AI-Sustainability Discourse Across Social Media: Sentiment, Topics, and Engagement on LinkedIn, X, and Bluesky

Kholoud Khalil Aldous[1], Md. Rafiul Biswas[2], Mabrouka Bessghaier[1]
Shimaa Ibrahim[1], Kais Attia[3], Wajdi Zaghouani[1]
[1]Northwestern University in Qatar, Qatar,
[2]Hamad Bin Khalifa University, Qatar
[3]Independent Researcher
{wajdi.zaghouani,kholoud.aldous}@northwestern.edu
{mabrouka.bessghaier,shimaa.ibrahim}@northwestern.edu
mbiswas@hbku.edu.qa, kais.attia.w@gmail.com

Abstract

We present the first cross-platform computational analysis of AI-sustainability discourse across LinkedIn, X, and Bluesky, using 52,102 posts collected over 15 months. Combining RoBERTa-based sentiment classification with BERTopic topic modeling, we identify 36 topics and compare sentiment, topic distributions, and engagement across platforms. Overall, 43.98% of posts are positive, 40.16% neutral, and 15.86% negative. Sentiment varies by topic rather than by platform: corporate and AI-potential themes are mostly positive, while resource-cost topics carry the highest negative sentiment. Topic distributions differ significantly across platforms: LinkedIn focuses on corporate sustainability, X on hardware and cryptocurrency, and Bluesky on environmental resource costs. Niche topics consistently attract more engagement than high-volume ones. These findings have implications for how organizations and policymakers communicate about AI-sustainability, as topic framing and platform audience composition shape public discourse more than platform selection alone. We release the annotated dataset to support future research.

1 Introduction

Artificial intelligence (AI), including generative AI (GenAI), is a major force shaping decision-making, organizations, and everyday life (Maslej et al., 2024; Polak and Anshari, 2024). In 2024, 78% of organizations reported the usage of AI, which is 23% higher than 2023 usage (55%) (Maslej et al., 2025). At the same time, rising concern about climate change has drawn attention to the environmental impact of digital technologies (Patterson et al., 2021; Strubell et al., 2019). The Information and Communication Technology sector accounts for 1.5-4% of global greenhouse gas emissions, comparable to the aviation industry, with data centers and transmission networks consuming 1-2% of the world's total electricity (Bieser et al., 2023). Training and deploying AI systems require significant energy and infrastructure, leading to growing debates about sustainability (Rillig et al., 2023; Liu and Yin, 2024). Particularly, AI systems shape climate discourse (Biswas et al., 2026), with recent work showing how large language models reflect particular opinions and framings about environmental issues (Locatelli et al., 2025; Zaghouani et al., 2026). This creates a recursive dynamic: AI both discusses and embodies sustainability challenges. Understanding how these concerns manifest in public discourse across social media platforms is essential for assessing public attitudes, understanding levels of concern, and informing responsible AI governance.

Social media platforms play a central role in shaping this discussion. They reflect public opinion and influence it through algorithms, design features, and user communities (Boyd and Ellison, 2007; Cinelli et al., 2021). Computational social science research has used social media data to explore public views on emerging technologies, applying methods such as sentiment analysis and topic modeling (Lazer et al., 2009; Tsur and Rappoport, 2012). Prior studies on AI discourse demonstrate both optimism about innovation and efficiency and concerns about ethics, labor, and governance (Fast and Horvitz, 2017; Cave et al., 2019). Concerns about climate change and related debates increasingly appear on social media, where longitudinal analyses show persistent themes and public concerns (Janaswamy and Blackburn, 2024). However, research examining AI-sustainability discourse often focuses on individual platforms (Wang et al., 2025), leaving unexplored how such conversations differ across platform ecologies. Platform differences can produce parallel and conflicting narratives (Madsen, 2022; Klinger et al., 2023), and newer platforms such as Bluesky further diversify this landscape by attracting different user communities with different norms and priorities.

Findings of the Association for Computational Linguistics: IJCNLP-AACL 2026, pages 3004–3018
November 6-10, 2026 ©2026 Association for Computational Linguistics

In this work, we focus on AI-sustainability: the intersection of environmental and sustainability concerns with artificial intelligence systems, including how users discuss the energy consumption, carbon footprint, and ecological impact of AI training and deployment. We analyze sentiment, topics, and engagement of 52,102 posts to understand how platforms shape public discourse in this domain. Specifically, we explore the following research questions:

- **RQ1:** How does sentiment toward AI-sustainability differ across platforms?
- **RQ2:** (a) What topics emerge around AI-sustainability? (b) Do these topics differ across social media platforms? (c) What is the sentiment associated with each topic?
- **RQ3:** Which topics generate the highest engagement, measured by likes, in AI sustainability discussions across platforms?

Our contributions are: (1) the first cross-platform computational analysis of AI sustainability discourse across LinkedIn, X, and Bluesky; (2) evidence that sentiment is fundamentally topic-dependent rather than platform-dependent, with a structural divide between infrastructure-cost topics (e.g., *AI Water Consumption Footprint*, *AI Energy Power Consumption*), which carry substantially higher negative sentiment (48.4% and 47.4%, respectively), and AI-potential topics, which remain predominantly positive; and (3) a dataset of 52,102 posts labelled with topic assignments and sentiment scores, made available for future research.

2 Related Work

Sentiment Analysis. A previous study by Miyazaki et al. (2024) analyzes a five-year dataset (2019–2023) of GenAI posts on X, showing that public discussions of GenAI are led by positive sentiment. Other Reddit-based studies show that sentiment changes differently within platform subgroups (Xu et al., 2024; Mentzer et al., 2024). Extending beyond single platforms, Appiah and Htait (2025) analyze 133K posts from X, YouTube, and Reddit (2017–2024), finding 46.09% negative, 39.29% neutral, and 14.62% positive sentiments, though without comparing distributions across platforms. Wang et al. (2025) analyze 104K X posts about data center energy use, showing an 87.4% increase in public attention from late 2022 and a dominant shift from positive to negative sentiment. Most directly related to our study, Longo and Longo (2025) analyze Reddit communities' views on AI and the environment, showing that community context significantly changes the emotional tone. Building on this, our work examines how public sentiment toward AI-sustainability differs across platforms through cross-platform comparison (RQ1).

Topical Analysis. Previous work shows that identical topics spread and are framed differently across platforms due to their structures and user cultures. Cinelli et al. (2020) find that identical posts distribute differently on X, Instagram, YouTube, and Reddit. Zannettou et al. (2017) track how topics move between platforms and communities. Classic affordance studies (Treem and Leonardi, 2013) link these differences to features such as visibility, persistence, and post connectivity. Recent work by Kloo et al. (2024) shows that similar topics are expressed differently depending on platform affordances. Xian et al. (2024) combine BERTopic with sentiment analysis to track temporal patterns in large-scale LLM discourse across global news sources. Aldous et al. (2019) further find that topic distribution of news content vary significantly across Facebook, Instagram, Twitter, YouTube, and Reddit. Our study extends this research to social media discussions of AI-Sustainability Discourse (RQ2).

User engagement. User engagement on social media depends on both content and platform features. Prior work shows that morally framed topics increase engagement (Berger and Milkman, 2012; Stieglitz and Dang-Xuan, 2013), and on LinkedIn, hashtags, mentions, post age, and follower count significantly predict reactions (Usera and Durham, 2025). Aldous et al. (2023) further show that engagement levels vary by both topic and platform. Cross-platform research shows distinct engagement dynamics (Cinelli et al., 2020), and newer platforms such as Bluesky, can match or exceed X in engagement for specific topics (Zheng et al., 2025; Arroyo-Machado et al., 2025). We compare user engagement per topic to identify which AI-sustainability topics attract the most interaction (RQ3).

3 Data Collection

We developed an English keyword search strategy using 36 terms in two categories. **Category 1 — AI Models and Technologies (15 terms):** "generative ai", "genai", "llm", "large language model", "gpt-3", "gpt-4", "claude", "gemini", "llama", "bert", "transformer model", "foundation model", "large models", "big ai model", "ai model training". **Category 2 — Environmental and Sustainability Metrics (21 terms):** "sustainability", "environmental impact", "carbon footprint", "carbon emissions", "carbon impact", "energy consumption", "energy use", "energy efficiency", "energy cost", "energy optimization", "energy reduction", "energy saving", "energy monitoring", "energy aware", "power consumption", "data center energy", "sustainable ai", "sustainable computing", "green ai", "eco-friendly ai", "low-carbon ai". Posts required $\geq$1 term from each category, resulting in 315 query combinations (15 $\times$ 21 keywords).

We used Meltwater (https://www.meltwater.com/), an enterprise media monitoring and social listening platform, to collect our data. Meltwater data has been used in previous research (Guo et al., 2023). We used a case-insensitive boolean search query combining keywords from both categories using AND logic. In total, we collected 42,769 social media datapoints from three social media platforms: LinkedIn (n=20,574), X (n=20,263), and Bluesky (n=1,932). These platforms were selected to capture different professional, mainstream, and emerging discourse around AI-sustainability. The collection period covered 15 months, from August 5, 2024 to November 5, 2025, inclusive.

To address the limited Bluesky coverage in the original Meltwater dataset (n=1,932), data collection was supplemented using the official AT Protocol API (Bluesky Social, 2024), specifically the `app.bsky.feed.searchPosts` endpoint, querying each of the 315 keyword combinations with the identical Boolean AND logic applied in Meltwater and paginating until no further results were returned for the same 15 months. Rate limits were respected using exponential backoff. This supplementary collection retrieved 14,665 unique posts not previously captured. Following deduplication across all collection rounds, the consolidated Bluesky corpus comprised 16,597 posts, and the dataset totaled **57,434** posts (LinkedIn: 20,574; X: 20,263; Bluesky: 16,597).

4 Methods

Sentiment Analysis. We used a RoBERTa-based sentiment classifier fine-tuned on social media text (Loureiro et al., 2022). Specifically, we applied the "cardiffnlp/twitter-roberta-base-sentiment-latest" model (Barbieri et al., 2020), which classifies content as *positive*, *neutral*, or *negative* and returns a confidence score for each label. This model was selected over lexicon-based approaches because transformer-based models have been shown to substantially outperform rule-based methods on short social media text (**?**); limitations related to cross-platform domain transfer are discussed in Section 8.

Topical Modeling. We modeled topics across all platforms together so that the same topic labels could be used to compare platforms directly. We applied BERTopic (Grootendorst, 2022) with `all-mpnet-base-v2` embeddings, UMAP dimensionality reduction (n_neighbors $=$ 15, cosine metric), and HDBSCAN clustering (min_cluster_size $=$ 260, $\approx$0.5% of corpus). Following initial clustering, 48.8% of documents were assigned to the outlier class (topic -1) and reassigned using BERTopic's `reduce_outliers` function with the `''embeddings''` strategy (`threshold=0.0`), which assigns each outlier to its nearest topic by cosine similarity between document and topic centroid embeddings. The final model contains **36 distinct topics** (CV coherence $=$ 0.66, topic diversity $=$ 0.91). As BERTopic outputs are sensitive to embedding choice and hyperparameter settings (Grootendorst, 2022), full configuration details are reported in Appendix A to support replication and alternative parameterisations. Topic quality was validated through author review of terms and representative documents.

A Pearson chi-squared test was applied to the 3 $\times$ 36 platform-by-topic contingency table to assess whether topic distributions differ significantly across platforms. One zero cell (*GenAI Carbon Footprint Concerns* on LinkedIn) was replaced with 0.5 prior to testing, as this topic appeared exclusively on Bluesky and X.

To improve the interpretability of the 36 topics, we used Claude Sonnet 4.6 to generate concise human-readable topic labels. For each topic, the model was prompted with its three representative documents and top 20 c-TF-IDF keywords using the following prompt: *"I have a topic that contains the following documents: [DOCUMENTS].*

The topic is described by the following keywords: [KEYWORDS]. Extract a short but highly descriptive topic label of at most 5 words. Format: topic: "<topic label>." The generated labels were reviewed and verified by the authors for accuracy and consistency. This approach follows recent practice of using large language models to automate the labelling of BERTopic topic outputs (Kozlowski et al., 2024).

Engagement Metrics. Due to the limited cross-platform engagement metrics available during data collection, and because each platform calculates engagement differently, we use the platform-native like signal as the primary engagement metric: Reactions for LinkedIn (functionally equivalent to likes) and Likes for X and Bluesky. Table 1 shows the descriptive statistics. LinkedIn generates the highest total engagement (495,041 Reactions; median = 7.0), followed by X (411,142 Likes; median = 1.0) and Bluesky (123,698 Likes; median = 1.0). The distributions are heavily right-skewed, with most posts receiving few likes and a small number of posts having high engagement. For RQ3, we compute the median likes per post within each platform and assess topic differences using the Kruskal-Wallis test (Cha et al., 2010). Engagement metrics are not normalised by follower count; raw counts may partly reflect account reach.

Initial Data Cleaning. The raw dataset comprised 57,434 posts (LinkedIn: 20,574; X: 20,263; Bluesky: 16,597). Table 2 summarises each cleaning step. Duplicates were removed by URL (568 posts) and by normalized text content (2,948 posts). Only English posts were retained. Posts with fewer than ten characters were excluded (108 posts). We removed posts published within five days of the collection end date (259 posts), setting the effective window to August 5, 2024 – October 31, 2025. Posts published within five days of the collection end date were removed because their engagement counts (likes, reactions) would still be accumulating at the time of collection, producing artificially low values relative to posts collected earlier in the window and biasing topic-level engagement comparisons. Accounts exceeding 1% of platform post volume were removed as bots (Varol et al., 2017): `@grok` on X (853 posts) and `bigearthdata.ai` on Bluesky (217 posts). In total, 4,953 posts were removed, resulting in 52,481 posts (LinkedIn: 18,956; X: 17,571; Bluesky: 15,954).

Preprocessing for topic modelling. For using BERTopic (Grootendorst, 2022), a more aggressive preprocessing pipeline was applied to reduce noise and improve embedding quality. Steps applied sequentially: (1) URLs removed; (2) user mentions removed; (3) the hashtag symbol removed while retaining the associated word; (4) emojis and non-Latin scripts removed; (5) accented Latin characters transliterated to ASCII equivalents using `unidecode`; (6) text lowercased; (7) punctuation removed; and (8) redundant whitespace normalized. Posts with fewer than ten words after preprocessing were excluded (379 posts removed). The final topic-ready corpus comprised 52,102 posts (LinkedIn: 18,952; X: 17,463; Bluesky: 15,687), with a mean post length of 53.0 words (median: 57.0; range: 10–92) and a mean of 362.0 characters (max: 432). The maximum post length reflects Meltwater's delivery of content snippets rather than full post text for LinkedIn, capping retrieved content at approximately 425 characters; this is discussed further as a limitation in Section 8.

Preprocessing for sentiment analysis. We followed the official NLP guidelines (Barbieri et al., 2020) for the *cardiffnlp/twitter-roberta-base-sentiment-latest* model (Loureiro et al., 2022): only URLs (replaced with `$URL$`) and mentions (replaced with `@user`) were substituted. All other features (emojis, hashtags, punctuation, casing) were retained to preserve the sentiment signal. No posts were removed at this step.

5 Results

RQ1: How does sentiment toward AI-sustainability differ across platforms?

Sentiment analysis of 52,102 posts showed significant differences across platforms ($\chi^2(4) = 10{,}549.01$, $p < 0.001$; Table 3). LinkedIn has the highest positive posts (63.49%), X has the highest neutral posts (46.89%), and Bluesky has the higes negative posts (37.22%). In total, positive sentiment is the most common (43.98%), followed by neutral (40.16%) and negative (15.86%).

In answer to RQ1, sentiment toward AI-sustainability differs significantly across platforms, with LinkedIn reflecting an optimistic professional outlook, X a more balanced perspective, and Bluesky a notably more critical stance.

Platform	# Posts	Signal	Total	Mean	Median	Std
LinkedIn	18,952	Reactions	495,041	26.1	7.0	206.2
X	17,463	Likes	411,142	23.5	1.0	877.6
Bluesky	15,687	Likes	123,698	7.9	1.0	110.3

Table 1: Posts and likes count per platform.

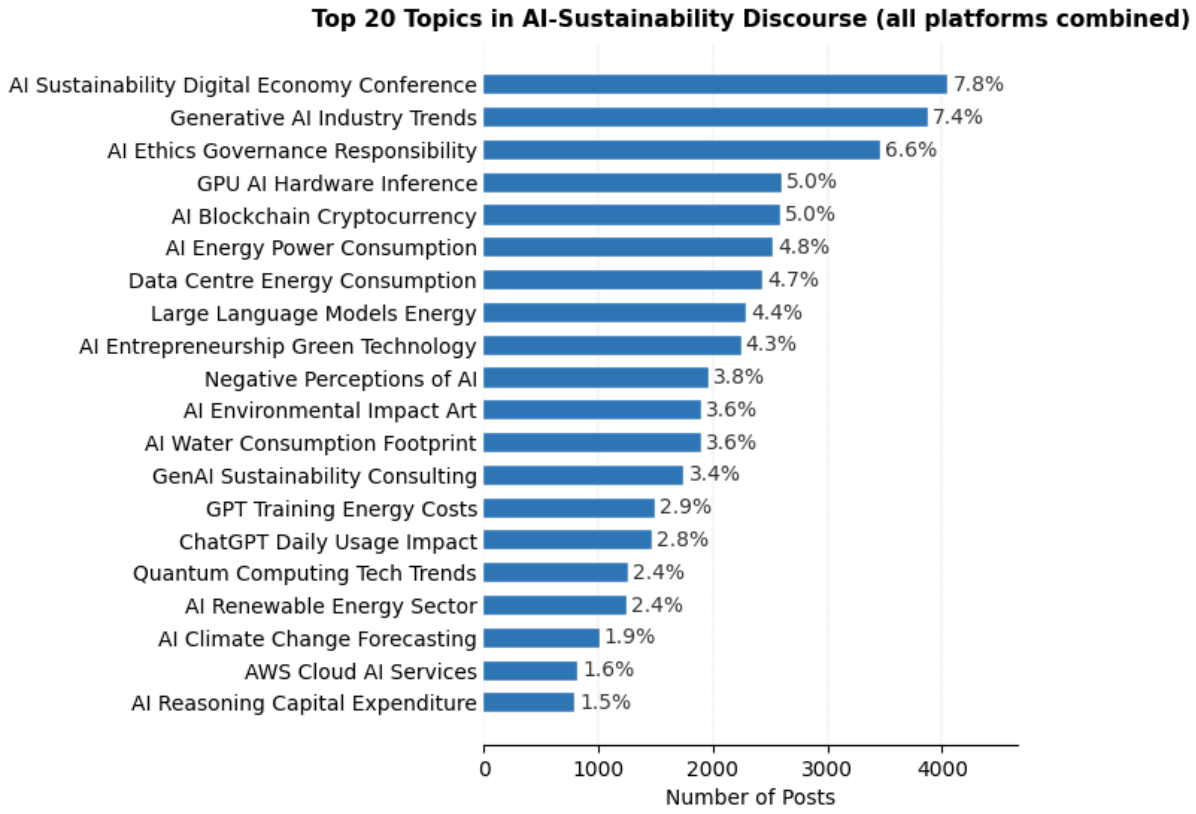


Figure 1: Top 20 topics in AI-sustainability discourse across all platforms (document count and % of dataset).

Cleaning Step	Removed	Remaining
Raw dataset	—	57,434
Duplicate removal (by URL)	568	56,866
Duplicate removal (by text content)	2,948	53,918
Non-English post removal	0	53,918
Unavailable/deleted post removal	0	53,918
Short post removal (<10 characters)	108	53,810
Recent post removal (>Oct 31, 2025)	259	53,551
Bot/spam account removal	1,070	52,481
Initial data cleaning	4,953	52,481
Preprocessing for topic modelling	379	**52,102**

Table 2: Data cleaning steps and dataset size.

RQ2: (a) What topics emerge around AI-sustainability?

Figure 1 shows the 20 largest topics by posts count. The complete topic list is in Appendix B. The 20 most common topics cover four broad thematic clusters: *corporate and governance framing* (e.g., *AI Sustainability Digital Economy Conference*, *AI Ethics Governance Responsibility*); *environmental resource costs* (e.g., *AI Energy Power Consumption*, *AI Water Consumption Footprint*); *critical and sceptical discourse* (*Negative Perceptions of AI*); and *emerging technology* (e.g., *GPU AI Hardware Inference*, *Quantum Computing Tech Trends*).

The most common topic is *AI Sustainability Digital Economy Conference* (7.8%, $n = 4{,}053$), covering corporate sustainability reporting, digital economy initiatives, and regional AI governance events. This is followed by *Generative AI Industry Trends* (7.4%, $n = 3{,}872$) and *AI Ethics Governance Responsibility* (6.6%, $n = 3{,}459$). These three topics account for 21.8% of the corpus, reflecting the strong presence of professional, corporate, and governance framings in AI-sustainability discourse.

Environmental resource topics are well-

Platform	n	Negative (%)	Neutral (%)	Positive (%)
LinkedIn	18,952	1.77	34.74	63.49
X	17,463	11.96	46.89	41.15
Bluesky	15,687	37.22	39.22	23.56
Total	**52,102**	**8,262 (15.86%)**	**20,926 (40.16%)**	**22,914 (43.98%)**

Table 3: Sentiment Distribution by Platform.

represented: *AI Energy Power Consumption* (4.8%, $n = 2{,}529$), *Data Centre Energy Consumption* (4.7%, $n = 2{,}440$), and *AI Water Consumption Footprint* (3.6%, $n = 1{,}904$) together account for 13.1% of all posts, confirming that infrastructure resource concerns are central to public AI-sustainability discourse. *Negative Perceptions of AI* (3.8%, $n = 1{,}960$) and *AI Environmental Impact Art* (3.6%, $n = 1{,}904$) further reflect the critical strand of the conversation. Topics such as *GPU AI Hardware Inference* (5.0%) and *Large Language Models Energy* (4.4%) reflect the technical side of the discourse. *AI Entrepreneurship Green Technology* (4.3%) sits at the intersection of business and sustainability concerns. Niche but emerging topics such as *Quantum Computing Tech Trends* (2.4%), *AI Renewable Energy Sector* (2.4%), and *AI Climate Change Forecasting* (1.9%) suggest that the discourse is expanding beyond generative AI into related technology areas.

In answer to RQ2(a), AI sustainability discourse centres on four themes: corporate and governance framing, environmental resource costs, critical discourse, and emerging technology, with professional and governance topics being the most common.

RQ2: (b) Do these topics differ across social media platforms?

Topic distributions differ significantly across platforms ($\chi^2(70) = 27{,}699.48$, $p < 0.001$), with each platform showing a distinct thematic profile, as shown in Table 4.

LinkedIn is driven by professional and corporate themes: *AI Sustainability Digital Economy Conference* (14.3%) and *Generative AI Industry Trends* (13.6%) are the two most prevalent topics, collectively accounting for 27.9% of LinkedIn posts, consistent with its professional networking context.

X is driven by technology and finance themes: *AI Blockchain Cryptocurrency* (13.3%) and *GPU AI Hardware Inference* (9.2%) are the two most prevalent topics, collectively accounting for 22.5% of X posts, consistent with its role as a space for technology and market discourse.

Bluesky is driven by environmental and critical themes: *AI Energy Power Consumption* (13.4%), *Negative Perceptions of AI* (10.9%), *AI Water Consumption Footprint* (10.8%), and *AI Ethics Governance Responsibility* (10.8%) are the four most prevalent topics, collectively accounting for 45.9% of Bluesky posts, consistent with its user base's focus on the environmental costs and governance concerns around AI.

In answer to RQ2(b), AI sustainability topics differ substantially across platforms, shaped by each platform's distinct user culture and context.

RQ2: (c) What is the sentiment associated with each topic?

Figure 2 presents the RoBERTa sentiment distribution across the top 20 topics. Sentiment varies substantially by topic, confirming that platform-level patterns reported in RQ1 are driven by topical composition rather than platform culture alone. Topics centred on environmental resource costs and AI criticism carry the highest negative proportions: *AI Environmental Impact Art* and *Negative Perceptions of AI* show the most negative distributions, followed by *AI Water Consumption Footprint*, *AI Energy Power Consumption*, *ChatGPT Daily Usage Impact*, *GPT Training Energy Costs*, and *Data Centre Energy Consumption*, all of which exceed 10% negative posts. By contrast, *AWS Cloud AI Services*, *Quantum Computing Tech Trends*, and *AI Renewable Energy Sector* are the most positively framed topics, each exceeding 60% positive, consistent with their association with technological opportunity and innovation. *Generative AI Industry Trends* and *AI Sustainability Digital Economy Conference* are also mostly positive, reflecting the corporate and professional framing dominant on LinkedIn. *AI Ethics Governance Responsibility* shows a more balanced distribution, reflecting the

Bluesky (n=15,687)		**LinkedIn** (n=18,952)		**X** (n=17,463)	
AI Energy Power Consumption	13.4	AI Sustainability Digital Econ. Conf.	14.3	AI Blockchain Cryptocurrency	13.3
Negative Perceptions of AI	10.9	Generative AI Industry Trends	13.6	GPU AI Hardware Inference	9.2
AI Water Consumption Footprint	10.8	AI Ethics Governance Responsibility	6.5	Generative AI Industry Trends	6.5
AI Ethics Governance Responsibility	10.8	AI Entrepreneurship Green Tech.	6.5	Large Language Models Energy	5.8
Data Centre Energy Consumption	7.9	GenAI Sustainability Consulting	4.6	AI Sustainability Digital Econ. Conf.	5.2

Table 4: Top 5 topics per platform by percentage of platform posts.

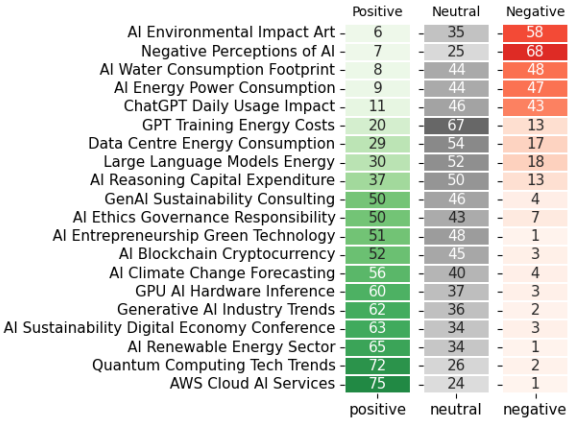


Figure 2: RoBERTa sentiment (%) distribution for the top 20 topics

contested nature of AI governance debates.

In answer to RQ2(c), sentiment is strongly tied to topic: topics framing AI as a technological opportunity are predominantly positive, and topics foregrounding environmental and resource costs are predominantly negative.

RQ3: Which topics generate the highest engagement, measured by likes, in AI sustainability discussions across platforms?

A Kruskal-Wallis test confirms that engagement differs significantly across topics on all three platforms: Bluesky ($H = 666.12$, $p < 0.001$), LinkedIn ($H = 785.60$, $p < 0.001$), and X ($H = 536.52$, $p < 0.001$). Table 5 shows the top 5 topics by median likes per platform.

On **Bluesky**, the three most engaging topics are *GenAI Carbon Footprint Concerns*, *GenAI Environmental Artist Impact*, and *AI Environmental Impact Art* (all median = 2.0), reflecting its user base's focus on environmental criticism.

On **X**, *AI Blockchain Cryptocurrency* generates the highest median likes (8.0), followed by *Quantum Computing Tech Trends* (6.0) and *AWS Cloud AI Services, I Reasoning Capital Expenditure, AI Healthcare Ecology Public Health* (all median = 4.0), reflecting its technology and finance audience.

On **LinkedIn**, *AI Circular Economy Packaging* is the most engaging topic (median = 18.5), followed by *AI Precision Agriculture Food* and *AI Sustainability Digital Economy Conference* (both median = 15.0). Notably, *AI Sustainability Digital Economy Conference*, the most common topic on LinkedIn (22.9%), ranks only third in engagement, confirming that topic popularity does not predict engagement (Aldous et al., 2019).

In answer to RQ3, the topics that generate the highest engagement are not the most common ones on any platform, suggesting that niche topics attract more interaction than broadly discussed ones. Engagement patterns are also platform-specific, with LinkedIn showing the highest median engagement overall and Bluesky the lowest, reflecting differences in audience size and platform culture.

6 Discussion

Our analysis of 52,102 posts across LinkedIn, X, and Bluesky resulted in three main findings: sentiment in AI-sustainability discourse is driven by topic rather than platform; each platform hosts a structurally distinct thematic profile; and niche topics consistently generate more engagement than high-volume ones.

Topic-Dependent Sentiment. Sentiment in AI-sustainability discourse is structured by topic rather than by platform. This extends Longo and Longo (2025), who showed on Reddit that community context shapes emotional tone: our cross-platform comparison demonstrates that even across communities with very different norms, topic remains the controlling driver. It also builds on Wang et al. (2025), who documented a shift toward negative sentiment in X posts about data center energy: our data show that this negativity extends beyond X and is tied to the topic itself rather than to any single platform. The overall positive aggregate (43.98%) should therefore not be taken as representative of the full discourse, as it reflects the high volume of positively-framed corporate and governance topics rather than a uniform positive outlook across all AI-sustainability themes.

Bluesky		X		LinkedIn	
GenAI Carbon Footprint Concerns	2.0	AI Blockchain Cryptocurrency	8.0	AI Circular Economy Packaging	18.5
GenAI Environmental Artist Impact	2.0	Quantum Computing Tech Trends	6.0	AI Precision Agriculture Food	15.0
AI Environmental Impact Art	2.0	AWS Cloud AI Services	4.0	AI Sustainability Digital Econ. Conf.	15.0
Quantum Computing Tech Trends	1.0	AI Reasoning Capital Expenditure	4.0	AI Healthcare Ecology Public Health	14.0
AI Precision Agriculture Food	1.0	AI Healthcare Ecology Public Health	4.0	AWS Cloud AI Services	12.0

Table 5: Top 5 topics by median likes per post for each platform. Med. = median likes.

Platform-Specific Topic Distributions. The significant difference in topic distributions across platforms ($\chi^2(70) = 27{,}699.48$, $p < .001$) directly explains the platform-level sentiment differences reported in RQ1. LinkedIn appears optimistic (63.49% positive) not because its users are more favorable toward AI sustainability, but because it hosts topics with mostly positive framing. Bluesky appears critical (23.56% positive) because it hosts topics with mostly negative framing. This builds on Kloo et al. (2024), who showed that similar topics are expressed differently across platforms: our results go further, showing that platforms host different conversations altogether rather than different framings of the same ones.

The Bluesky composition is partially explained by documented user migration. Chang et al. (2023) showed that nearly 50% of environmentally-oriented X users became inactive following the platform's ownership change, and Kleppmann et al. (2024) and Balduf et al. (2024) document that many migrated to Bluesky. A direct consequence is that business leaders monitoring AI-sustainability discourse through LinkedIn and X are exposed to a more optimistic information environment than the full cross-platform conversation reflects, with potential consequences for risk assessment.

Niche Topics Drive Engagement. Across all three platforms, niche topics consistently generate higher per-post engagement than high-volume ones, a pattern that is statistically robust (Kruskal-Wallis: Bluesky $H = 666.12$; LinkedIn $H = 785.60$; X $H = 536.52$; all $p < .001$) and consistent with Aldous et al. (2019, 2023), who found that topic popularity does not predict engagement, and with Bruns et al. (2016), who showed that specialized content outperforms generic content. Two mechanisms likely explain this. First, saturation: high-volume topics are discussed so frequently that individual posts offer diminishing novelty (Astvansh, 2024). Second, specificity: niche topics deliver targeted value to focused communities, producing stronger per-post interaction.

Practical Implications. Three implications follow directly from the results. First, topic framing is the primary lever for shaping audience sentiment, not platform selection: framing AI energy use around industry trends resulted in 62% positive posting across platforms, while framing it around resource costs resulted in 47–48% negative posting regardless of platform. Second, platform audiences differ in their baseline information environment, LinkedIn and X users encounter a more optimistic discourse than Bluesky users, which practitioners and policymakers should account for when selecting platforms for outreach or sampling public opinion. Third, sector-specific content outperforms generic AI-sustainability content in per-post engagement, suggesting targeted strategies: circular economy and agriculture framing on LinkedIn, technology and finance angles on X, and environmental impact framing on Bluesky.

7 Conclusion

We analyzed 52,102 posts across LinkedIn, X, and Bluesky over 15 months. Three findings stand out. Sentiment is topic-dependent, not platform-dependent: corporate and AI-potential topics are positive everywhere; infrastructure-cost topics are negative everywhere. Topic distributions differ significantly across platforms: LinkedIn favours corporate sustainability, X hardware and cryptocurrency, and Bluesky environmental resource concerns. Niche topics consistently outperform high-volume topics in engagement, supporting platform-specific over uniform content strategies. These findings ground communication strategies for AI-sustainability discourse.

8 Limitations

Platform and Language Scope. The corpus covers only English-language content from three platforms (X, Bluesky, and LinkedIn) over 15 months, limiting generalizability to other platforms, non-Western contexts, and non-English AI-sustainability discourse. Bluesky's user com-

position skews toward technology-critical and environmentally-conscious communities, which may overrepresent environmental concern relative to the broader public. Expanding to platforms such as Facebook, TikTok, or Reddit could result different topical and sentiment patterns.

Data Collection. Posts were collected using Meltwater and supplemented by the Bluesky AT Protocol API, and the two routes rely on different crawling infrastructures with their own indexing and rate limits, meaning coverage is not guaranteed to be equivalent and may introduce platform-specific collection biases. The AT Protocol collection may over-represent posts from users with high post volumes, and any posts deleted between publication and retrieval would be absent from the corpus. We used only publicly available data; private posts and restricted accounts are not represented, which may skew the observed distributions. Furthermore, the keyword list may not have captured all relevant discourse, as posts discussing AI sustainability using different terminology or framing may have been missed, introducing a selection bias toward more explicitly labelled content. Finally, Meltwater delivers content snippets rather than full post text, capping retrieved content at 91 words after preprocessing. This may systematically exclude substantive content from longer LinkedIn posts, where native post length can reach approximately 3,000 characters, potentially affecting both topic assignments and sentiment classifications for that platform.

Engagement Metrics. Engagement is measured using like counts only, as comments, shares, and views were not consistently available across all three platforms. This limits the scope of engagement analysis and may not capture the full range of audience interaction. Furthermore, engagement varies substantially by account reach, but follower counts were not available for all posts, and not all posts were associated with identifiable account names, precluding follower-normalised engagement analysis.

Methodological Choices. The 36-topic granularity reflects a deliberate hyperparameter choice ($\text{min_cluster_size} = 260$) rather than a ground truth partition of the discourse; different parameter values would yield different topic structures (Grootendorst, 2022). Engagement is used as a proxy for content resonance, though high engagement may reflect controversy or algorithmic amplification rather than genuine audience interest. Automated account detection beyond the 1% volume threshold applied during preprocessing was not performed, which may allow residual automated content to influence engagement and sentiment distributions.

Sentiment Classifier. The cardiffnlp/twitter-roberta-base-sentiment-latest model was trained on Twitter data and is applied here to LinkedIn and Bluesky without domain adaptation or recalibration (Barbieri et al., 2020; Loureiro et al., 2022). LinkedIn's formal professional register differs substantially from Twitter conventions, and the observed LinkedIn negative rate (1.77%) may partly reflect underprediction of negative sentiment in formal text rather than platform sentiment alone. Platform-stratified human annotation was not feasible within the scope of this study; future work should validate classifier performance per platform and assess whether the reported topic-level sentiment patterns are robust to recalibration.

9 Ethical Considerations

This study was conducted under institutional ethics review board approval. We analyzed publicly available social media posts related to AI-sustainability discourse. Data were collected via Meltwater in accordance with platform terms of service and institutional licensing agreements, without web scraping or circumvention of platform safeguards. We distinguish between platform-level consent, granted through users' acceptance of platform terms, and informed consent in the research ethics sense. Because this study involves large-scale secondary analysis of publicly accessible content, obtaining individual consent was not feasible. To minimize potential harms, analyses are conducted exclusively at the aggregate level. We do not identify, profile, or quote individual users in ways that would enable re-identification, and findings are reported at the level of platforms and topics rather than individuals.

Data are stored securely with restricted access and are used solely for analysis and verification. Publicly shared materials do not include raw or user-level data. Our dataset contains publicly posted content, some of which includes self-disclosed personal information (e.g., locations). Although posts were public at the time of collection, we follow these guidelines: (1) presenting only aggregate statistics, (2) not sharing raw data publicly to prevent re-identification, and (3) storing data securely with limited access. Users may not have

anticipated research use; we minimize privacy risks through aggregation and anonymization in reporting.

Beyond privacy, we considered potential downstream risks of this work. Results could be misused to optimize greenwashing strategies or amplify misinformation. We mitigate these risks by contextualizing findings, avoiding prescriptive engagement optimization guidance, and emphasizing that engagement metrics do not imply factual accuracy.

Although individual participants did not directly consent to this research, the study serves a broader public interest by documenting how environmental concerns about AI are expressed across different online communities. The findings contribute to informed debate on AI sustainability and responsible technology governance.

10 Data Availability

The dataset used in this study was collected via the Meltwater platform from LinkedIn, X, and Bluesky using the keyword query described in Section 3. The dataset can be requested through a controlled-access process at https://forms.gle/YUFdA16R6HkSZjp88

11 Acknowledgements

This work was supported by NPRP Grant No. 14C-0916-210015 from the Qatar National Research Fund (QNRF), a member of the Qatar Research, Development and Innovation (QRDI) Council. The findings and conclusions presented herein are solely those of the authors.

During the preparation of this work, the author(s) used large language models (including OpenAI's ChatGPT and Anthropic's Claude Sonnet 4.6) to assist with data analysis, visualisation code development, and writing. The author(s) subsequently reviewed, verified, and edited all content and take full responsibility for the final version of the manuscript.

References

Kholoud Khalil Aldous, Jisun An, and Bernard J. Jansen. 2019. View, like, comment, post: Analyzing user engagement by topic at 4 levels across 5 social media platforms for 53 news organizations. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 47–57.

Kholoud Khalil Aldous, Jisun An, and Bernard J. Jansen. 2023. What really matters?: characterising and predicting user engagement of news postings using multiple platforms, sentiments and topics. *Behaviour & Information Technology*, 42(5):545–568.

Ezekiel Appiah and Amal Htait. 2025. Ai in the public eye: Analysing social media sentiment and opinion on artificial intelligence. In *Proceedings of the UK AI Conference 2024*, volume 295 of *Proceedings of Machine Learning Research*, pages 1–15. PMLR.

Wenceslao Arroyo-Machado, Nicolas Robinson-Garcia, and Daniel Torres-Salinas. 2025. Are there stars in bluesky? a comparative exploratory analysis of altmetric mentions between x and bluesky. *Journal of Informetrics*, 19(3):101700.

Vivek Astvansh. 2024. Recycling firm-generated content on social media platforms: phenomenon and research propositions. *Marketing Intelligence & Planning*, 42(8):1407–1432.

Leonhard Balduf, Saidu Sokoto, Onur Ascigil, Gareth Tyson, Björn Scheuermann, Maciej Korczyński, Ignacio Castro, and Michał Król. 2024. Looking AT the Blue Skies of Bluesky. In *Proceedings of the 2024 ACM on Internet Measurement Conference*, pages 76–91, Madrid Spain. ACM.

Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the association for computational linguistics: EMNLP 2020*, pages 1644–1650.

Jonah Berger and Katherine L. Milkman. 2012. What makes online content viral? *Journal of Marketing Research*, 49(2):192–205.

Jan C. Bieser, Ralph Hintemann, Lorenz M. Hilty, and Severin Beucker. 2023. A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology. *Environmental Impact Assessment Review*, 99:107033.

Md. Rafiul Biswas, Mabrouka Bessghaier, Shimaa Ibrahim, George Mikros, and Wajdi Zaghouani. 2026. Longitudinal trends in global climate change discourse on facebook. In *Proceedings of the ACM Web Conference 2026*, WWW '26, page 9473–9481, New York, NY, USA. Association for Computing Machinery.

Bluesky Social. 2024. AT Protocol: Authenticated Transfer Protocol. https://atproto.com. Accessed: April 21, 2026.

Danah M. Boyd and Nicole B. Ellison. 2007. Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication*, 13(1):210–230.

Axel Bruns, Brenda Moon, Avijit Paul, and Felix Münch. 2016. Towards a typology of hashtag publics: a large-scale comparative study of user engagement across trending topics. *Communication Research and Practice*, 2(1):20–46.

Stephen Cave, Kate Coughlan, and Kanta Dihal. 2019. "scary robots": Examining public responses to ai. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, page 331–337. Association for Computing Machinery.

Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and Krishna Gummadi. 2010. Measuring user influence in twitter: The million follower fallacy. In *Proceedings of the international AAAI conference on web and social media*, 1, pages 10–17.

Charlotte H. Chang, Nikhil R. Deshmukh, Paul R. Armsworth, and Yuta J. Masuda. 2023. Environmental users abandoned Twitter after Musk takeover. *Trends in Ecology & Evolution*, 38(10):893–895.

Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):e2023301118. Publisher: Proceedings of the National Academy of Sciences.

Matteo Cinelli, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoli, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The COVID-19 social media infodemic. *Sci Rep*, 10(1):16598.

Ethan Fast and Eric Horvitz. 2017. Long-term trends in the public perception of artificial intelligence. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI'17, page 963–969. AAAI Press.

Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint*. ArXiv:2203.05794 [cs].

Jinsong Guo, Aditya Jami, Markus Kröll, Lukas Schweizer, Sergey Paramonov, Eric Aichinger, Stefano Sferrazza, Mattia Scaccia, Stéphane Reissfelder, Eda Cicek, Giovanni Grasso, and Georg Gottlob. 2023. When Automatic Filtering Comes to the Rescue: Pre-Computing Company Competitor Pairs in Owler. *Proceedings of the ACM on Management of Data*, 1(2):1–23.

Smriti Janaswamy and Jeremy Blackburn. 2024. Exploring climate change discourse: Measurements and analysis of reddit data. In *Workshop Proceedings of the International AAAI Conference on Web and Social Media*. NOCAPS 2024: Networks and Opinions on Climate Action in the Public Sphere.

Martin Kleppmann, Paul Frazee, Jake Gold, Jay Graber, Daniel Holmgren, Devin Ivy, Jeromy Johnson, Bryan Newbold, and Jaz Volpert. 2024. Bluesky and the AT Protocol: Usable Decentralized Social Media. In *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*, DIN '24, pages 1–7, New York, NY, USA. Association for Computing Machinery.

Ulrike Klinger, Daniel Kreiss, and Bruce Mutsvairo. 2023. *Platforms, power, and politics: An introduction to political communication in the digital age*. John Wiley & Sons.

Ian Kloo, Iain J Cruickshank, and Kathleen M Carley. 2024. A cross-platform topic analysis of the nazi narrative on twitter and telegram during the 2022 russian invasion of ukraine. In *Proceedings of the international AAAI conference on web and social media*, volume 18, pages 839–850.

Diego Kozlowski, Carolina Pradier, and Pierre Benz. 2024. Generative ai for automatic topic labelling. *arXiv preprint arXiv:2408.07003*.

David Lazer, Alex Pentland, Lada A. Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, and et al. 2009. Computational social science. *Science*, 323(5915):721–723.

Vivian Liu and Yiqiao Yin. 2024. Green ai: exploring carbon footprints, mitigation strategies, and trade offs in large language model training. *Discover Artificial Intelligence*, 4(1):49.

Marcelo Sartori Locatelli, Pedro Dutenhefner, Arthur Buzelin, Pedro Loures Alzamora, Yan Aquino, Pedro Augusto Torres Bento, Samira Malaquias, Victoria Estanislau, Caio Santana, Lucas Dayrell, and 1 others. 2025. Ai and climate change discourse: What opinions do large language models present? In *Proceedings of the 2nd Workshop on Natural Language Processing Meets Climate Change (ClimateNLP 2025)*, pages 113–125.

Arianna Longo and Alessandro Y. Longo. 2025. No AI on a dead planet: Sentiment and emotion analysis across Reddit communities on AI and the environment. In *Proceedings of the 1st Workshop on Ecology, Environment, and Natural Language Processing (NLP4Ecology2025)*, pages 77–83. University of Tartu Library.

Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. TimeLMs: Diachronic language models from Twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland. Association for Computational Linguistics.

Vibeke Madsen. 2022. Making sense of internal social media and participatory communication: exploring the employee perspective in a change process in a public sector organization. *Journal of Communication Management*, 26.

Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Jack Clark, and Michael C. Horowitz. 2025. The 2025 ai index report. Technical report, Stanford Institute for Human-Centered Artificial Intelligence.

Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, and 1 others. 2024. Artificial intelligence index report 2024. *arXiv preprint arXiv:2405.19522*.

Leland McInnes, John Healy, Steve Astels, and 1 others. 2017. hdbscan: Hierarchical density based clustering. *J. Open Source Softw.*, 2(11):205.

Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.

K. Mentzer and 1 others. 2024. Analyzing reddit discourse surrounding generative ai. *Issues in Information Systems*, pages 277–292.

Kunihiro Miyazaki, Taichi Murayama, Takayuki Uchiba, Jisun An, and Haewoon Kwak. 2024. Public perception of generative ai on twitter: an empirical study based on occupation and usage. *EPJ Data Science*, 13(1):2.

David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. Carbon emissions and large neural network training. *arXiv*.

Petr Polak and Muhammad Anshari. 2024. Exploring the multifaceted impacts of artificial intelligence on public organizations, business, and society. *Humanities and Social Sciences Communications*, 11(1):1–3.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Matthias C Rillig, Marlene Ågerstrand, Mohan Bi, Kenneth A Gould, and Uli Sauerland. 2023. Risks and benefits of large language models for the environment. *Environmental science & technology*, 57(9):3464–3466.

Stefan Stieglitz and Linh Dang-Xuan. 2013. Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. *Journal of Management Information Systems*, 29(4):217–248.

Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in nlp. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 3645–3650.

Jeffrey W. Treem and Paul M. Leonardi. 2013. Social media use in organizations: Exploring the affordances of visibility, editability, persistence, and association. *Annals of the International Communication Association*, 36(1):143–189.

Oren Tsur and Ari Rappoport. 2012. What's in a hashtag? content based prediction of the spread of ideas in microblogging communities. In *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*, WSDM '12, page 643–652. Association for Computing Machinery.

Daniel Usera and Natalie Durham. 2025. What predicts engagement on linkedin? engagement-boosting strategies for professionals. *Business and Professional Communication Quarterly*, page 23294906251336710.

Onur Varol, Emilio Ferrara, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2017. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 280–289.

Hongyu Wang, Weiqi Hua, Jinqing Peng, and Maomao Hu. 2025. Public sentiment analysis of data center energy consumption using social media data and large language models. *Energy and Buildings*, 341:115802.

Lu Xian, Lingyao Li, Yiwei Xu, Ben Zefeng Zhang, and Libby Hemphill. 2024. Landscape of large language models in global english news: Topics, sentiments, and spatiotemporal analysis. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 1661–1673.

Zhaoxiang Xu, Qingguo Fang, Yanbo Huang, and Mingjian Xie. 2024. The public attitude towards chatgpt on reddit: A study based on unsupervised learning from sentiment analysis and topic modeling. *PLOS ONE*, 19(5):1–26.

Wajdi Zaghouani, Md. Rafiul Biswas, Mabrouka Bessghaier, Shimaa Amer Ibrahim, and George Mikros. 2026. ClimateChat-300K: A multi-modal Facebook dataset for understanding diverse perspectives in climate communication. In *Proceedings of the Fifteenth Language Resources and Evaluation Conference*, pages 11221–11229, Palma de Mallorca, Spain. ELRA Language Resource Association.

Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Nicolas Kourtelris, Ilias Leontiadis, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2017. The web centipede: understanding how web communities influence each other through the lens of mainstream and alternative news sources. In *Proceedings of the 2017 Internet Measurement Conference*, IMC '17, page 405–417. Association for Computing Machinery.

Er-Te Zheng, Xiaorui Jiang, Zhichao Fang, and Mike Thelwall. 2025. How is science discussed on bluesky? *Preprint*, arXiv:2507.18840.

A BERTopic Implementation Details

Documents were embedded using the `all-mpnet-base-v2` sentence transformer (Reimers and Gurevych, 2019), which produces 768-dimensional embeddings. Embeddings were reduced to five dimensions using UMAP (McInnes et al., 2018) with n_neighbors $= 15$, min_dist $= 0.0$, cosine similarity metric, and random_state $= 42$. Clustering was performed with HDBSCAN (McInnes et al., 2017) with min_cluster_size $= 260$, min_samples $= 10$, Euclidean metric, and EOM cluster selection. Topic representations were refined using a `CountVectorizer` with English stop words, min_df $= 5$, max_df $= 0.85$, and ngram_range $= (1, 2)$. Additional representations were generated via KeyBERTInspired, Maximal Marginal Relevance (MMR, diversity$= 0.3$), and Part-of-Speech tagging using spaCy.

B Complete Topic List

Table 6 presents the full set of 36 topics identified by BERTopic, ordered by document count. Topic labels were generated using Claude Sonnet 4.6 by prompting the model with each topic's three representative documents and top 20 c-TF-IDF keywords. All labels were reviewed and verified by the authors for accuracy and domain relevance. The top five c-TF-IDF keywords per topic are provided for reproducibility. Platform columns report the percentage of each platform's posts assigned to that topic, enabling cross-platform comparison of topical prevalence.

Table 6: Complete list of 36 topics identified by BERTopic. Topics are sorted by document count.

ID	Claude Topic Label	N	%	Top Keywords	LinkedIn %	Bluesky %	X %
0	AI Sustainability Digital Economy Conference	4,053	7.8	sarawak, idecs25, genai sustainability, digitaleconomy, sustainability greentech	14.3	2.7	5.2
7	Generative AI Industry Trends	3,872	7.4	shaping, ai generative, skills, content creation, ai trends	13.6	1.0	6.5
8	AI Ethics Governance Responsibility	3,459	6.6	governance, ethical ai, responsible ai, suno, song suno	6.5	10.8	3.0
4	GPU AI Hardware Inference	2,601	5.0	gpu, inference, gpus, chip, memory	3.7	1.8	9.2
1	AI Blockchain Cryptocurrency	2,590	5.0	crypto, decentralized, web3, btc, earn	0.7	0.9	13.3
10	AI Energy Power Consumption	2,529	4.8	ai energy, power consumption, electricity, nuclear, fossil	0.8	13.4	1.6
5	Data Centre Energy Consumption	2,440	4.7	centres, data centres, electricity, datacenters, center energy	3.1	7.9	3.5
2	Large Language Models Energy	2,288	4.4	models llms, slms, small language, smaller, energy cost	3.6	3.8	5.8
24	AI Entrepreneurship Green Technology	2,245	4.3	entrepreneurship, entrepreneurship businessstrategy, datadriven greentech, iot bigdata, futureready datadriven	6.5	1.2	4.7
11	Negative Perceptions of AI	1,960	3.8	bad, theft, negative, water usage, hate	0.5	10.9	0.8
3	AI Water Consumption Footprint	1,904	3.6	water usage, water use, water consumption, ai water, freshwater	0.5	10.8	0.6
9	AI Environmental Impact Art	1,904	3.6	ai environmental, artists, ai art, explained generative, explained	0.8	5.5	5.1
22	GenAI Sustainability Consulting	1,750	3.4	genai sustainability, generativeai genai, capgemini, sustainability generative, energy demands	4.6	1.0	4.1
17	GPT Training Energy Costs	1,494	2.9	training large, ai energy, like gpt, training gpt, consumes	1.7	2.8	4.2
6	ChatGPT Daily Usage Impact	1,464	2.8	query, use chatgpt, coworker, told, google search	0.7	6.0	2.2
14	Quantum Computing Tech Trends	1,261	2.4	quantum, quantum computing, tech trends, quantumcomputing, shaping	3.2	0.9	2.9
15	AI Renewable Energy Sector	1,248	2.4	energy sector, grid, solar, ai energy, nuclear	3.1	2.3	1.6
32	AI Climate Change Forecasting	1,010	1.9	climate change, ai climate, weather, geospatial, forecasting	2.0	2.8	1.1
35	AWS Cloud AI Services	818	1.6	aws, palantir, amazon web, web services, pltr	2.7	0.3	1.5
33	AI Reasoning Capital Expenditure	798	1.5	capex, reasoning, asi really, reasoning model, dyson spheres	0.8	0.4	3.3
16	GenAI Environmental Artist Impact	797	1.5	impact genai, use genai, using genai, genai environmental, bad	0.6	2.6	1.6
13	AI Healthcare Ecology Public Health	793	1.5	ecology, publichealth, ai healthcare, medicine, cancer	1.8	1.8	1.0
21	AI Retail Consumer Personalisation	761	1.5	retail, brands, customer, brand, experiences	3.0	0.5	0.7
26	AI Manufacturing Robotics	751	1.4	maintenance, manufacturers, robotics, robots, siemens	2.8	0.4	0.9
19	AI Inference Efficiency Benchmarking	704	1.4	text prompt, median, efficiency important, gemini apps, ai efficiency	0.4	0.8	2.9
12	AI Precision Agriculture Food	683	1.3	agriculture, food, farming, farmers, agricultural	2.1	1.1	0.6
18	DeepSeek Chinese AI Models	680	1.3	chinese, r1, deepseek r1, deepseek ai, world brain	1.1	0.9	1.9
34	Mobile AI Chip Efficiency	671	1.3	apple, chip, flash, mobile, gemini ai	1.4	0.4	1.9

Continued on next page

Table 6 continued

ID	Claude Topic Label	N	%	Top Keywords	LinkedIn %	Bluesky %	X %
20	AI Supply Chain Logistics	662	1.3	supply chain, procurement, logistics, supply chains, chains	2.8	0.4	0.4
30	AI Circular Economy Packaging	649	1.2	packaging, recycling, bert, circular, chemistry	2.3	0.7	0.6
28	AI Finance Fintech Banking	638	1.2	finance, fintech, banking, insurance, financial services	2.2	0.3	0.9
27	AI Smart Cities Urban Mobility	630	1.2	mobility, cities, urban, transportation, city	1.9	0.7	1.0
23	AI Construction Real Estate	600	1.1	construction, estate, real estate, buildings, architecture	2.1	0.6	0.7
25	AI Fashion Sustainability	527	1.0	fashion, cutt, cutt ly, ly, luxury	1.5	1.2	0.4
31	GenAI Carbon Footprint Concerns	451	0.9	scary, aside, footprint generative, attracts, guy	0.0	0.2	2.3
29	1-Bit LLM CPU Inference	417	0.8	bitnet, cpp, sourced, blazing, blazing fast	0.5	0.1	1.8