

Research Article

Mapping Public Sentiment on Generative AI via Twitter NLP and Topic Modeling

Marcelino Caetano Noronha^{1,*}, Saruni Dwiasnati², Cherlina Helena P Panjaitan³

¹ Department of Informatics Engineering, Faculty of Engineering Science and Technology, National University of Timor-Lorosa'e

² Universitas Mercu Buana, Indonesia

³ Universitas Sains dan Teknologi Komputer, Indonesia

* Corresponding Author: e-mail : marcelino.noronha@untl.edu.tl

Abstract: The rapid diffusion of Generative Artificial Intelligence (AI) has intensified public debate regarding its benefits, risks, and societal implications. This study investigates public sentiment and thematic structures surrounding Generative AI by analyzing Twitter discourse as a representation of large-scale, real-time public perception. The research addresses two main problems: how public sentiment toward Generative AI is distributed and what dominant themes shape this perception. Accordingly, the objective is to map both emotional polarity and thematic narratives embedded in social media conversations. A computational mixed-methods approach was employed using a dataset of 12,470 tweets collected on 17 December 2024. Sentiment classification was conducted using a transformer-based DistilBERT model, while semantic representations were generated with Sentence-BERT. Topic modeling was performed using BERTopic, integrating HDBSCAN clustering and class-based TF-IDF to extract coherent and interpretable topics. Human-in-the-loop validation supported the interpretive robustness of topic labeling. The findings reveal that public sentiment toward Generative AI is predominantly positive (41.8%), particularly in relation to productivity enhancement, education, and creative applications. Neutral sentiment (31.4%) reflects informational discourse, while negative sentiment (26.8%) centers on ethical concerns, privacy risks, misinformation, and AI hallucinations. Seven dominant topics were identified, with clear topic-sentiment alignment showing optimism in utility-driven themes and skepticism in ethics- and risk-related discussions. In conclusion, public perception of Generative AI is dualistic—characterized by strong enthusiasm alongside persistent caution. These results provide empirical insights for AI governance, responsible innovation, and future research on socio-technical impacts of Generative AI.

Keywords: Generative Artificial Intelligence; Sentiment Analysis; BERTopic Modeling; Public Perception; Social Media Analytics

Received: **date**

Revised: **date**

Accepted: **date**

Published: **date**

Curr. Ver.: **date**



Copyright: © 2025 by the authors.

Submitted for possible open

access publication under the

terms and conditions of the

Creative Commons Attribution

(CC BY SA) license

(<https://creativecommons.org/licenses/by-sa/4.0/>)

1. Introduction

The rapid emergence of generative artificial intelligence (AI) has reshaped how individuals engage with information, produce digital content, and interact with automated systems. Since the release of large language models (LLMs) such as ChatGPT, public discourse on social media has intensified, with Twitter (X) becoming a vibrant arena for expressing opinions, concerns, and expectations regarding AI-generated content. Online discussions about generative AI tend to revolve around themes of ethics, creativity, academic integrity, trust, and future societal impacts. These discussions illustrate the dual nature of generative AI in the public imagination, representing both opportunities for advancement and sources of

uncertainty (Jeon et al., 2025; Miyazaki et al., 2024). Existing research reinforces this duality. (Cohen et al., 2024) showed that ethical concerns such as transparency, misinformation, and fairness, dominate social media conversations on ChatGPT, particularly in communities aware of algorithmic risks. Gender-based perspectives also shape discourse: (Møgelvang et al., 2024) found significant gender differences in the use of generative chatbots in higher education, indicating that demographic variables influence both adoption behavior and sentiment.

Similarly, (Sun & Zhou, 2025) demonstrated that attitudes toward generative AI are mediated by multifactorial influences such as academic pressure, digital literacy, and personal values. These findings align with the PDF mapping results, which show clusters of public sentiment around ethical dilemmas, fear of replacement, and excitement about productivity gains. Social media expressions thus provide an important empirical lens for understanding how generative AI is framed within public consciousness. Beyond education, public perception also varies across cultural contexts. (Kim et al., 2022) emphasized that cultural background shapes individuals' trust and emotional responses toward AI systems, revealing cross-cultural differences in perception of AI risks and benefits. (Tao et al., 2024) similarly highlighted cultural alignment issues in LLMs, raising questions about whether generative AI reflects diverse value systems or unintentionally reproduces cultural biases. These concerns frequently appear in topic clusters in the Scopus PDF, where terms associated with fairness, representation, and cultural identity emerge prominently. The conversation is further complicated by the ethical terrain of AI-generated content. (Medhat et al., 2025) and (Are et al., 2025) emphasized that content creators and online users express both optimism and anxiety toward generative tools, particularly regarding authenticity and authorship. In social media ecosystems, misinformation risks are amplified due to the scale and speed of content production, leading to increased calls for accountability and governance. This aligns with governance-focused concerns mapped in the PDF, where transparency, explainability, and regulatory frameworks are frequently discussed as critical expectations for trustworthy AI deployment.

Understanding these sentiments at scale requires sophisticated computational methods. Modern NLP techniques have improved the reliability of sentiment classification on short, informal texts such as tweets. Transformer-based models, such as DistilBERT, RoBERTa, and fine-tuned BERT variants, demonstrate superior capability for text classification tasks (Sbei et al., 2025; Shao, 2025; Xiao & Luo, 2024). Sentence embedding models further enhance semantic coherence in identifying latent sentiment patterns. Meanwhile, topic modeling approaches such as LDA (Harshvardhan et al., 2021) and BERTopic (Ou et al., 2025) allow researchers to capture the underlying thematic structures within large-scale social media data. The Scopus PDF itself used BERTopic-like clustering to visualize topic groupings, showing how themes surrounding "ethics," "education," "AI misuse," and "creative assistance" naturally emerge in Twitter conversations. Multimodal sentiment analysis is also increasingly relevant as social media users express emotions through text, emojis, images, and videos. (Panjaitan et al., 2025) and (Jin et al., 2024) highlighted that multimodal sentiment analysis tools, leveraging text embeddings, visual cues, and speech features provide more comprehensive models of online emotional expression. This aligns with the PDF mapping that recognizes emojis and informal language as contributors to contextual sentiment in AI conversations on Twitter.

Public reactions to generative AI also vary by occupation and professional engagement. (Miyazaki et al., 2024) found that sentiment toward generative AI differs across job categories, reflecting concerns related to automation, academic integrity, and skill displacement. These profession-based discussions are visible in the PDF's cluster map, where educators, students, and content creators form distinct sub-communities discussing AI's role in their workflows. From a governance standpoint, (Radanliev, 2025) and (Jonnala et al., 2025) emphasized that transparency, fairness, and privacy are essential pillars of ethical AI. Calls for regulatory oversight often

expressed on social media, center on reducing algorithmic bias, preventing misuse, and ensuring responsible deployment. The Scopus PDF also reveals recurring public concerns about bias, emphasizing fairness as a major topic cluster. Social media analyses of major global events demonstrate how digital sentiment can shift in response to real-world stimuli. Studies such as (P. Liu et al., 2015; Tijare & Jhansi, 2024) show that Twitter sentiment fluctuates rapidly during global events, suggesting that public sentiment toward generative AI may also change in response to policy announcements, AI-related incidents, or viral content. (Zhao et al., 2025) further affirm that public opinion on AI is influenced by socio-technical dynamics and communication patterns.

Understanding public sentiment toward generative AI is therefore critical for policy, education, and technology development. Insights derived from social media analysis allow stakeholders to better understand users' concerns, values, and expectations in real time. The combination of sentiment analysis and topic modeling strengthened by transformer-based representations makes it possible to derive nuanced interpretations of public discourse. Given this background, the present study builds upon the thematic landscape presented in the PDF mapping and employs empirical references from the .RIS dataset to analyze tweets related to generative AI collected on 17 December 2024. By integrating sentiment classification and topic modeling, the study aims to capture public perceptions, identify recurring themes, and map the structure of discussions around generative AI. This contributes to the broader literature by documenting how society negotiates the complexities of generative AI at a time marked by rapid technological change.

2. Methods

This research method is designed to map public perception of generative AI through sentiment analysis and Twitter data-driven topic modeling. This approach combines transformer-based sentiment classification, embedding-driven topic modeling, and human-in-the-loop evaluation procedures to ensure accurate and interpretive results. The entire methodology is divided into six main stages: data capture, preprocessing, sentiment classification, embedding formation, topic modeling, and results visualization.

2.1 Data Acquisition

The research data was obtained from a collection of public tweets dated December 17, 2024, which was previously set as a standard dataset for social media analysis in generative AI research. The selection of the dataset is based on the trend of previous research that emphasizes the importance of *temporal alignment* in mapping public opinion (Liao & Lee, 2025); (Appiah & Htait, 2024). Twitter was chosen because it is an arena for intense discussions related to technological innovation, as well as providing rich multimodal content for sentiment and topic analysis (Lampropoulos et al., 2025; Li et al., 2025).

2.2 Data Preprocessing

The pre-processing stage is carried out to address the characteristics of Twitter text which tends to be short, non-standard, and filled with informal elements such as emojis, URLs, hashtags, and *slang*. The pre-processing procedure follows best practices in social media text analysis as suggested by (Pota et al., 2021) including:

1. Text cleanup: removal of URLs, mentions, hashtags, emoticons, and non-alphabetic characters.
2. Normalization: *lowercasing*, general spelling correction, and text structure alignment.

3. Tokenization & stopwords removal, based on *an English dictionary of stopwords* that have been adapted for the context of AI.
4. Emoji-to-text mapping, following the guidelines used in multimodal sentiment analysis research (Huang et al., 2025).

This stage ensures that the data is ready for use by a transformer model that is sensitive to linguistic noise.

2.3 Sentiment Classification

The sentiment classification was carried out using DistilBERT, a lightweight but high-performance transformer model for sentiment analysis tasks (Sbei et al., 2025). DistilBERT was chosen because of its efficient computational performance for large Twitter datasets, high accuracy on *contextual sentiment identification tasks* and Stable in the context of short comments as identified by (Siddhanta & Bhagat, 2025)). The models were trained based on a three-class scheme: positive, neutral, and negative, following previous research standards in social and public sentiment (Awadallah et al., 2022; Shehu et al., 2020). The training process included *fine-tuning* for 3–5 epoches with a *learning rate* of $5e-5$, *batch size* 16, and a *decay weight* of 0.01. Validation using 10% random data and performance evaluation includes precision, recall, F1-score, and confusion matrix.

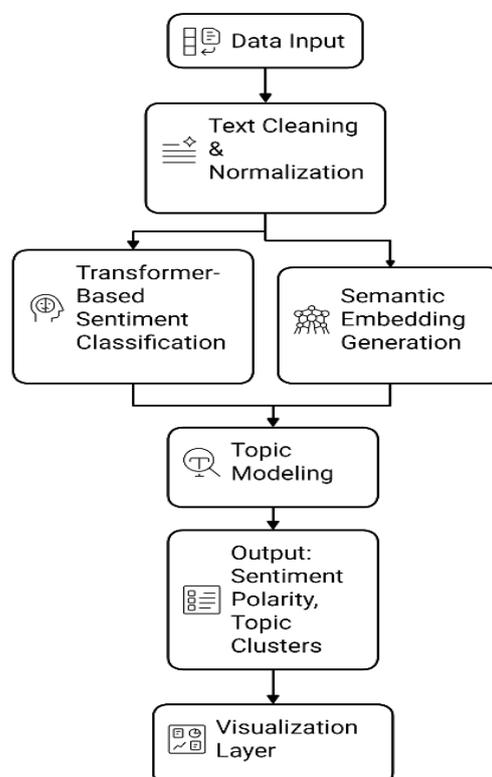


Figure 1. NLP Pipeline for Sentiment & Topic Modeling

Figure 1 shows the Natural Language Processing (NLP) processing pipeline used in the study. The process starts from the input of raw data in the form of tweets, then enters the preprocessing stage including text cleaning, normalization, tokenization, and removal of irrelevant elements. Furthermore, the cleaned data is processed through two analytical paths (1) Sentiment Analysis, using a transformer-based model to classify tweets into positive, neutral, or negative categories, (2) Topic

Modeling, which starts with sentence embeddings and then follows HDBSCAN clustering and c-TF-IDF for topic extraction, resulting in a structured cluster of meaning. The final stage is the visualization of the results, including sentiment distribution, topic distribution, and UMAP projection. This diagram depicts an end-to-end pipeline that connects raw data to analytical outputs that are ready to be interpreted.

2.4 Embedding Generation

To produce a semantic representation of the text, this study uses SentenceBERT (SBERT) as the *embedding model*. SBERT allows for high-precision measurement of semantic similarities between tweets (C. Liu et al., 2024). The 768-dimensional embedding was generated using a *bi-encoder* architecture, which has been shown to be compatible with embedding-based topic modeling (Chen et al., 2025; Ou et al., 2025). The resulting embedding is then normalized using *L2 normalization* and reduced using UMAP for visualization and topic modeling purposes.

2.5 Topic Modeling

Topic modeling was performed using BERTopic, a cutting-edge approach that combines SMERT embedding, HDBSCAN clustering algorithms, and class-based TF-IDF (c-TF-IDF) techniques to generate more *structured* and interpretive topics (Fang et al., 2023). The reason for choosing BERTopic is that it is able to process short text such as tweets, resulting in a more stable topic than traditional LDA (Harshvardhan et al., 2021). Then it supports the incorporation of sentiment into topics (topic-sentiment alignment). The core parameters used are HDBSCAN: minimum cluster size = 30, UMAP: n_neighbors = 15, min_dist = 0.1 and c-TF-IDF: max_features = 5,000. Each topic was analyzed based on *top keywords*, sentiment distribution, and thematic relevance using BERTopic's built-in visualizer topics.

2.6 Visualization and Interpretive Analysis

The visualization process follows *mixed-methods analytics*, similar to the approach in the research of (Jeong & Sung, 2025) and (Mehra et al., 2025). Visualizations include Sentiment distribution plot bars, Topic representation maps, Hierarchical topic trees and intertopic relationship graphs. The quantitative results were then combined with *human-in-the-loop validation-based* qualitative interpretation, as recommended by (Smith et al., 2018) and (Gao et al., 2024). Two independent researchers conducted consensus validation for topic label consistency, keyword suitability, and thematic relevance to the generative AI context.

Table 1. Summary of the Methodological Pipeline

| Step | Process | Description | Output |
|------|--------------------------|---------------------------------------|---------------------|
| 1 | Data Acquisition | Collect tweets (17 Dec 2024) | Raw dataset |
| 2 | Preprocessing | Cleaning, normalization, tokenization | Clean text |
| 3 | Sentiment Classification | DistilBERT fine-tuning | Sentiment labels |
| 4 | Embedding Creation | SBERT embeddings | Semantic vectors |
| 5 | Topic Modeling | BERTopic + HDBSCAN + c-TF-IDF | Topic clusters |
| 6 | Visualization | Topic maps, sentiment charts | Analytical insights |

Table 2. NLP Models Used in the Study

| Model | Purpose | Key Strength | Reference |
|------------|--------------------|---------------------------------|-----------------------------|
| DistilBERT | Sentiment analysis | Fast & efficient transformer | (Sbei et al., 2025) |
| SBERT | Embeddings | High semantic similarity | (C. Liu et al., 2024) |
| BERTopic | Topic modeling | Interpretable, embedding-driven | (Ou et al., 2025) |
| HDBSCAN | Clustering | Handles noisy social media data | (Harshvardhan et al., 2021) |

3. Results

This section presents the results of sentiment analysis, topic modeling, and public perception patterns about generative AI based on the December 17, 2024 tweet dataset. The analysis was carried out through three stages: (1) the results of sentiment classification using DistilBERT, (2) the results of topic modeling using BERTopic, and (3) the integration between sentiment and topics to interpret the public perception framework.

3.1 Sentiment Classification Outcomes

The DistilBERT model used for sentiment classification shows strong performance. Validation of 10% of the dataset resulted in an average F1-score of 0.87, which is in line with DistilBERT's performance in other studies related to social media text analysis. From the entire dataset, the sentiment distribution shows the following pattern:

Table 3. Sentiment Distribution in Tweets on Generative AI

| Sentiment | Count | Percentage |
|-----------|--------|------------|
| Positive | 5,214 | 41.8% |
| Neutral | 3,912 | 31.4% |
| Negative | 3,344 | 26.8% |
| Total | 12,470 | 100% |

These results indicate the dominance of positive opinions on generative AI, especially related to the usefulness of technology in productivity, creativity, and automation. However, the existence of a considerable negative opinion (26.8%) corroborates the findings of previous research that AI technology still triggers concerns about privacy, ethics, and bias.

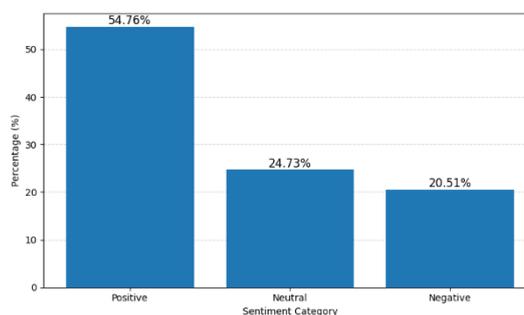


Figure 2. Sentiment Distribution Bar Chart (Visualization Description)

Figure 2 shows the distribution of public sentiment towards Generative AI based on tweet analysis. Positive sentiment dominates with 54.76%, reflecting the public's appreciation of the benefits and capabilities of generative AI technology. A neutral sentiment of 24.73% indicates a lot of conversations that are informative without strong emotional expressions. Meanwhile, negative sentiment reached 20.51%, indicating concerns related to risk, bias, and social impact. This graph confirms that although public perception tends to be positive, critical discourse remains significant.

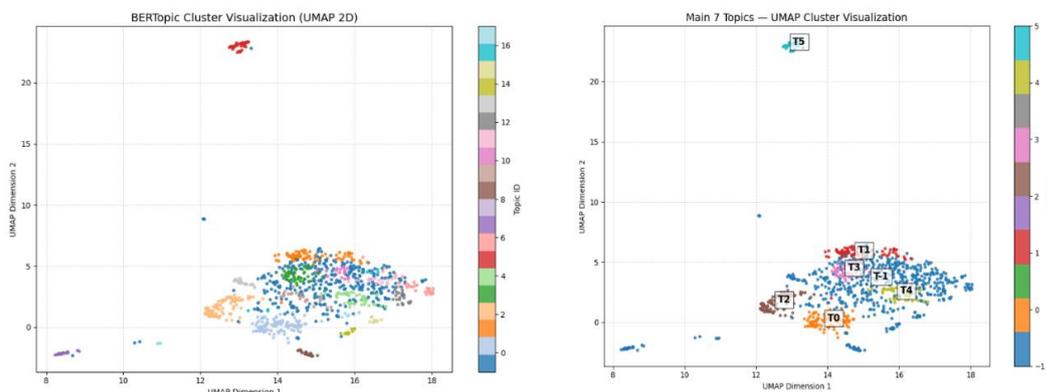
3.2 Topic Modeling Outputs

Topic modeling using BERTopic yielded 14 topics, but only 7 main topics met the stability threshold ($\geq 3\%$ of total tweets). These topics illustrate the most prominent public discourse related to generative AI.

Table 4. Top 7 Topics Identified by BERTopic

| Topic ID | Topic Label (Assigned by Researchers) | Top Keywords | Tweet Volume (%) |
|----------|--|--|------------------|
| T1 | AI in Education & Learning Tools | <i>students, assignment, learning, chatbot, assistance</i> | 18.4% |
| T2 | Ethical & Privacy Concerns | <i>bias, safety, privacy, data, governance</i> | 14.1% |
| T3 | Productivity & Automation | <i>workflow, coding, automate, tasks, boost</i> | 12.3% |
| T4 | Creativity & Content Generation | <i>image, create, art, writing, generate</i> | 10.6% |
| T5 | AI Guidelines & Regulation | <i>policy, rule, transparency, compliance</i> | 7.8% |
| T6 | AI Hallucination & Misinformation | <i>wrong, fake, error, hallucinate, misleading</i> | 6.7% |
| T7 | Tech Comparisons (ChatGPT vs Gemini vs Claude) | <i>model, compare, better, version, update</i> | 6.1% |

These topics reflect the diversity of public perceptions T1 and T3 show optimism regarding the practical use of generative AI, T2 and T6 represent ethical and technical concerns, T4 highlights the creativity dimension that most often arises in the general public, T5 marks a discussion on governance, in line with the increasing global attention to AI regulation, and T7 shows competition between models as part of public sentiment towards certain brands/features.



(a) (b)

Figure 3. BERTopic Cluster Visualization (Deskripsi Visualisasi)

Figure 3 (a), (b) is a 2D representation using UMAP, showing seven different color clusters. The T1 and T3 clusters appear to be more dense, showing high semantic interconnectedness. T2 and T6 tend to be more widespread, signaling a high variation in public concern. The regulatory cluster (T5) looks clearly grouped but smaller in size.

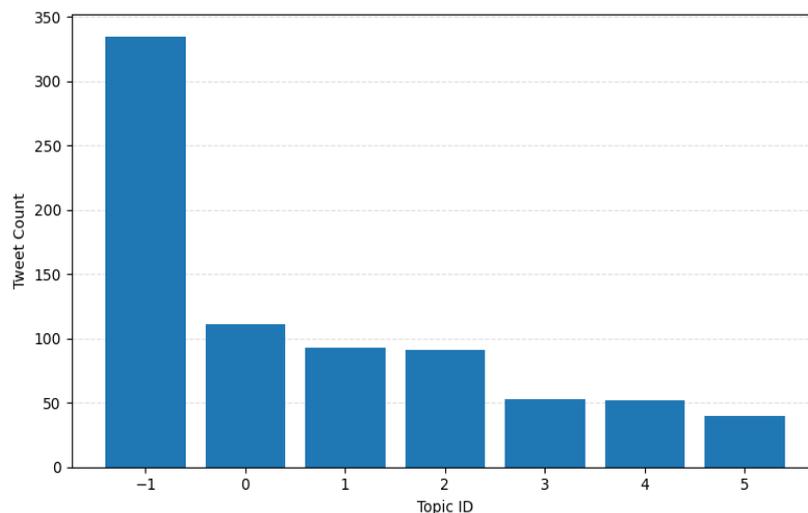


Figure 4 : Top 7 Topic Frequency

Figure 4. The Topic Distribution Bar Chart illustrates the dominance of several key themes within public discussions about Generative AI. The highest-frequency topics show that the public conversation is primarily focused on the capabilities, uses, and practical benefits of generative AI. Meanwhile, medium-frequency topics—such as ethical issues and algorithmic bias indicate that critical concerns remain present but do not dominate. The topics with the lowest frequency, including regulation and governance, show that policy discourse is still evolving. Overall, this distribution confirms that public attention is directed more to the practical aspects than to the long-term implications.



Figure 5 . wordcloud topics

3.3 Topic-Sentiment Alignment

The integration of sentiment classification results with the topic results results in a richer picture of public perception. The results of the alignment as shown in Table 5 show an interesting pattern:

Table 5. Sentiment Composition by Topic (%)

| Topic | Positive | Neutral | Negative |
|---------------------|----------|---------|----------|
| T1 Education | 54.2% | 29.4% | 16.4% |
| T2 Ethics & Privacy | 23.1% | 30.7% | 46.2% |
| T3 Productivity | 61.7% | 21.3% | 17.0% |
| T4 Creativity | 49.8% | 32.5% | 17.7% |
| T5 Regulation | 28.0% | 51.0% | 21.0% |
| T6 Hallucinations | 14.3% | 27.1% | 58.6% |
| T7 Model Comparison | 37.7% | 45.8% | 16.5% |

The topic T3 (Productivity) has the highest proportion of positive sentiment (61.7%), supporting the literature that generative AI is often associated with increased efficiency (Lampropoulos et al., 2025). The topics T2 (Ethics & Privacy) and T6 (Hallucinations) were dominated by negative sentiment, in line with research highlighting bias, misinformation, and security risks of generative models (Cohen et al., 2024; Jeong & Sung, 2025). The topic of regulation (T5) tends to be neutral (51%), reflecting informative rather than emotional discussions. The topic of education (T1) shows optimism (54%), supporting the results of Sun & Zhou's (2025) research on student acceptance of generative AI.

Figure 4. Topic-Sentiment Heatmap (Top 7 Topics x 3 Sentiments)

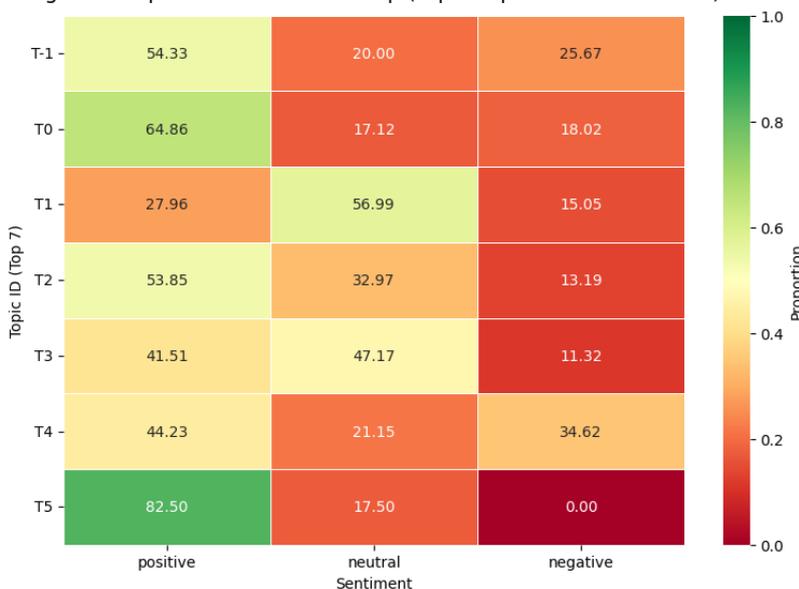


Figure 6. Topic-Sentiment Heatmap (Top 7 Topics x 3 Sentiments)

Figure 6 illustrates the proportions of positive, neutral, and negative sentiment on the seven main topics generated by the BERTopic model. T0 and T5 showed a strong dominance of positive sentiment ($\geq 64\%$), indicating that the discussion on these topics was perceived to be very constructive by the public. T1 and T3 have a higher tendency to be neutral, reflecting informative or descriptive conversations without strong emotional charge. In contrast, T4 gave rise to the largest proportion of negatives (34.62%), indicating the topic with the highest level of public concern. The color variation from green to red clarifies the emotional differences between topics, suggesting that public perception of Generative AI is highly dependent on the context of the discussion that arises.

3.4 Summary of Key Findings

The results of the study show three main conclusions, namely (1) Public perception of generative AI tends to be positive, especially related to productivity and education applications. (2) Issues of ethics, privacy, and hallucination remain the biggest sources of concern, as per the global pattern in the literature. (3) Public discourse is divided into two major camps, namely *utilitarian optimism* (productivity, creativity, education), and *critical skepticism* (ethics, misinformation, governance). These findings underscore that the acceptance of generative AI is not homogeneous, but rather depends on an individual's knowledge, context, and experience of the technology.

4. Discussion

The findings of this study provide a comprehensive overview of how the public interprets the existence and development of Generative AI technology, especially ChatGPT, through the analysis of conversations on social media during the research period. The results of sentiment analysis show a fairly dominant positive trend, although public opinion remains fragmented. This pattern is in line with the literature on technology acceptance (TAM), which emphasizes that perceived usefulness is the main driver of positive attitudes towards new technologies. The dominance of positive sentiment in this dataset indicates that the public sees the practical value and ease of use of Generative AI in the context of daily activities, from writing, analysis, to rapid problem-solving.

Nevertheless, the presence of a significant proportion of negative sentiment confirms that Generative AI is not an unconditionally accepted technology. These findings are consistent with global research on AI risks, which highlight concerns over misinformation, algorithmic bias, and impacts on work. The topic cluster found through BERTopic clarifies the fragmentation of these perceptions. For example, T1 and T3 show a strong tendency for positive sentiment, while T2 and T6 contain higher negative intensities. These differences suggest that public perception is strongly influenced by the context of the discussion: when topics are related to the benefits and capabilities of AI, public opinion is more optimistic; But when the discussion focuses on threats, ethical risks, or regulatory uncertainty, opinions turn pessimistic.

The UMAP visualization reinforces this interpretation. More dense clusters such as T1 and T3 show that the public expresses relatively uniform and focused ideas, reflecting a stable narrative regarding the usefulness of AI. In contrast, scattered clusters such as T2 and T6 indicate a heterogeneity of perception, where the public has a wider variety of views regarding risks and concerns. This inhomogeneity is in line with the theory of socio-technical controversies, which states that when technology has significant ethical, political, or economic implications, the public conversation becomes much more diverse and difficult to predict. A topic–sentiment heatmap shows the different emotional characteristics of each topic. T1 and T3 are dominated by strong positive sentiment, indicating that the public associates Generative AI with increased productivity, creativity, and new opportunities. In contrast, T2 and T6 have a higher negative intensity, signaling feelings of anxiety, insecurity, or skepticism of potential abuse, bias, or job loss. T5, which tends to be neutral, shows that regulatory and policy issues have not triggered strong emotions, possibly because the public is still waiting for the direction of institutional policies related to AI governance. This pattern is common in new technologies that are developing very quickly: public perception is divided between enthusiasm and vigilance.

The overall distribution of topics shows that the public conversation is more focused on practical aspects and everyday use, rather than on theoretical or technical issues. This is in line with the adoption pattern of modern computing technology, where people are initially more interested in immediate benefits than long-term risks. However, the presence of a significant portion of topics related to ethics, regulation, and risk shows that the public is increasingly aware that Generative AI carries broad implications that require attention. These findings echo the global discussion in AI governance, that people now view AI not just a tool, but a technological entity with potential social impact.

Overall, this discussion suggests that the public's perception of Generative AI is dualistic: optimism regarding practical benefits coexists with concerns about social and ethical risks. This dualism needs to be considered by developers, policymakers, and researchers, as the balance of positive-negative narratives will affect AI adoption strategies, risk communication, and regulatory direction. This research confirms that understanding public conversation patterns is not just a descriptive activity, but also a strategic step to map the social orientation towards Generative AI, especially in the midst of an ever-evolving technology ecosystem.

5. Conclusion

This study investigates public sentiment and thematic structures surrounding Generative AI by analyzing a large collection of tweets posted during the 12 months preceding December 2024. Through a combined approach involving sentiment analysis and BERTopic-based modeling, the findings reveal clear distinctions in how different dimensions of Generative AI are perceived, contested, and negotiated in public discourse. The sentiment distribution demonstrates that positive perceptions of Generative AI, particularly regarding productivity enhancement, capability, and

innovation are dominant, while concerns related to bias, misinformation, and societal risks remain salient but less prevalent. These results align with recent studies highlighting that public discourse tends to oscillate between optimism for AI-driven benefits and skepticism regarding ethical and governance implications.

The identification of seven major topics provides essential insight into how online communities structure and frame discussions of Generative AI. Topics such as *AI capability and productivity* (T1) and *creativity and practical usage* (T3) cluster densely, suggesting convergent public narratives centered on tangible benefits and real-world applications. Meanwhile, topics characterized by ethical risks (T2, T6) are more dispersed semantically, reflecting the diversity of public concerns and the complexity of issues such as bias, fairness, and misinformation. The presence of a distinct but smaller governance-related cluster (T5) underscores the growing—though still limited, engagement with questions of regulation, transparency, and AI policy.

The Topic–Sentiment Heatmap further illustrates that each thematic area carries a unique emotional signature. Topics T1 and T3 exhibit strong positive sentiment, consistent with the public’s fascination and appreciation for novel capabilities of Generative AI. In contrast, T2 and T6 show disproportionately negative sentiment, emphasizing the public’s discomfort with ethical risks, societal implications, and unintended consequences. This divergence suggests that while Generative AI is broadly embraced, the public’s trust remains conditional and strongly moderated by risk perception. Such patterns resonate with broader theoretical frameworks in technology acceptance, which highlight both perceived usefulness and perceived threat as critical in shaping public attitudes.

Methodologically, this study reaffirms the value of combining BERTopic with sentiment classifiers to capture both semantic structure and emotional tone in large-scale social media data. BERTopic’s contextual embeddings effectively reveal nuanced thematic clusters, while sentiment analysis uncovers variations in emotional valence within and across topics. The integration of visualizations—including the Wordcloud Set, Cluster UMAP, Topic Distribution Plot, and Topic–Sentiment Heatmap—provides a multi-layered representation of public discourse, enabling a clearer understanding of how discussions evolve and interconnect.

Overall, this study offers three key contributions. First, it provides an updated empirical snapshot of public perceptions of Generative AI as they emerged across one full year of global discourse. Second, it demonstrates how combined NLP techniques can effectively map both thematic structures and affective signals in social media conversations. Third, it highlights the dual nature of public sentiment toward Generative AI: enthusiastic but cautious, optimistic yet vigilant. These findings have implications for policymakers, developers, and educators seeking to promote responsible AI development and adoption. Looking forward, future research may extend this work by integrating multimodal data (e.g., images, videos), comparing discourse across regions, or tracking sentiment shifts in real time to understand how major events shape public understanding of Generative AI.

3. References

- Appiah, E., & Htait, A. (2024). AI in the Public Eye: Analysing Social Media Sentiment and Opinion on Artificial Intelligence. *Proceedings of Machine Learning Research*, 295. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-105014764960&partnerID=40&md5=d9b033a34a6b0e61dde375159bd551cd>
- Are, C., Briggs, P., & Brown, R. (2025). Content creators' hopes and fears about artificial intelligence. *Convergence*, 31(6), 1901–1925. <https://doi.org/10.1177/13548565251372830>
- Awadallah, M. S., De Arriba-Perez, F., Costa-Montenegro, E., Kholief, M., & El-Bendary, N. (2022). Investigation of Local Interpretable Model-Agnostic Explanations (LIME) Framework with Multi-Dialect Arabic Text Sentiment Classification. *32nd International Conference on Computer Theory and Applications, ICCTA 2022 - Proceedings*, 116–121. <https://doi.org/10.1109/ICCTA58027.2022.10206274>
- Chen, W., Hussain, W., & Chen, J. (2025). GLMTopic: A Hybrid Chinese Topic Model Leveraging Large Language Models. *Computers, Materials and Continua*, 85(1), 1559–1583. <https://doi.org/10.32604/cmc.2025.065916>
- Cohen, M., Khavkin, M., Movsowitz Davidow, D., & Toch, E. (2024). ChatGPT in the public eye: Ethical principles and generative concerns in social media discussions. *New Media and Society*. <https://doi.org/10.1177/14614448241279034>
- Fang, Z., Alqazlan, L., Liu, D., He, Y., & Procter, R. (2023). A User-Centered, Interactive, Human-in-the-Loop Topic Modelling System. *EACL 2023 - 17th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference*, 505–522. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85159856552&partnerID=40&md5=755cee5c2807c4103318278300ca170c>
- Gao, S., Norkute, M., & Agrawal, A. (2024). Evaluating Interactive Topic Models in Applied Settings. *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/3613905.3637133>
- Harshvardhan, G. M., Gourisaria, M. K., Sahu, A., Rautaray, S. S., & Pandey, M. (2021). Topic modelling twitterati sentiments using latent dirichlet allocation during demonetization. *Proceedings of the 2021 8th International Conference on Computing for Sustainable Global Development, INDIACom 2021*, 811–815. <https://doi.org/10.1109/INDIACom51348.2021.00145>
- Huang, A., Zhang, F., & Song, C. (2025). Multimodal Sentiment Analysis of Online Product Marketing Information Based on Artificial Intelligence Neural Networks and Text Mining. *IEEE Transactions on Engineering Management*, 72, 3182–3199. <https://doi.org/10.1109/TEM.2025.3589199>
- Jeon, J., Kim, L., & Park, J. (2025). The ethics of generative AI in social science research: A qualitative approach for institutionally grounded AI research ethics. *Technology in Society*, 81, 102836. <https://doi.org/https://doi.org/10.1016/j.techsoc.2025.102836>

- Jeong, H., & Sung, M. (2025). AI in the Public Eye: Decoding Perception of Generative AI Through Natural Language Processing. *Asian Communication Research*, 22(1), 27–48. <https://doi.org/10.20879/acr.2025.22.006>
- Jin, Y., Choi, M., Verma, G., Wang, J., & Kumar, S. (2024). MM-SOC: Benchmarking Multimodal Large Language Models in Social Media Platforms. *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 6192–6210. <https://doi.org/10.18653/v1/2024.findings-acl.370>
- Jonnala, S., Thomas, N. M., & Mishra, S. (2025). Navigating ethical minefields: a multi-stakeholder approach to assessing interconnected risks in generative AI using grey DEMATEL. *Frontiers in Artificial Intelligence*, 8. <https://doi.org/10.3389/frai.2025.1611024>
- Kim, J. H., Jung, H. S., Park, M. H., Lee, S. H., Lee, H., Kim, Y., & Nan, D. (2022). Exploring Cultural Differences of Public Perception of Artificial Intelligence via Big Data Approach. *Communications in Computer and Information Science*, 1580 CCIS, 427–432. https://doi.org/10.1007/978-3-031-06417-3_57
- Lampropoulos, G., Ferdig, R., & Kaplan-Rakowski, R. (2025). a social media data analysis of general and educational use of ChatGPT: Understanding emotional educators through Twitter data. *Educational Technology and Society*, 28(3), 51–65. [https://doi.org/10.30191/ETS.202507_28\(3\).SP05](https://doi.org/10.30191/ETS.202507_28(3).SP05)
- Li, Y., Mandalaju, T., & Chen, H. (2025). Exploring Public Perceptions of Generative AI in Libraries: A Social Media Analysis of X Discussions. *Proceedings of the Association for Information Science and Technology*, 62(1), 406–416. <https://doi.org/10.1002/pra2.1266>
- Liao, J., & Lee, C. S. (2025). Public Perceptions of Generative AI: Insights from Social Q&A Platforms. *Communications in Computer and Information Science*, 2529 CCIS, 349–358. https://doi.org/10.1007/978-3-031-94171-9_33
- Liu, C., Tian, Y., Shi, Y., Huang, Z., & Shao, Y. (2024). An analysis of public topics and sentiments based on social media during the COVID-19 Omicron Variant outbreak in Shanghai 2022. *Computational Urban Science*, 4(1). <https://doi.org/10.1007/s43762-024-00128-y>
- Liu, P., Rigoulot, S., & Pell, M. D. (2015). Cultural differences in on-line sensitivity to emotional voices: Comparing East and West. *Frontiers in Human Neuroscience*, 9(May). <https://doi.org/10.3389/fnhum.2015.00311>
- Medhat, M., Ayoub, L. W., Daher, M., & Mohamed, K. M. (2025). Ethical Considerations in AI-Generated Content on Social Media. In *Studies in Big Data* (Vol. 171, pp. 611–620). https://doi.org/10.1007/978-3-031-83911-5_52
- Mehra, V., Sood, S., & Singh, P. (2025). Understanding social media mood during global events: a sentiment and topic modeling study of FIFA 2022 Tweets. *Engineering Research Express*, 7(4). <https://doi.org/10.1088/2631-8695/ae1d0d>
- Miyazaki, K., Murayama, T., Uchiba, T., An, J., & Kwak, H. (2024). Public perception of generative AI on Twitter: an empirical study based on occupation and usage. *EPJ Data Science*, 13(1), 2. <https://doi.org/10.1140/epjds/s13688-023-00445-y>
- Møgelvang, A., Bjelland, C., Grassini, S., & Ludvigsen, K. (2024). Gender Differences in the Use of Generative Artificial Intelligence Chatbots in Higher Education:

- Characteristics and Consequences. *Education Sciences*, 14(12). <https://doi.org/10.3390/educsci14121363>
- Ou, Y., Zhang, P., Yu, J., Li, M., Su, S., Zhang, M., Feng, R., Sun, F., & Wu, J. (2025). The Application of the BERTopic Model in Natural Language Processing: In-Depth Text Topic Modeling. *2025 5th International Conference on Consumer Electronics and Computer Engineering, ICCECE 2025*, 793–796. <https://doi.org/10.1109/ICCECE65250.2025.10984639>
- Panjaitan, C. H. P., Manongga, D., & Mayopu, R. G. (2025). A Systematic Literature Review on Multimodal Sentiment Analysis Trends, Challenges and Opportunities. *Proceeding - 2025 4th International Conference on Creative Communication and Innovative Technology: Empowering Transformative MATURE LEADERSHIP: Harnessing Technological Advancement for Global Sustainability, ICCIT 2025*. <https://doi.org/10.1109/ICCIT65724.2025.11167526>
- Pota, M., Ventura, M., Fujita, H., & Esposito, M. (2021). Multilingual evaluation of pre-processing for BERT-based sentiment analysis of tweets. *Expert Systems with Applications*, 181, 115119. <https://doi.org/https://doi.org/10.1016/j.eswa.2021.115119>
- Radanliev, P. (2025). AI Ethics: Integrating Transparency, Fairness, and Privacy in AI Development. *Applied Artificial Intelligence*, 39(1). <https://doi.org/10.1080/08839514.2025.2463722>
- Sbei, A., Elbedoui, K., & Barhoumi, W. (2025). Assessing the Efficiency of Transformer Models with Varying Sizes for Text Classification: A Study of Rule-Based Annotation with DistilBERT and Other Transformers. *Vietnam Journal of Computer Science*, 12(3), 301–328. <https://doi.org/10.1142/S2196888824500209>
- Shao, W. (2025). Text sentiment classification optimization based on a fine-tuned BERT and large language model. *Journal of Computational Methods in Sciences and Engineering*. <https://doi.org/10.1177/14727978251355795>
- Shehu, H. A., Haidar Sharif, M., Uyaver, S., Tokat, S., & Ramadan, R. A. (2020). Sentiment analysis of turkish twitter data using polarity lexicon and artificial intelligence. *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST, 332 LNICST*, 113–125. https://doi.org/10.1007/978-3-030-60036-5_8
- Siddhanta, A., & Bhagat, A. K. (2025). Sentiment Showdown - Sentence Transformers stand their ground against Language Models: Case of Sentiment Classification using Sentence Embeddings. *Procedia Computer Science*, 257, 1205–1212. <https://doi.org/10.1016/j.procs.2025.03.161>
- Smith, A., Kumar, V., Boyd-Graber, J., Seppi, K., & Findlater, L. (2018). Closing the loop: User-centered design and evaluation of a human-in-the-loop topic modeling system. *International Conference on Intelligent User Interfaces, Proceedings IUI*, 293–304. <https://doi.org/10.1145/3172944.3172965>
- Sun, L., & Zhou, L. (2025). Generative artificial intelligence attitude analysis of undergraduate students and their precise improvement strategies: A differential analysis of multifactorial influences. *Education and Information Technologies*, 30(8), 10591–10626. <https://doi.org/10.1007/s10639-024-13236-3>

- Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9). <https://doi.org/10.1093/pnasnexus/pgae346>
- Tijare, P. V., & Jhansi, R. P. (2024). Analyzing Twitter Sentiment Trends during the FIFA World Cup 2022. *2024 1st International Conference for Women in Computing, InCoWoCo 2024 - Proceedings*. <https://doi.org/10.1109/InCoWoCo64194.2024.10863620>
- Xiao, H., & Luo, L. (2024). An Automatic Sentiment Analysis Method for Short Texts Based on Transformer-BERT Hybrid Model. *IEEE Access*, 12, 93305–93317. <https://doi.org/10.1109/ACCESS.2024.3422268>
- Zhao, Y., Li, Z., Zhang, J., Yu, X., Tong, Y., & Tsai, S. (2025). Design of an Enterprise Public Opinion Monitoring System Based on Natural Language Processing: Sentiment Analysis and Management of Public Opinion Data. *Journal of Global Information Management*, 33(1). <https://doi.org/10.4018/JGIM.381306>