifically, `Meta-Llama-3-8B-Instruct`[9] — to reclassify all tweets under a consistent and climate-specific sentiment framework. We designed a prompt (Fig. 1) that instructs the model to assess each tweet's sentiment solely in relation to climate change. The model then outputs one of the three standardized sentiment labels.

Sentiment Annotation Prompt for Climate Change Texts

You are tasked with performing sentiment analysis on texts referencing climate change (e.g., "climate," "weather," "global warming"). Assign one of the following sentiment labels:
Positive: Uses affirmative language about climate change without negation. Examples: "Global warming is harmful," "World leaders must act on climate change," or supportive figurative expressions.
Negative: Denies or downplays climate change. Includes negation (e.g., "There is no global warming"), or sarcasm dismissing concerns.
Neutral: Ambiguous or non-committal stance (e.g., "I don't know if climate change is taking place") or lacking clear opinion.
Label each sample with only one word: **Positive**, **Negative**, or **Neutral**.
Text to analyze: `{text}`

Figure 1. Prompt used for annotating sentiment in climate change-related texts.

This process allowed us to harmonize sentiment interpretation across both datasets and ensured that all subsequent analyses reflect genuine sentiment toward climate change, regardless of the dataset's original structure or annotation scheme.

By applying this unified sentiment framework to tweets clustered by topic—across both datasets—we are able to track how public sentiment toward each climate-related theme evolved between the earlier (2015–2018) and later (2019) periods. This sets the foundation for our temporal analysis of climate opinion dynamics, which is presented in the following section.

7 Results

7.1 *Estimating the Number of Topics*

As a preliminary step in the analysis, we aimed to estimate a suitable number of dominant topics represented in the combined tweet corpus. To that end, we applied the K-Means clustering algorithm with varying values of K, representing different potential topic counts.

To qualitatively assess the coherence and separability of the resulting clusters, we visualized the tweet embeddings using t-distributed Stochastic Neighbor Embedding (t-SNE). This dimensionality reduction technique maps high-dimensional vectors into a two-dimensional space while preserving the local structure of the data, thereby allowing us to visually evaluate the clustering structure.

[9] https://llama.meta.com

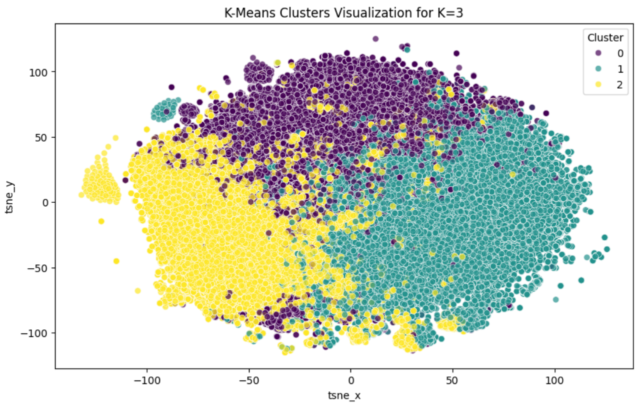

Figure 2. t-SNE visualization of tweet embeddings clustered using K-Means with $K = 3$.

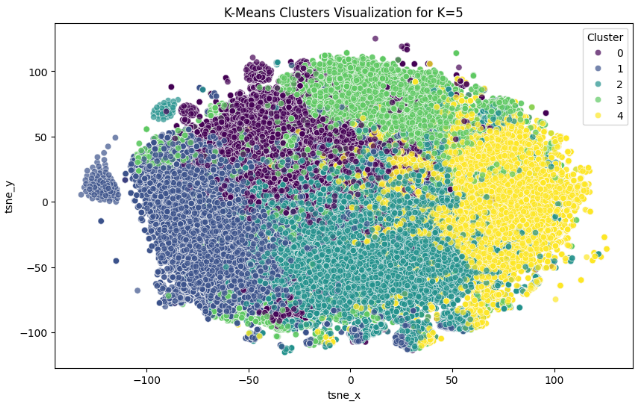

Figure 3. t-SNE visualization of tweet embeddings clustered using K-Means with $K = 5$.

Figures 2 and 3 present the clustering results for $K = 3$ and $K = 5$, respectively. The visualization for $K = 3$ reveals relatively well-separated clusters, suggesting distinct thematic groupings in the data. In contrast, the $K = 5$ configuration appears more fragmented, with less distinct boundaries between some clusters.

These visualizations served as an initial guide in selecting an appropriate number of topics for further analysis.

The distribution of tweets across the three clusters identified through K-Means clustering with $K = 3$ is as follows: Cluster 0 accounts for approximately 28.8% of the tweets, Cluster 1 for 38.8%, and Cluster$\tilde{2}$ for 32.4%. The specific thematic content of each cluster is interpreted in the next section.

7.2 *Topic Interpretation Based on Clustered Keywords*

To interpret the nature of each identified topic, we analyzed the most (30) frequent and representative words within each cluster. These clusters were extracted from the combined dataset of both periods. Table 1 lists the top terms for each cluster when applying K-Means with $K = 3$. The vocabulary associated with each cluster offers insight into the thematic content represented by the tweets within it.

Based on the prominent words in each cluster, we propose the following thematic interpretations:

- **Cluster 0 – Climate Change Controversy and Politics:** This cluster includes terms such as *trump*, *believe*, *real*, *president*, and *science*, suggesting a discourse focused on political debate and skepticism or affirmation regarding climate change. The presence of *epa* and *scientists* further indicates that this topic may center on the legitimacy of climate science and institutional responses.

Table 1. Top Words per Cluster ($K = 3$), extracted from the combined dataset of both periods.

Cluster 0	Cluster 1	Cluster 2
climate	climatechange	climate
change	climate	change
global	climatestrike	climatestrike
warming	climateaction	global
trump	climatecrisis	warming
climatechange	change	greennewdeal
climatestrike	amp	climatechange
amp	sustainability	today
world	environment	people
environment	globalwarming	amp
new	today	climateaction
people	people	like
real	greennewdeal	globalwarming
believe	world	believe
like	action	trump
science	new	environment
president	need	going
says	future	world
climatecrisis	global	need
epa	time	time
climateaction	planet	new
need	earth	actonclimate
fight	great	future
action	savetheplanet	day
scientists	youth	planet
time	day	la
energy	like	real
going	energy	strike
years	actonclimate	thank
planet	green	right

- **Cluster 1 – Climate Activism and Global Sustainability:** This cluster features frequent hashtags and terms associated with organized movements and environmental advocacy, such as *climatestrike*, *climateaction*, *sustainability*, *greennewdeal*, and *savetheplanet*. It likely represents the narrative of global climate activism, particularly involving youth movements and calls for systemic change.
- **Cluster 2 – Climate Events and Urgency:** With terms like *today*, *going*, *strike*, *thank*, and *right*, this cluster appears to reflect real-time reactions to climate-related events, including public demonstrations and environmental crises. The co-occurrence of *greennewdeal* and *actonclimate* supports the interpretation that this topic captures urgent, event-driven discourse.

7.3 *Sentiment Dynamics Within Topics*

The stacked bar chart (Figure 4 illustrates a comparative analysis of sentiment distribution across three distinct topics (labeled 0, 1, and 2) between two distinct data sources (referred to as Source 1 and Source 2). Each bar represents a specific topic within a given data source, segmented to display the proportions of negative, neutral, and positive sentiment. A general observation reveals a prevalence of positive sentiment across most topic-source combinations. However, notable variations emerge upon closer examination. For Topic 0, Source 1 exhibits a considerably higher proportion of negative sentiment compared to Source 2, which displays a markedly lower negative sentiment and a correspondingly higher positive sentiment. In contrast, Topic 1 demonstrates a relatively consistent sentiment distribution across both sources, with a dominant positive sentiment

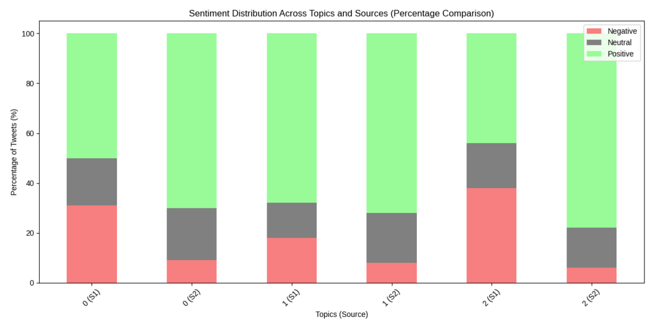

Figure 4. Sentiment distribution across three topics for two data sources, highlighting temporal differences in public sentiment.

and similar, albeit slightly elevated in Source 1, levels of negative and neutral sentiment. The most pronounced difference is observed in Topic 2, where Source 1 is characterized by the highest proportion of negative sentiment among all combinations, alongside a substantial neutral sentiment. Conversely, Source 2 for Topic 2 shows a significant reduction in negative sentiment and a dominance of positive sentiment. In conclusion, the observed variations in sentiment distribution between Source 1 and Source 2 for each topic suggest a potential temporal evolution in public sentiment, as each source represents a distinct time period.

To determine whether the distribution of sentiment categories (positive, neutral, negative) significantly differs between the two sources within each topic cluster, we conducted a Chi-Square Test of Independence for each cluster separately. This test evaluates whether the proportions of sentiment labels are independent of the data source, while accounting for the differences in sample size between sources. The results revealed statistically significant differences in sentiment distribution across all three clusters. In each case, the test yielded a p-value less than 0.01, indicating that the proportion of sentiment types (e.g., the ratio of positive to negative tweets) changed meaningfully between the two time periods represented by the sources. These findings suggest a temporal shift in sentiment expression within each identified topic.

8 Discussion and conclusions

The main aim of this study was to investigate how public opinions expressed in terms of sentiment toward climate change evolved over time. As is known sentiment analysis can be performed at various granularities, each offering different insights, such as the word level, document-, sentence- or aspect-levels. The closest brief description of our methodology employed in the study is a hybrid approach, i.e., a combination of some of these levels. We first specified particular dominant climate-related topics that evolved over time in the two datasets we explored by employing a BERT-based clustering approach to extract semantically coherent topics from unstructured texts. Each tweet was encoded into a BERT-related vector, capturing contextually relevant semantic information. The problem with the two datasets we used was that they were not identical in scope and in the contextual information annotation. Some of the sentiment values were assigned not to general topics or main - climate change - targets, but to some lexical entities, peripheral from the point of view of a particular text opinion sentiment. The BERT transformer advanced the study, more effectively capturing long-range dependencies and contextual relationships.

The analysis revealed dynamic patterns in climate change discourse, frequently characterized by polarization and shifts in dominant narratives. A notable observation was how public expressions of sentiment modulates, often reflecting concern, anxiety, or fear, probably intensified during extreme weather events or following significant climate-related news. This suggests a direct link between real-world occurrences and shifts in public opinion, demonstrating the dynamic, responsive nature of public sentiment and highlighting the importance of real-time monitoring for climate communication strategies.

An important part of our study was a distinction between lexical emotion sentiment and more general, opinion thematic sentiments that might often contradict. This layering of opinion via the opinion sentiment targets might be an important site of opinion and attitude sentiment distinction in our further studies. There are new proposals to extend the effectiveness of sentiment identification in data such as e.g., semantic analysis based on AI (SemAI) for assessing and categorizing user views from social data mentioned in the Literature review secton here, which applies algorithms in the stages of feature extraction, optimization, and sentiment prediction [16] with The Degree of Correlation Network Model (DCNM) and the Heap based Optimization model, supported by Self-Attention based Deep Analyzing Network (SA-DAN). And yet, despite significant advancements, opinion-related sentiment analysis continues to face several challenges that would be undertaken in the future studies. They are primarily related to implicit meanings such as figurative language or sarcasm and irony detection. Another hurdle, often requiring deep contextual and commonsense understanding, is handling the scope of negation and intensification. As is also encountered in our analyzed data, there are still problems with context and domain dependency as the opinion sentiment can change across contexts and domains as well as target orientation that can vary depending on whether we identify sentiments towards particular entities expressed in lexical units or whether the sentiment of the whole entity (e.g., post or document).

An important methodological asset of our study lies in the integration of lexicon-based tools, human annotation, and LLM-powered reclassification within a unified framework for sentiment analysis. The use of BERT embeddings allowed for a semantically nuanced clustering of topics beyond surface-level lexical similarity, while LLMs enabled a more contextualized and climate-specific reinterpretation of sentiment. This computational pipeline ensured interpretability and adaptability across datasets with divergent structures.

Comparing Source 1 and Source 2 climate change Twitter data, some of the opinion sentiment marking agrees with human contextual sentiment recognition (see Table 2). For example:

> @GretaThunberg Thank you for standing up #ClimateStrike

The two sentiment models—VADER and LLMs—mark the positive sentiment of the above opinion consistently.

However, several challenges remain: sentiment models, especially lexicon-based ones like VADER, are sensitive to contextual ambiguity, figurative language, and sarcasm. Both VADER and even LLMs, despite their contextual strength, may struggle with stance disambiguation when multiple sentiment targets (e.g., politicians vs. climate change) co-occur. An example of such opinion sentiment recognition problems can be encountered in the following case:

> @tiniebeany climate change is an interesting hustle as it was global warming but the planet stopped warming for 15 yes while the suv boom (Dataset 1)

Such and similar posts are indirect and are interpreted as denying the reality of global climate change, based on the uninformed

Dataset	#	Example	Sentiment	LLM	VADER
1	1	Fabulous! Leonardo #DiCaprio's film on #climate change is brilliant!!! Do watch. `https://t.co/7rV6BrmxjW` via @youtube	Positive	Positive	Positive
1	2	@tiniebeany climate change is an interesting hustle as it was global warming but the planet stopped warming for 15 yes while the suv boom	False	Neutral	Positive
2	1	The people on #ClimateStrike in DC are marching to confront the banks funding the climate crisis! #FireDrillFriday #ShutDownDC	Negative	Positive	Negative
2	2	@realdonaldtrump is presiding over a mass extinction and climate catastrophe. We need to take action now! #ClimateStrike shutdowndc923 @Trump International Hotel Washington, D.C.	Negative	Positive	Negative
2	3	@GretaThunberg Thank you for standing up #ClimateStrike	Positive	Positive	Positive

Table 2. Sentiment comparison between LLM and VADER across two datasets

or disinformed types of opinion. The opinion sentiment is **Negative**, even though the positive lexicon sentiment lexis—"interesting" or "stopped warming"—is used. Such contextual clues may be misleading and cause LLM **Neutral** marking and **Positive** by VADER. Problems with LLM-generated correct identification of implicit, as well as some of the indirectly conveyed opinions, were first reported in Liebeskind and Lewandowska-Tomaszczyk [14]. In our present paper more nuanced opinion recognition problems are noted. Neither VADER nor LLM recognized the anti-climate change opinion sentiment, although the lexical form "hustle" may function as an interpretative indicator.

Another example (Dataset 2) also presents questionable climate change sentiment recognition:

> The people on #ClimateStrike in DC are marching to confront the banks funding the climate crisis! #FireDrillFriday #ShutDownDC

VADER marks the post sentiment as **Negative**, and LLM as **Positive**, while the post is immersed in a more complex outside context—#ClimateStrike confronting bank funding of investments which accelerate climate change. In other words, the post includes two opinion sentiment stances: climate change recognition (**Positive**) and opposition to the institutions contributing to the process (**Negative**).

Thus, while our methodology enables a richer analysis of opinion dynamics, it also underscores the need for further refinement in handling subtle pragmatic cues and layered sentiment targets. At the same time, this research unequivocally demonstrates that universal opinion is not solely shaped by one mode of persuasion, but rather by a dynamic interplay of appeals. While logos-based arguments provide the foundational framework of factual evidence and logical reasoning that underpins informed perspectives, they are often insufficient on their own to influence widespread belief. We found that emotional appeals are equally critical, tapping into shared human experiences, values, and sentiments to foster resonance and engagement. Furthermore, the ethos of the source—its credibility, authority, and perceived trustworthiness—plays a pivotal role in determining the receptiveness of an audience to any given message such as eg. in the case of Greta Thunberg.

Our analysis, while comprehensive within its defined parameters, may be limited by its focus on specific social media platforms (Twitter) and a fairly short time frame. This can affect the generalizability of findings across more diverse textual sources (e.g. news articles, forums), different cultural contexts, or linguistic variation. One can also indicate a general lack of consideration for user demographics (inaccessible in social media data) and entirely external factors that shape opinions in climate change discourse, a limitation that may also apply to aspects of this study. Yet, in our attempt to use topic cluster modeling and fairly intensive scrutiny, some of these threats may have been overcome.

There are also specific linguistic-pragmatic features that signal subtle shifts in meaning or stance [8]. They are hedging devices (e.g., "perhaps," "it seems") and modal verbs (e.g., "may," "might," "could"), which signal caution, uncertainty, or indirectness in expressing claims, first-person pronouns, indicating a subjective stance and authorial presence, nominalizations (such as "global warming") and passive voice constructions (e.g., "it was observed"), which contribute to informational density and objectivity, often favored in factual or scientific discourse. Furthermore, the pragmatic analysis can delve more deeply into the contextual and consequential dimensions of language (presupposition, implicature, speaker's intent) and, finally, implicit meanings detection (figurative language, irony, sarcasm).

Needless to say, our analysis mirrored to a large extent all those important political events that took place in the period between 2015 - 2019. The generated thematic clusters contained items such as TRUMP and GRETA as their main headwords. Hence, future investigations should also extend research towards more effective modeling of real-world debates on complex societal issues. Conducting comprehensive cross-cultural and cross-linguistic analyses of climate narratives and opinion change is also vital, exploring how language, cultural norms, and communication styles influence opinion dynamics globally.

Incorporating more realistic social network structures into opinion dynamics simulations will better reflect how opinions propagate and change within diverse social groups. Furthermore, a significant leap from merely detecting opinion change to understanding how to influence it responsibly involves investigating causal relationships between narratives, linguistic shifts, and real-world outcomes, moving beyond mere correlation to understand how specific narratives directly influence behavior and societal change.

Future research might also prioritize refining LLMs by fine-tuning them with more diverse and real-world human discourse data, including content that reflects polarized or fact-resistant opinions. This is crucial to overcome current LLM biases, negatvely influencing e.g., word embeddings, thereby enabling more accurate simulations and analyses of complex opinion dynamics.

Further exploration of hybrid computational-linguistic models is warranted, combining the strengths of advanced LLMs with explicit linguistic and pragmatic theories to enhance accuracy, interpretability, and detection of subtle opinion shifts. Developing more nuanced and multidimensional representations of opinion, moving beyond scalar sentiment or stance to capture the complexity of human beliefs and attitudes, is also a key direction.

Following this line of enquiry, our further research in opinion-related sentiment analysis is thus likely to use a hybrid linguistic

NLP research methodology and focus on more robust contextual understanding that can better capture implicit and nuanced language. Furthermore, it goes without saying, that incorporating deeper aspects of common sense reasoning would improve opinion sentiment understanding.

For future work, exploring alternative topic modeling approaches beyond BERT-based clustering, such as Latent Dirichlet Allocation (LDA), could offer complementary insights into the underlying thematic structures of climate change discourse. While BERT excels in capturing semantic nuances through contextual embeddings, probabilistic topic models like LDA could provide a different perspective on topic composition by identifying latent topics based on word co-occurrence patterns, potentially revealing distinct and interpretable themes. This comparative analysis would enrich our understanding of topic dynamics and validate the robustness of identified themes across different methodological paradigms.

Acknowwledgements The study was prepared in the Linguistics group of COST Action CA 21129 *What are Opinions? Integrating Theory and Methods for Automatically Analyzing Opinionated Communication* (OPINION). The authors thank Avigail Haliva and Shilat Shimon for their significant assistance in implementing the code and running the experiments. Their contribution played an important role in the technical execution of the study.

References

[1] A. Addawood, J. Schneider, and M. Bashir. Stance classification of twitter debates: The encryption debate as a use case. In *Proceedings of the 8th international conference on social media & society*, pages 1–10, 2017.

[2] H. J. Alantari, I. S. Currim, Y. Deng, and S. Singh. An empirical comparison of machine learning methods for text-based sentiment analysis of online consumer reviews. *International Journal of Research in Marketing*, 39(1):1–19, 2022.

[3] D. Albarracin and S. Shavitt. Attitudes and attitude change. *Annual review of psychology*, 69(1):299–327, 2018.

[4] A. Chourasiya, A. Khan, K. Bajaj, M. Tomar, T. Kohli, and D. Chauhan. A review of sentiment analysis and emotion detection from text using different models. 2025.

[5] G. Colombetti. *The feeling body: Affective science meets the enactive mind.* MIT press, 2014.

[6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. 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, volume 1 (long and short papers)*, pages 4171–4186, 2019.

[7] M. E. Doherty and E. M. Kurz. Social judgement theory. *Thinking & Reasoning*, 2(2-3):109–140, 1996.

[8] J. W. Du Bois. The stance triangle. In *Stancetaking in discourse: Subjectivity, evaluation, interaction*, pages 139–182. John Benjamins Publishing Company, 2008.

[9] L.-W. Ku, C.-Y. Lee, and H.-H. Chen. Identification of opinion holders. In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 14, Number 4, December 2009*, 2009.

[10] B. Lewandowska-Tomaszczyk. Online interconnectivity and negative emotion patterning. *Sociedad de la Información*, 44:76–109, 2013.

[11] B. Lewandowska-Tomaszczyk and C. Liebeskind. Opinion events and stance types: advances in llm performance with chatgpt and gemini. *Lodz Papers in Pragmatics*, 20(2):413–432, 2024. doi: doi:10.1515/lpp-2024-0039. URL https://doi.org/10.1515/lpp-2024-0039.

[12] B. Lewandowska-Tomaszczyk, C. Liebeskind, A. Bączkowska, J. Ruzaite, A. Dylgjeri, L. Kazazi, and E. Lombart. Opinion events: Types and opinion markers in english social media discourse. *Lodz Papers in Pragmatics*, 19(2):447–481, 2023.

[13] C. Liebeskind and B. Lewandowska-Tomaszczyk. Navigating opinion space: A study of explicit and implicit opinion generation in language models. In *Proceedings of the First LUHME Workshop*, pages 28–34, 2024.

[14] C. Liebeskind and B. Lewandowska-Tomaszczyk. Opinion identification using a conversational large language mode. In *The International FLAIRS Conference Proceedings*, volume 37, 2024.

[15] Y. Mao, Q. Liu, and Y. Zhang. Sentiment analysis methods, applications, and challenges: A systematic literature review. *Journal of King Saud University-Computer and Information Sciences*, page 102048, 2024.

[16] J. Maruthupandi, S. Sivakumar, V. S. Kumar, and P. B. Srikaanth. Semai: A unique framework for sentiment analysis and opinion mining using social network data. *SN Computer Science*, 6(2):99, 2025.

[17] H. Zhao, Z. Liu, X. Yao, and Q. Yang. A machine learning-based sentiment analysis of online product reviews with a novel term weighting and feature selection approach. *Information Processing & Management*, 58 (5):102656, 2021.