From Rumors to Risk: Mapping and Modeling Climate-Disaster Misinformation

Tristan Ballard
Sust Global
San Francisco, California

Abstract

Recent years have seen a surge in climate-disaster misinformation, with social media amplifying unfounded claims in the lead-up to and aftermath of major disasters. This misinformation has hindered disaster preparation and recovery while fueling harassment against meteorologists and government officials, eroding trust in scientific institutions. While tools exist for analyzing general climate-change misinformation, current datasets often overlook the rapidly shifting narratives tied to specific events like wildfires, floods, or hurricanes. This proposal addresses that gap by developing a dynamic, evolving dataset on climate-disaster misinformation. Built through targeted social media data collection and rigorous labeling, the dataset will adapt alongside AI/ML advancements through iterative feedback from model performance and emerging trends. This openly accessible resource will enable researchers and practitioners to refine detection algorithms, design interventions, and inform crisis communication strategies—ensuring both data and models remain aligned with the shifting misinformation landscape. Ultimately, this work seeks to clarify key drivers of misinformation propagation and support more effective climate disaster response.

1 Introduction

Misinformation surrounding climate and weather is far from new (1). Throughout history, there have been claims of governments manipulating rain patterns, and climate change misinformation has been prevalent for decades(2). However, recent years have seen a dramatic escalation in both the volume and impact of climate-disaster misinformation. Social media has enabled false narratives to spread further and faster than ever before, fueling distrust in scientific institutions and even leading to harassment and threats against meteorologists (3)(4)(5).

During Hurricanes Helene and Milton in 2024, for example, rumors spread that the National Oceanic and Atmospheric Administration (NOAA) had artificially strengthened the storms to influence election outcomes. In the aftermath, misinformation alleged mismanagement of relief efforts (3)(4). An analysis from the Institute for Strategic Dialogue found that in the wake of Helene, just 33 debunked tweets generated over 160 million views, and misinformation led to credible threats of violence against FEMA officials (3). In response, FEMA established dedicated Rumor Response Pages (6), underscoring the growing impact of such theories on disaster preparedness, public safety, and trust in institutions.

The dangers of widespread climate-disaster misinformation are manifold (7)(2). In the immediate wake of a disaster, distorted claims about relief measures or evacuation routes can jeopardize rescue efforts, while longer-term exposure to conspiracy theories undermines public confidence in meteorological and climate science—fostering skepticism at a time when coordinated action is more critical than ever(8). Despite the urgent need for robust detection and intervention strategies, research on climate-disaster misinformation remains limited.

Several initiatives have introduced specialized tools and datasets for studying climate misinformation, demonstrating the promise of advanced analytical approaches. For instance, the 2024 ClimateMiSt corpus (9) contains 146,670 tweets with veracity and stance annotations, and CLIMATE-FEVER (10) is designed to verify real-world climate claims. Meanwhile, another project revealed climate change misinformation can spread widely on LinkedIn (11).

Nevertheless, most of these projects and research discoveries concentrate on overarching climate change skepticism and denialism rather than the unique and varied misinformation that surfaces during and after specific weather disasters. While broader climate misinformation often disputes scientific consensus, event-specific rumors can focus on crisis details, resource allocation, and official instructions. Further, the publicly available datasets tend to remain static once published, limiting their ability to capture newly emergent rumors and adapt to shifting social media patterns. A dynamic, evolving resource is needed to capture these changes and align with advances in AI modeling (12).

This proposal addresses that gap by developing a dedicated, dynamic dataset on climate-disaster misinformation and applying AI/ML techniques to analyze it. By tracking narratives around specific disasters—such as wildfires and hurricanes—and mapping their spread before and after crises, we aim to generate actionable insights for countering false claims, restoring public trust, and improving disaster response and recovery.

2 Generating a Public Dataset on Climate-Disaster Misinformation

2.1 Dataset Scope and Objectives

The goal is to develop a dynamic, regularly updated dataset (semiannually to annually) that captures climate-disaster misinformation. It will focus on the lead-up to and aftermath of specific events rather than broader climate change misinformation. Initial efforts will target recent events—such as the 2025 Los Angeles wildfires or Hurricanes Helene and Milton—using rigorous labeling protocols. Over time, historical and global incidents will be integrated to track misinformation trends across regions and their evolution (13). To keep the dataset aligned with AI/ML advancements, iterative feedback from model performance and emerging social media trends will continuously refine both the dataset and analytical approaches (13).

2.2 Data Collection Strategy

Social media content from platforms like Twitter and Reddit will be compiled using relevant keywords, hashtags, and location-based filters. News articles, press releases, and fact-checking resources (6) can provide context and ground truth for labeling. A modular design will enable continuous updates to cover new events, capturing shifting narratives and maintaining the dataset's relevance for ongoing research.

2.3 Labeling and Annotation

Domain experts will label each entry according to its accuracy, the type of misinformation (e.g., conspiracy or false attribution), and metadata such as source credibility. This approach builds on the foundational methodologies seen in LIAR (PolitiFact-labeled political statements)(14), and also draws from more recent climate-focused datasets like CLIMATE-FEVER (10) or ClimateMiSt (9). By tailoring the labeling schema to event-specific, crisis-related dimensions, this dataset will more accurately reflect the unique characteristics of misinformation that emerge during and after climate disasters. Advances in natural language processing may enable these experts to scale the labeling process efficiently (15).

2.4 Public Accessibility and Research Impact

Central to this effort is making the dataset freely available through established data-sharing platforms, allowing researchers worldwide to replicate findings, develop innovative detection models, and contribute new annotations. This open, collaborative model is expected to drive progress in real-time monitoring tools, early warning systems, and evidence-based policy interventions for misinformation in disaster scenarios.

3 Modeling and Insights

Applying advanced AI/ML methods to the new dataset can reveal how false narratives form and spread. Previous successes with resources like LIAR (14) show that combining textual classification and network analysis provides comprehensive insights into misinformation dynamics, including user characteristics that drive propagation. Making this dataset publicly accessible will empower external researchers and industry experts to develop and refine such models. This open collaboration model aims to maximize innovation in detecting, understanding, and mitigating the spread of climate-disaster misinformation.

3.1 Textual Analysis and Classification

A key first step is to train supervised or semi-supervised classifiers to identify misinformation. Transformer-based models such as BERT or RoBERTa—already effective at detecting false claims in political domains (16)—could be adapted for climate-disaster content. As with LIAR's reliance on PolitiFact labels (14), fact-checking data and domain-expert annotations would guide these models, helping to detect patterns in conspiracy-oriented or pseudo-scientific language that emerge around climate-disaster events. Over time, classifiers could improve through fine-tuning and multi-class classification techniques, distinguishing between outright fabrications and subtler forms of misleading information.

3.2 Network Analysis and Influencer Identification

Graph-based techniques are essential for mapping how misinformation moves across social networks. Prior studies have shown that a small group of "supersharers" can produce a disproportionately large impact on the overall spread of false content. For example, an analysis of 2020 U.S. election interference misinformation found that just 2,107 accounts were responsible for 80% of misinformation content (17). Constructing interaction networks from retweets, replies, or shared URLs would help isolate these high-influence nodes. Community detection algorithms (e.g., Louvain (18)) and influence propagation models (e.g., Independent Cascade (19)) could simulate how misinformation might travel under different conditions. When demographic or geographic data is available, deeper insights might emerge about the users most likely to initiate or amplify climate-disaster conspiracies.

3.3 Feedback Loops for Dataset Enrichment

An iterative approach allows model findings to shape continued data gathering. If analyses expose new misinformation narratives or reveal unexpected clusters, the dataset can be updated to capture these developments. This dynamic loop—where results guide subsequent data collection—has been effective in earlier misinformation research, ensuring that both the dataset and models remain relevant as false narratives evolve (20).

3.4 Expected Insights and Applications

Combined textual and network analyses may highlight which communities are most vulnerable to particular conspiracies and identify common user traits behind high-impact misinformation (17). Recognizing such patterns lays the groundwork for targeted interventions and real-time monitoring tools. In the long run, these techniques have the potential to drive early-warning systems, empower policymakers to allocate resources more efficiently, and support more informed public communications during climate-disaster events.

4 Challenges and Limitations

Despite its promise, the proposed framework faces certain obstacles. First, data access may be restricted by platform API limits or privacy policies, limiting both historical and real-time collection. Second, ensuring data quality poses challenges, as user-generated content can be noisy and may require extensive preprocessing. Third, labeling consistency can be difficult to maintain, especially when dealing with varied disaster scenarios and misinformation narratives, necessitating standardized annotation protocols. Fourth, ongoing funding and resource allocation will be essential for sustaining the dataset's infrastructure and expert oversight. Finally, generalizability of models trained on specific events may not transfer seamlessly to new or shifting narratives, underscoring the need for strategic data partnerships and continual updates.

REFERENCES

[1] Joseph E Uscinski, Karen Douglas, and Stephan Lewandowsky. Climate change conspiracy theories. *Climate Science*, pages 1–35, 2017.

[2] Karen M Douglas and Robbie M Sutton. Climate change: Why the conspiracy theories are dangerous. *Bulletin of the Atomic Scientists*, 71(2):98–106, 2015.

[3] Institute for Strategic Dialogue. Hurricane helene brews up storm of online falsehoods and threats. Website, October 2024. Accessed: 2025-02-03 at `https://www.isdglobal.org/digital_dispatches/hurricane-helene-brews-up-storm-of-online-falsehoods-and-threats/`.

[4] Lorena O'Neil. Meteorologists get death threats as hurricane milton conspiracy theories thrive. Rolling Stone, October 2024. Accessed: 2025-02-03 at https://www.rollingstone.com/culture/culture-features/hurricane-milton-misinformation-meteorlogist-death-threats-1235130352/.

[5] Rachel Leingang. 'a flood of disinformation': Rumors and lies abound amid los angeles wildfires. The Guardian, January 2025. Accessed: 2025-02-03 at https://www.theguardian.com/us-news/2025/jan/16/disinformation-los-angeles-wildfires.

[6] Federal Emergency Management Agency. Common disaster-related rumors. Website. Accessed: 2025-02-03 at https://www.fema.gov/disaster/recover/rumor-response.

[7] Mikey Biddlestone, Flavio Azevedo, and Sander van der Linden. Climate of conspiracy: A meta-analysis of the consequences of belief in conspiracy theories about climate change. *Current Opinion in Psychology*, 46:101390, 2022.

[8] Kim-Pong Tam and Hoi-Wing Chan. Conspiracy theories and climate change: A systematic review. *Journal of Environmental Psychology*, page 102129, 2023.

[9] YeonJung Choi, Lanyu Shang, and Dong Wang. Climatemist: Climate change misinformation and stance detection dataset. In *International Conference on Advances in Social Networks Analysis and Mining*, pages 321–330. Springer, 2024.

[10] Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. Climate-fever: A dataset for verification of real-world climate claims. *arXiv preprint arXiv:2012.00614*, 2020.

[11] Texas A&M Arts & Sciences Marketing & Communications. Atmospheric scientist tackles climate change misinformation on linkedin. Website, October 2024. Accessed: 2025-02-03 at https://artsci.tamu.edu/news/2024/10/atmospheric-scientist-tackles-climate-change-misinformation-on-linkedin.html.

[12] Hunt Allcott, Matthew Gentzkow, and Chuan Yu. Trends in the diffusion of misinformation on social media. *Research & Politics*, 6(2):2053168019848554, 2019.

[13] Matthew J Hornsey, Emily A Harris, and Kelly S Fielding. Relationships among conspiratorial beliefs, conservatism and climate scepticism across nations. *Nature Climate Change*, 8(7):614–620, 2018.

[14] William Yang Wang. " liar, liar pants on fire": A new benchmark dataset for fake news detection. *arXiv preprint arXiv:1705.00648*, 2017.

[15] Shaina Raza, Drai Paulen-Patterson, and Chen Ding. Fake news detection: comparative evaluation of bert-like models and large language models with generative ai-annotated data. *Knowledge and Information Systems*, pages 1–26, 2025.

[16] Heejung Jwa, Dongsuk Oh, Kinam Park, Jang Mook Kang, and Heuiseok Lim. exbake: Automatic fake news detection model based on bidirectional encoder representations from transformers (bert). *Applied Sciences*, 9(19):4062, 2019.

[17] Sahar Baribi-Bartov, Briony Swire-Thompson, and Nir Grinberg. Supersharers of fake news on twitter. *Science*, 384(6699):979–982, 2024.

[18] Pasquale De Meo, Emilio Ferrara, Giacomo Fiumara, and Alessandro Provetti. Generalized louvain method for community detection in large networks. In *2011 11th international conference on intelligent systems design and applications*, pages 88–93. IEEE, 2011.

[19] Chi Wang, Wei Chen, and Yajun Wang. Scalable influence maximization for independent cascade model in large-scale social networks. *Data Mining and Knowledge Discovery*, 25:545–576, 2012.

[20] Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. Detection and resolution of rumours in social media: A survey. *Acm Computing Surveys (Csur)*, 51(2):1–36, 2018.

A Appendix: Proposed Timeline

The proposed timeline outlines key phases for dataset development, model training, and long-term sustainability.

Month 1–3: Project Setup and Data Acquisition

- Finalize data collection agreements and access rights (e.g., platform APIs).
- Begin gathering initial social media and fact-checking content related to select disasters.

Month 4–6: Dataset Curation and Annotation

- Clean and preprocess collected data.
- Establish labeling protocols and train annotators for consistent quality.
- Generate the first publicly accessible version of the dataset.

Month 7–9: Preliminary Modeling and Analysis

- Develop and test initial classification models (e.g., using BERT or GPT).
- Begin network analyses to identify influential users or "super-spreaders."

Month 10–12: Refinement and Iteration

- Incorporate feedback from early users and domain experts.
- Update dataset with additional events or newly surfaced misinformation narratives.
- Refine models and network analyses based on expanded data.

Month 13–18: Scaling and Collaboration

- Expand dataset coverage to include international disasters.
- Integrate findings into pilot monitoring tools or dashboards for policymakers.
- Explore partnerships for sustained dataset growth and improved labeling.

Month 19–24: Evaluation and Long-Term Sustainability

- Conduct comprehensive evaluation of model performance and data quality.
- Publish results, lessons learned, and best practices.
- Secure ongoing funding or institutional support for dataset maintenance and updates.