

MASH: A Multiplatform and Multimodal Annotated Dataset for Societal Impact of Hurricane

Ruichen Yao¹, Aslanbek Murzakhmetov^{1,2*}, Raaghav Pillai^{1*}, Aliya Maussymbayeva^{1,3*}, Zelin Li¹, Yifan Liu¹, Yaokun Liu¹, Lanyu Shang⁴, Yang Zhang⁵, Na Wei¹, Ximing Cai¹, Dong Wang¹

¹University of Illinois Urbana-Champaign, Illinois, USA

²M.Kh. Dulaty Taraz University, Jambyl Region, Kazakhstan

³Abylkas Saginov Karaganda Technical University, Karaganda Region, Kazakhstan

⁴Loyola Marymount University, California, USA

⁵Miami University, Ohio, USA

¹{ryao8, aslanbek, pillai7, aliyam, zelin3, yifan40, yaokun2, nawei2, xmcai, dwang24}@illinois.edu

²an.murzakhmetov@dulaty.kz, ³a.maussymbayeva@ktu.edu.kz, ⁴lanyu.shang@lmu.edu, ⁵zhang981@miamioh.edu

Abstract

Natural disasters cause multidimensional threats to human societies, with hurricanes exemplifying one of the most disruptive events that not only caused severe physical damage but also sparked widespread discussion on social media platforms. Existing datasets for studying societal impacts of hurricanes often focus on outdated hurricanes and are limited to a single social media platform, failing to capture the broader societal impact in today's diverse social media environment. Moreover, existing datasets annotate visual and textual content of the post separately, failing to account for the multimodal nature of social media posts. To address these gaps, we present a multiplatform and Multimodal Annotated Dataset for Societal Impact of Hurricane (MASH) that includes 59,607 relevant social media data posts from Reddit, TikTok, and YouTube. In addition, all relevant social media data posts are annotated in a multimodal approach that considers both textual and visual content on three dimensions: Humanitarian Classes, Bias Classes, and Information Integrity Classes. To our best knowledge, MASH is the first large-scale, multi-platform, multimodal, and multi-dimensionally annotated dataset centered on hurricane disasters. In addition, we introduce an online platform that supports interactive data exploration, provides preliminary analytical results, and allows users to share their insights regarding the societal impacts of hurricanes. We envision that MASH can contribute to the study of hurricanes' impact on society, such as disaster severity classification, public sentiment analysis, disaster policy making, and bias identification.

Datasets — <https://huggingface.co/datasets/YRC10/MASH>

Online Platform — <https://hurricane.web.illinois.edu>

Introduction

Natural disasters have been a persistent threat to human societies, with hurricanes standing out as one of the most representative disasters that often cause severe destruction and

*These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

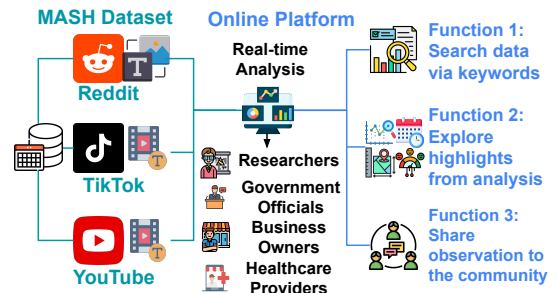


Figure 1: MASH Dataset and Online Platform

significant societal impacts (Shang et al. 2024, 2025; Alam et al. 2021; Marshall and Wang 2016). For example, Category 4 Hurricane Helene and Category 5 Hurricane Milton struck the United States in 2024, resulting in at least 250 casualties and over \$300 billion in economic losses (Paciorek et al. 2024; Wilcox and Jacobs 2024). More importantly, beyond their physical destruction, these hurricanes triggered extensive social media interactions that shaped public understanding, trust in official information, and coordination of community aid, especially among socially vulnerable groups such as low-income households, older adults, and people with disabilities. In this context, the social impact of the disaster centers on how social media communication shapes public risk perception, trust in authorities, and equitable access to relief during emergency response. In this paper, we present a comprehensive dataset that consolidates multimodal posts (e.g., text, images, videos) from multiple social media platforms to study the societal impact of hurricanes. The dataset is annotated across multiple dimensions (e.g., Humanitarian, Bias, and Information Integrity) to facilitate research on the societal impacts of hurricanes. By providing labeled data, this dataset seeks to support future studies on disaster response such as early detection of false information, identification of vulnerable communities in need of aid, and humanitarian response during disasters.

Dataset	Collection Time	Hurricanes Anno. Amount	Multi Modal	Human. Anno.	Bias Anno.	Info. Integ. Anno.	Online Platform
Sandy 2012 (Imran et al. 2013)	2012	2,000	✗	✓	✗	✗	✗
CrisisMMD (Alam, Ofli, and Imran 2018)	2017	12,041	✓	✓	✗	✗	✗
HIM-Twitter (Alam et al. 2018)	2017	50,217	✓	✓	✗	✗	✗
TweetDIS (Tekumalla and Banda 2022)	2012 - 2017	5,962	✗	✓	✗	✗	✗
Eyewitness Messages (Zahra, Imran, and Ostermann 2020)	2016 - 2018	2,000	✗	✓	✗	✗	✗
MEDIC (Alam et al. 2022)	2017 - 2018	6,686	✗	✓	✗	✗	✗
Natural Hazards (Meng and Dong 2020)	2012 - 2019	7,823	✗	✗	✗	✗	✗
HumAID (Alam et al. 2021)	2016 - 2019	45,581	✗	✓	✗	✗	✗
MASH (Ours)	2024 (latest)	59,607	✓	✓	✓	✓	✓

Table 1: Comparison with Existing Hurricane Disaster Datasets

Existing datasets often focus on hurricanes and other natural disasters that formed several years ago. For instance, many widely used datasets such as CrisisMMD (Alam, Ofli, and Imran 2018), HIM-Twitter (Alam et al. 2018), and Eyewitness Messages (Zahra, Imran, and Ostermann 2020) focus on the 2017 hurricane season, which is 8 years ago. However, social media was far less developed at the time than it is today, with a much greater number of users and platform diversity than before. As a result, these datasets have limited effectiveness in studying the societal impacts of hurricanes in the current context. In addition, previous datasets predominantly collected posts from a single platform, X (previously known as Twitter), while neglecting other widely used platforms. By excluding posts from diverse platforms, such datasets overlook critical variations in user behavior and communication styles across different social media. This narrow focus fails to capture the broader spectrum of social media activities, leading to datasets that are insufficient for analyzing the comprehensive societal impact of hurricanes. Furthermore, existing datasets either focus exclusively on the textual components of social media posts or treat textual and visual elements independently (Imran et al. 2013; Alam, Ofli, and Imran 2018; Alam et al. 2018; Tekumalla and Banda 2022; Alam et al. 2022, 2021). This fragmented treatment overlooks the inherently multi-modal nature of social media posts, thereby impeding a comprehensive multimodal understanding of social media content. This gap highlights the need for more comprehensively annotated and up-to-date datasets that incorporate multi-modal content from multiple platforms to provide a comprehensive understanding of hurricanes’ societal effects.

Motivated by the identified gaps in existing datasets, we introduce a novel multiplatform and Multimodal Annotated Dataset for Societal Impact of Hurricane (**MASH**). Figure 1 shows the overview of the MASH dataset and the online platform. In particular, the MASH dataset consists of 59,607

relevant *social media posts* collected from three platforms: Reddit, TikTok, and YouTube since they are widely used during hurricanes and provide API for large-scale data collection. All relevant posts are annotated along three dimensions: *Humanitarian Classes*, *Bias Classes*, and *Information Integrity Classes* in a multi-modal approach that jointly considers textual and visual content, providing a rich labeled dataset for in-depth analysis. The data collection and annotation framework can also be deployed to other disasters like wildfires and earthquakes. We also develop an online platform to facilitate the data exploration and real-time analysis for end users such as researchers, government officials, business owners, and healthcare providers. Specifically, the online platform has the following functions. First, it contains an insight page that can retrieve social media and news media data via customizable keywords to support the investigation of emerging disasters. Second, it contains a highlight page to show some preliminary analysis results (e.g., sentiments, word clouds) of the MASH dataset. Third, it allows users to share observations, interpretations, and community insights about the dataset and the analytical findings.

The development of the MASH dataset is the joint effort of an interdisciplinary research team from information science, computer science, hydrology, and environmental engineering. To the best of our knowledge, MASH is the first large-scale, multimodal, multi-platform, and multi-dimensionally annotated dataset centered on hurricane disasters. Beyond data construction, we further demonstrate the usefulness of MASH through a series of analyses, including correlation analysis, temporal analysis, spatial analysis, baseline model evaluation, and sentiment analysis. We envision that the comprehensiveness of platforms, the multimodality of the data, and the multidimensionality of the annotations will contribute to future studies on disaster response, bias detection, truth discovery, sentiment trajectory analysis, and disaster policy making during natural disasters.

Related Works

With the growth of web-based communication platforms and the ubiquity of the Internet, social media has become a rich resource for exploring the impact of specific events on human society. Recently, several datasets have been collected to study the societal impact of social media posts in the context of hurricanes. Sandy 2012 (Imran et al. 2013), TweetDIS (Tekumalla and Banda 2022), and NaturalHazards (Meng and Dong 2020) are datasets that cover tweets during the 2012 Hurricane Season, especially during Hurricane Sandy. In addition, CrisisMMD (Alam, Ofli, and Imran 2018), HIM-Twitter (Alam et al. 2018), TweetDIS (Tekumalla and Banda 2022), Eyewitness Messages (Zahra, Imran, and Ostermann 2020), MEDIC (Alam et al. 2022), Natural Hazards (Meng and Dong 2020), and HumAID (Alam et al. 2021) primarily capture social media discussions during the 2017 hurricane season, especially during Hurricanes Harvey, Irma, and Maria. Table 1 shows the comparison between the MASH dataset and prior datasets on hurricanes. The hurricanes that prior datasets focus on are all 8 years ago or even more than 10 years ago. With the progress of society and the development of technology, a single social media platform (X) at that time can no longer well reflect the impact of hurricanes on society. In contrast, in this study, we collected posts from three social media platforms (Reddit, TikTok, and YouTube) and focused on the recent 2024 hurricane season, especially the category 4 Hurricane Helene and category 5 Hurricane Milton. In addition, prior datasets are limited to Humanitarian Class annotation and labeled textual and visual content separately (Imran et al. 2013; Alam, Ofli, and Imran 2018; Alam et al. 2018; Tekumalla and Banda 2022; Zahra, Imran, and Ostermann 2020; Alam et al. 2022, 2021), which restricts their ability to support comprehensive analyses of hurricanes’ societal impact. In comparison, we annotated social media posts in three dimensions (Humanitarian Class, Bias Class, and Information Integrity Class) by jointly analyzing textual and visual content, enabling more robust and multimodal machine learning models for disaster-related social media analysis. Moreover, previous datasets only contain a limited set of analyses such as temporal analysis, baseline model evaluation, and sentiment analysis (Alam et al. 2022; Tekumalla and Banda 2022; Alam et al. 2018; Wang et al. 2012; Meng and Dong 2020). On the contrary, we perform systematic analyses of the MASH dataset, including correlation analysis, temporal analysis, spatial analysis, baseline model evaluation, and sentiment analysis.

Data Collection

The MASH dataset contains social media data that focuses on recent hurricanes, especially the two major hurricanes that hit the United States during the fall of 2024: *Hurricane Helene* and *Hurricane Milton*, to obtain comprehensive information about the societal impact of hurricanes. According to National Hurricane Center, Hurricane Helene was a Category 4 Storm formed on September 24, 2024 and Hurricane Milton was a Category 5 Storm formed on October 5, 2024. The combined damage from these two consecutive

	Reddit	TikTok	YouTube
Collected Raw Posts	12,301	67,027	2,767
Relevant and Annotated Posts	9,928	47,915	1,764
Avg. #Words per Annotated Post	119	21	18
#Images on Annotated Posts	1,579	—	—
#Videos on Annotated Posts	878	47,915	1,764

Table 2: Distribution of Collected and Annotated Posts

hurricanes is estimated to be more than \$300 billion (Wilcox and Jacobs 2024). To ensure timely and relevant coverage of the hurricane-related discussions, the data collection period spans from September 1 to November 30, 2024. This collection window spans approximately one month before the hurricanes and one month after the main impact period, enabling us to capture pre-disaster signals, peak impacts, and recovery-related discussions.

We collected social media posts from three platforms: Reddit, TikTok, and YouTube with a total of 82,095 posts. These three platforms are widely used by diverse demographics and communities as they provide users with different ways to disseminate information in the multi-modality form of text, images, and videos (Sihag et al. 2023). The social media data collection adopted the Reddit API, TikTok Research Tools, and YouTube API. We did not include data from X because the platform discontinued free academic API access in 2023 and replaced it with expensive paid tiers, making large-scale data collection financially infeasible for this project (Murtfeldt et al. 2024; Pehlivan et al. 2025; Bisiani et al. 2025). We utilized three keywords: *Hurricane*, *Hurricane Helene*, and *Hurricane Milton* to retrieve relevant content. To ensure the ethics of the research and protect users’ privacy, we only collected post information (i.e., post ID, post content, post time, and post location) without users’ identity information. Table 2 presents the number of posts collected from three social media platforms. For text-centric platforms like Reddit, we not only collected the title and description of the post but also collected the attached images and videos. For video-centric platforms such as TikTok and YouTube, we collected the video together with the textual data. To reduce the cost and workload for further cleaning and annotation tasks, we only collected videos from TikTok and YouTube that are less than five minutes long.

Data Annotation

The MASH dataset contains annotations for social media posts in three dimensions: *Humanitarian Class*, *Bias Class*, and *Information Integrity Class*. In natural disasters, the main purpose of humanitarian assistance is to save lives, reduce suffering, and rebuild affected communities (Alam et al. 2021). For the Humanitarian Class annotations, we defined 7 categories to classify the content of posts according to the textual and visual context of the post. The content of social media posts may contain bias that can affect user perception (Wessel et al. 2023; Liu, Li, and Wang 2024). The task of bias class annotation is to identify the presence of defined biases or discriminatory elements in 5 categories

within the post. False claims arise in environments when people face a scarcity of needed information (Muhammed T and Mathew 2022; Liu et al. 2025a). The information integrity category assesses the veracity of posts in the face of hurricane disasters, distinguishing between content with verifiable facts and incorrect content. Together, these annotations provide a comprehensive foundation for analyzing the societal impact of hurricanes on social media posts, enabling research in areas such as detection of false claims, bias, and critical events; physi-social impact forecasting; and vulnerability assessment. Prior to annotation, we performed a cleaning step to filter out irrelevant posts, ensuring that the subsequent labeling process focuses only on meaningful content. In this section, we introduce the annotation methodology, as well as the label distributions and findings.

Annotation Methodology

To reduce the time and cost of manual annotation, we leverage a Multimodal Large Language Model (MLLM) to support the annotation. We note that a single judgment by the MLLM may contain errors and not be reliable. Therefore, we design a human-MLLM collaborative framework for both data cleaning and annotation, which includes consistency checks based on multiple rounds of MLLM-generated labels and human verification of MLLM-generated labels.

In the first stage, inspired by the work of Wang et al. (2023) and Liu et al. (2026), the MLLM is queried three times per post to generate three separate labels along with corresponding explanations. The input to the MLLM consists of task-specific prompts (available in the Appendix Figures 7 to 9) and the full content of each social media post, which includes textual descriptions, and includes Images and videos if they are present. Different from prior datasets (e.g., CrisisMMD (Alam, Ofi, and Imran 2018)), which annotate textual and visual content *independently*, we treat each post as a holistic unit and assign labels based on the *joint interpretation* of text and visuals. To protect user privacy, all tagged usernames, starting with “@”, are replaced with “@user”. This multi-modal input enables the MLLM to generate labels based on a comprehensive understanding of the post across all available modalities. To measure the consistency of the MLLM’s outputs, we compute the entropy of the three MLLM-generated-labels. A zero entropy score, indicating all three predictions are identical, reflects a high degree of confidence in the MLLM’s judgment. In such cases, the label is considered high-confidence and directly adopted.

On the other hand, if the label entropy computed in the first stage is non-zero, indicating that the MLLM exhibits uncertainty regarding the task, we proceed to the second stage. In the second stage, we provide the MLLM with all three sets of labels and corresponding reasons generated in the first stage. The MLLM is prompted (available in the Appendix Figure 10) to compare, analyze, and determine which label and justification are more reasonable. This setup encourages the model to perform consistency resolution and preference selection, aligning with recent work demonstrating that MLLM benefits from self-consistency (Wang et al. 2023) and iterative self-refinement strategies (Madaan et al. 2023; Shinn et al. 2023) to improve label reliability and rea-

son quality. The MLLM subsequently generates three new sets of labels and reasons, from which the new entropy is computed. If the new entropy value is zero, the new label is also considered as a high-confidence label and accepted accordingly. If the new entropy remains non-zero, indicating persistent uncertainty, the post is escalated to the third stage for human annotation. In this stage, human annotators from diverse disciplines in our team collaboratively review the post and determine the appropriate label. Through open discussion, annotators exchange perspectives and reach consensus on the final label. This human annotation process ensures that ambiguous cases are carefully reviewed.

Furthermore, we conduct two separate human verification tasks to evaluate the reliability of the MLLM-generated labels. Firstly, for the data cleaning task, we randomly sample 500 posts from each social media platform (1,500 posts in total) from the raw collected dataset, and three annotators in our team independently verify whether each post is relevant to the disaster. Secondly, for the annotation task, we randomly sample another 500 posts from each platform (1,500 posts in total) from the cleaned relevant dataset. For every sampled post, three annotators independently verify all labels generated by the MLLM. The verification results for both tasks show high consistency and are reported in Section Annotation Consistency. We used Gemini-2.0-Flash as the MLLM in both data cleaning and annotation tasks because of its cost efficiency and advanced capabilities over multimodal inputs, which is suitable for our large-scale dataset (Hirosawa et al. 2024; Team et al. 2024).

Data Cleaning

We observe that the collected raw social media posts contain irrelevant and off-topic posts, which could reduce the quality of the dataset. For example, using the keyword “hurricane” in API search may retrieve posts about the Miami Hurricanes football team or Carolina Hurricanes hockey team, which are not related to the hurricane disaster. Additionally, some posts use hurricane-related hashtags solely for promotional purposes, such as adding #hurricane to increase exposure and attract more attention to their advertisement. These practices often result in posts with irrelevant content being included in the dataset. To address this issue, we implement the human-MLLM collaborative annotation framework to keep posts related to actual North American hurricane disasters (particularly Hurricane Helene and Hurricane Milton) and exclude content that literally mentions “hurricane” but is not relevant to an actual disaster, such as sports, entertainment, and metaphors. The detailed prompt can be found in the Appendix Figure 7. Table 2 presents the distribution of relevant posts for each social media platform after cleaning. A total of 59,607 posts were identified as relevant and subsequently annotated across the three dimensions. The results in Table 6 show high agreement both among human annotators and between human and the MLLM across all platforms, indicating the reliability of data cleaning.

Humanitarian Class Annotation

The purpose of the Humanitarian Class annotation is to classify social media posts based on their themes before, during,

	Reddit	TikTok	YouTube
Casualty	891 (9%)	4,788 (10%)	310 (18%)
Evacuation	2,475 (25%)	15,009 (31%)	728 (41%)
Damage	4,340 (44%)	22,898 (48%)	1,364 (77%)
Advice	2,079 (21%)	14,461 (30%)	334 (19%)
Request	2,750 (27%)	8,804 (18%)	416 (24%)
Assistance	3,653 (37%)	23,940 (50%)	1,015 (58%)
Recovery	4,280 (43%)	22,437 (47%)	1,122 (64%)
Other Useful Info	2,511 (25%)	8,991 (19%)	145 (8%)

Table 3: Distribution of Humanitarian Classes

and after the hurricane disaster, facilitating a semantic understanding of the data sample. This classification also supports the analysis of public communication and information dissemination during natural disasters like hurricanes. Based on previous datasets (Alam, Offi, and Imran 2018; Zahra, Imran, and Ostermann 2020; Gebru et al. 2021; Alam et al. 2018; Imran et al. 2013; Tekumalla and Banda 2022; Alam et al. 2021), we classify the social media data samples into 7 categories: **Casualty** (Cslt), **Evacuation** (Evac), **Damage** (Dmg), **Advice** (Advc), **Request** (Rqst), **Assistance** (Aid), and **Recovery** (Rcv). The detailed definition of these classes is available in the Appendix Table 9. If the content of the post does not fit into any of these categories we defined, it is classified as **Other Useful Information**.

We utilized the MLLM-human collaborative framework introduced in Section to annotate the Humanitarian Class of each post. The prompt is presented in the Appendix Figure 8 and the human verification is presented in Section Annotation Consistency. Table 3 reports the distribution of Humanitarian Classes of posts from each social media platform, with the largest value in each platform highlights in bold and the second-largest value indicates with underline. We note that a social media post may contain information that belongs to multiple Humanitarian Classes. For example, a post describing disaster situations might mention both casualties and damage of infrastructure. Consequently, the cumulative percentages of posts across all Humanitarian Classes exceed 100% in Table 3. From Table 3 we observe that the categories Damage and Recovery have high prevalence across all platforms, suggesting that users frequently upload and discuss content related to the physical destruction caused by hurricanes and the subsequent recovery efforts.

Bias Class Annotation

Bias Class annotation aims to identify the existence of bias in social media posts. Gu, Guo, and Zhuang (2021) and Liu et al. (2025b) point out that social media users tend to express more polarized and biased opinions during emergencies like natural disasters because of tension and anxiety. Based on previous studies (Wessel et al. 2023; Liu, Li, and Wang 2024; Spinde et al. 2023), we define 5 post-level bias types, **Linguistic Bias** (LB), **Political Bias** (PB), **Gender Bias** (GB), **Hate Speech** (HS), and **Racial Bias** (RB). The purpose of the Bias Class annotation is to characterize how biased or discriminatory narratives emerge during disasters,

	Reddit	TikTok	YouTube
Linguistic Bias	1,914 (19%)	13,371 (28%)	241 (14%)
Political Bias	<u>1,798 (18%)</u>	<u>5,405 (11%)</u>	<u>238 (13%)</u>
Gender Bias	190 (2%)	2,891 (6%)	55 (3%)
Hate Speech	106 (1%)	1,249 (3%)	17 (1%)
Racial Bias	144 (1%)	1,444 (3%)	27 (2%)
Undefined Bias	7,141 (72%)	31,508 (66%)	1,334 (76%)

Table 4: Distribution of Bias Classes

enabling researchers to study their patterns, impacts, and potential consequences for different communities. The detailed definition of these classes can be found in the Appendix Table 10. If the content of the post does not fit into any of the biases we defined, it is classified as **Undefined Bias**.

Since the only difference between the Bias and Humanitarian annotations is the definitions of categories, we also adopt the human-MLLM collaborative framework introduced in Section to annotate Bias Classes. The prompt is presented in the Appendix Figure 8 and the human verification is presented in Section Annotation Consistency. Table 4 reports the distribution of Bias Classes of posts from each social media platform with the largest value in each platform highlights in bold and the second-largest value indicates with underline. As illustrated in Table 4, the majority of posts are free from defined bias, while a subset of posts exhibits one or multiple types of bias. Table 4 demonstrates that all three platforms exhibit a relatively high ratio of posts classified as Linguistic Bias and Political Bias. This trend may be attributed to the strong emotional reactions triggered by hurricane disasters, which often prompt people to express their views in more intense or confrontational language, thus leading to linguistic bias. Additionally, posts about government relief efforts often reflect political leanings, especially given the close timing of Hurricanes Helene and Milton and the 2024 U.S. presidential election. This overlap in time could exacerbate political discussions and criticisms, further exacerbating the presence of political bias in these posts.

Information Integrity Class Annotation

The primary goal of the Information Integrity Class annotation task is to determine whether the content described in a post is factual or constitutes false arguments. The Information Integrity Class includes three distinct labels: *True Information*, *False Information*, and *Unverifiable Information*. The inclusion of the Unverifiable Information label accounts for posts that share personal or subjective opinions, which cannot be objectively verified. Different from Humanitarian and Bias class annotation, information integrity annotation raises unique challenges as it requires factual verification that cannot be reliably performed by the MLLM alone. Additionally, the MLLM we adopted in the annotation framework, Gemini-2.0-flash, does not contain up-to-date knowledge beyond August 2024, making it difficult to verify the factuality of posts published between September to November 2024. To address this limitation, we enabled Gemini’s online function (i.e., Grounding with Google Search) to sup-

	Reddit	TikTok	YouTube
True Info	7,304 (74%)	33,440 (70%)	1,352 (77%)
False Info	746 (7%)	8,682 (18%)	324 (18%)
Unverifiable Info	1,878 (19%)	5,793 (12%)	88 (5%)

Table 5: Distribution of Information Integrity Classes

port truth discovery. This functionality allows the model to retrieve external content based on the content of the input post, such as government announcements, truth discovery websites, and news media reports. The retrieved information serves as an external source of evidence, enabling the MLLM to assess the factuality of the post based on real-time and authoritative content. A post is labeled as *True Information* only when its core claims are explicitly supported by consistent external evidence. Posts are labeled as *False Information* when authoritative sources directly contradict the claims. When retrieved evidence is absent or insufficient to support a clear judgment, the post is labeled as *Unverifiable Information*. The prompt for the Information Integrity annotation is present in the Appendix Figure 9 and the human verification is presented in Section Annotation Consistency. Due to the strict rate limits imposed by the Grounding with Google Search, we limited the MLLM to a single round.

Table 5 reports the distribution of Information Integrity Class for each social media. Across the three platforms, the percentage of posts labeled as factual information is consistently around 75%, demonstrating a notable level of uniformity. The proportion of posts labeled as false claims on TikTok and YouTube is consistent around 18%. In contrast, Reddit exhibits a lower proportion of false claims and accompanied by a corresponding increase in the proportion of posts categorized as unverifiable information. This pattern suggests that Reddit may contain fewer false claims and more content that lacks sufficient evidence for verification.

Annotation Consistency

Similar to Manakul, Liusie, and Gales (2023), we invite three human annotators from different disciplines to manually verify the reliability of the MLLM-generated outputs. We design two tasks to verify the reliability: data cleaning verification and data annotation verification. In the data cleaning verification, for each platform, three annotators independently review 500 randomly sampled posts from the raw collected dataset (1,500 posts in total) to confirm whether the posts are related to the hurricane disaster. In the annotation verification, for each platform, we randomly sample another 500 posts (also 1,500 posts in total) from the cleaned relevant dataset. For each sampled post, annotators verify all MLLM-assigned labels across the Humanitarian, Bias, and Information Integrity Classes. For the Information Integrity Class, we focus on true and false claims, thereby reducing all evaluation tasks to binary classification. To evaluate the reliability, we measure both inter-annotator agreement among the human annotators and the agreement between the human majority vote and the output generated by MLLM. We adopt Accuracy and Fleiss’ Kappa Score for

inter-annotator evaluation. Accuracy calculates the proportion of samples that are completely consistent among all annotators, while Fleiss’ Kappa Score calculates the overall consistency after random consistency correction between annotators. Fleiss’ Kappa is suitable for measuring agreement among three or more raters, making it appropriate for evaluating inter-annotator reliability (Fleiss 1971). Subsequently, we apply majority voting to aggregate human annotations. We compute Accuracy and Cohen’s Kappa to evaluate the consistency between human consensus and MLLM’s annotations. Different from Fleiss’ Kappa, Cohen’s Kappa quantifies agreement between two raters by correcting for the proportion of agreement that would be expected purely by chance, making it more appropriate for pairwise agreement assessment compared to Fleiss’ Kappa (Cohen 1960). Table 6 shows the consistency of the annotations among the three annotators and the consistency between the MLLMs’ annotations and the human consensus. We find that human inter-annotator agreement is strong, with Fleiss’ Kappa consistently near or above 0.8. Similarly, the agreement between MLLM’s label and human consensus is also high, with Cohen’s Kappa above 0.8 and accuracy exceeding 0.9, which reflects the reliability of MLLM’s annotations. We observe that categories such as Hate Speech and Gender Bias show relatively lower agreement. A plausible explanation is their highly imbalanced label distributions, which make the metrics more sensitive to small annotation disagreements.

Dataset Availability

We follow the FAIR Data Principles (FORCE11 2020) to release the MASH dataset by only providing the IDs and annotations of the posts. This approach not only ensures that the data obtained from IDs is up-to-date, but also protects user privacy by preventing the disclosure of users’ personal identifiable information. The MASH dataset is findable with a unique Digital Object Identifier (DOI): <https://doi.org/10.5281/zenodo.17682231>. The MASH dataset is reusable under the Creative Commons Attribution 4.0 (CC BY 4.0) license, with a README file in the repository to explain the proper usage.

Preliminary Analyses

In this section, we comprehensively analyze the collected social media data and their annotations across various dimensions. Specifically, we present Correlation Analysis and Multi-modal Baseline Model Evaluation in this section. Due to page limitations, the Temporal Analysis, Spatial Analysis, Sentiment Analysis, and Text-based Baseline Model Evaluation are presented in the Appendix.

Correlation Analysis

In correlation analysis, we calculate and analyze the correlations across different annotation dimensions. In particular, we focus on understanding the connections between annotated labels in different dimensions to uncover potential patterns and interactions within the dataset. We utilize Cramer’s V to measure the correlations between the various annotated labels, where higher values indicate stronger correlations.

Category	Inter-Annotator Consistency						Human-MLLM Consistency					
	Reddit (Acc / Fleiss' κ)		TikTok (Acc / Fleiss' κ)		YouTube (Acc / Fleiss' κ)		Reddit (Acc / Cohen's κ)		TikTok (Acc / Cohen's κ)		YouTube (Acc / Cohen's κ)	
Clean	0.93	0.87	0.99	0.98	0.97	0.78	0.99	0.99	1.00	1.00	0.99	0.84
Casualty	0.93	0.81	0.98	0.92	0.97	0.92	0.98	0.92	0.99	0.94	0.98	0.94
Evacuation	0.93	0.88	0.94	0.88	0.97	0.95	0.97	0.91	0.96	0.90	0.98	0.95
Damage	0.95	0.94	0.96	0.94	0.96	0.93	0.98	0.96	0.98	0.95	0.97	0.92
Advice	0.92	0.84	0.96	0.94	0.97	0.93	0.98	0.94	0.99	0.96	0.98	0.95
Request	0.93	0.87	0.97	0.92	0.97	0.93	0.98	0.93	0.99	0.97	0.98	0.95
Assistance	0.94	0.90	0.95	0.93	0.97	0.96	0.97	0.94	0.97	0.94	0.98	0.97
Recovery	0.93	0.91	0.96	0.94	0.96	0.95	0.97	0.93	0.97	0.94	0.98	0.96
Linguistic Bias	0.93	0.85	0.94	0.89	0.97	0.91	0.99	0.97	0.97	0.91	0.99	0.94
Political Bias	0.93	0.85	0.99	0.95	0.98	0.94	0.99	0.96	0.99	0.99	0.99	0.93
Gender Bias	0.99	0.82	0.99	0.90	0.99	0.87	0.99	0.91	0.99	0.92	0.99	0.79
Hate Speech	0.99	0.66	0.99	0.95	0.99	0.83	0.99	0.86	1.00	1.00	0.99	0.86
Racial Bias	0.99	0.79	0.99	0.86	0.99	0.92	1.00	1.00	0.99	0.94	0.99	0.95
Info Integrity	0.91	0.85	0.88	0.81	0.95	0.91	0.95	0.88	0.92	0.82	0.97	0.93

Table 6: Inter-Annotator and Human-MLLM Consistency by each Category

Reddit			TikTok			YouTube											
Cslt	0.090	0.070	0.040	0.040	0.06	Cslt	0.140	0.250	0.030	0.110	0.10	Cslt	0.110	0.210	0.020	0.020	0.05
Evac	0.020	0.030	0.030	0.010	0.02	Evac	0.090	0.210	0.030	0.060	0.07	Evac	0.110	0.250	0.030	0.040	0.06
Dmg	0.070	0.080	0.020	0.000	0.00	Dmg	0.010	0.130	0.020	0.030	0.03	Dmg	0.020	0.080	0.020	0.010	0.00
Adv	0.080	0.090	0.000	0.010	0.00	Adv	0.010	0.020	0.030	0.000	0.00	Adv	0.010	0.050	0.000	0.000	0.01
Rqst	0.080	0.120	0.000	0.010	0.01	Rqst	0.060	0.220	0.020	0.060	0.06	Rqst	0.050	0.240	0.030	0.010	0.00
Aid	0.090	0.020	0.010	0.020	0.00	Aid	0.090	0.130	0.030	0.000	0.00	Aid	0.060	0.160	0.010	0.020	0.07
Rcv	0.120	0.060	0.000	0.000	0.00	Rcv	0.050	0.170	0.030	0.030	0.04	Rcv	0.030	0.140	0.000	0.000	0.04
	LB PB GB HS RB		LB PB GB HS RB		LB PB GB HS RB												

Figure 2: Humanitarian Class vs. Bias Class Correlation

	Red	TT	YT
Cslt	0.07	0.17	0.13
Evac	0.04	0.11	0.00
Dmg	0.08	0.08	0.04
Adv	0.06	0.00	0.03
Rqst	0.10	0.08	0.01
Aid	0.09	0.03	0.08
Rcv	0.12	0.02	0.09
	FI	FI	FI

(a) False Info vs. Humanitarian

	Red	TT	YT
LB	0.30	0.17	0.17
PB	0.28	0.21	0.07
GB	0.06	0.02	0.06
HS	0.12	0.12	0.02
RB	0.07	0.10	0.09
	FI	FI	FI

(b) False Info vs. Bias

Figure 3: False Info vs. Humanitarian and Bias Classes

Specifically, we calculate the correlation based on the co-occurrence of multiple labels in the same post, rather than relying on the overall distribution of posts. This approach ensures that correlation reflects the actual relationship between labels and avoids the bias caused by comparing a large number of posts in different categories (Liu, Li, and Wang 2024). Figure 2 indicates the correlation relationship between Humanitarian Classes and Bias Classes. We observe that the Request category is strongly associated with Political Bias across all platforms. One possible reason is that users may criticize government responses or call for policy-level actions to support victims. Similarly, the Casualty tends to correlate with both Political and Linguistic Bias, implying that posts discussing human loss often incorporate emotional and

polarized expression. In contrast, bias categories such as Gender Bias, Hate Speech, and Racial Bias demonstrate little correlation with humanitarian classes, likely due to their relatively low prevalence in the dataset.

Figure 3 (a) illustrates the correlations between false information and the various Humanitarian Classes. We observe that posts related to the Casualty class exhibit the strongest correlation with false information across all platforms. This highlights the prevalence of false information surrounding death and injury reports in the aftermath of disasters. Figure 3 (b) presents the correlations between false information and the Bias Classes. We observe that Linguistic and Political Biases exhibit the strongest correlations with false information across platforms, particularly on Reddit. This points out that false information is often communicated through subjective or ideologically driven language. The above observations also suggest that the presence of Linguistic and Political Bias can serve as an indicator for identifying false information in social media discussions. In contrast, Gender Bias shows consistently low correlation with false content across all platforms, indicating that gender-related content is largely absent from false information in the disaster context.

Multi-modal Baseline Model Evaluation

We evaluate the zero-shot performance of three multi-modal baseline models on our dataset to establish benchmark performance and assess their capacity to handle disaster-related multimodal posts. Specifically, we evaluate Phi-4-Multimodal-Instruct (Abouelenin et al. 2025), LLaVA-Onevision-7B (Li et al. 2024), and Qwen2.5-VL-7B-Instruct (Team 2025) on the annotated labels. These models are selected due to their strong performance on diverse vision-language tasks, as well as their popularity and adoption as standard baselines in prior works (Ataallah et al. 2024a,b; Abouelenin et al. 2025). We conduct binary classification evaluation for every category across the Humanitarian

Category	Phi-4-Multimodal-Instruct		LLaVA-Onevision-7B		Qwen2.5-VL-7B-Instruct	
	Accuracy	Macro F1	Accuracy	Macro F1	Accuracy	Macro F1
Casualty	0.9070	0.5658	0.8730	0.6117	0.9110	0.5890
Evacuation	0.7113	0.4997	0.6475	0.5780	0.7116	0.4719
Damage	0.5848	0.4794	0.6809	0.6668	0.6090	0.5218
Advice	0.7364	0.5141	0.7001	0.6052	0.7410	0.5104
Request	0.8130	0.5917	0.6915	0.5829	0.8176	0.6022
Assistance	0.6075	0.5214	0.6675	0.6625	0.6089	0.5221
Recovery	0.5934	0.4815	0.6904	0.6808	0.5980	0.4868
Linguistic Bias	0.7579	0.5518	0.7239	0.5804	0.7627	0.5428
Political Bias	0.9026	0.7108	0.8695	0.7066	0.9081	0.6919
Gender Bias	0.9368	0.5281	0.9102	0.5457	0.9461	0.5230
Hate Speech	0.9721	0.5802	0.9475	0.5725	0.9754	0.5580
Racial Bias	0.9657	0.5916	0.9241	0.5646	0.9726	0.5631
True/False Information	0.5892	0.5257	0.8353	0.5557	0.4472	0.4449

Table 7: Accuracy and Macro F1 for Different Multi-modal Baseline Models in Zero-shot Setting

tarian, Bias, and Information Integrity Classes. We adopt a set of modality-specific preprocessing steps across text, images, and videos to ensure consistent multimodal representation. For the textual content, similar to the work of Shang et al. (2024), we remove all URL links and emojis from the content. For images, we resize each image to a uniform resolution of 512×512 following standard vision-language preprocessing conventions (Liu et al. 2023; Dai et al. 2023). For videos, we follow the common practice in video-language research by sampling frames at 1 fps and selecting 8 representative keyframes for downstream analysis (Miech et al. 2019; Xu et al. 2021; Ataallah et al. 2024b). We evaluate the zero-shot performance of all baseline models, meaning that the models are not fine-tuned on our dataset and are instead prompted only with the task definitions. This setup allows us to assess the inherent capability of general-purpose multimodal models to interpret disaster-related posts without any task-specific adaptation. The experiments are conducted on NVIDIA L40 GPU, and the results are presented in Table 7. From Tables 3 to 5, we observe that most of the binary classification tasks exhibit data imbalance, with one class accounting for less than 40% of the labeled samples. Therefore, we employ two metrics for evaluation: accuracy and macro F1 score, where accuracy focuses on overall performance and the macro F1 score focuses on minority group.

For most categories, the baseline models achieve accuracy in the 0.6–0.9 range and macro F1 scores in the 0.5–0.7 range, suggesting moderate multimodal understanding but still leaving substantial room for improvement. One possible reason is that these baseline models are not fine-tuned on disaster-specific content. This limitation is most evident in the true/false information classification task, where the absence of knowledge regarding the 2024 hurricanes leads Phi-4 and Qwen2.5 to perform like random guessing. Moreover, these baseline models are smaller than the Gemini-2.0-Flash that we used for the annotation task, which together contribute to relatively low performance. Another factor lies in the data processing pipeline where we only adopt 8 keyframes to represent each video due to computational con-

straints, which leads to the loss of visual features. In addition, the lack of audio information further limits the model’s ability to capture important contextual signals, which are often important in video-based content. These limitations highlight the need for future research on disaster-specific multimodal models that can better leverage multimodal inputs to accurately interpret complex crisis scenarios. In particular, developing smaller but more efficient architectures that can natively integrate text, images, video, and audio will contribute to a more comprehensive and resource-focused understanding of disaster crises. We additionally report a supervised baseline trained on the textual content of posts in the Appendix. The results demonstrate strong performance across multiple classification tasks, indicating that the annotations are coherent and that model performance can be improved through task-specific supervision.

Conclusion and Discussion

This paper introduces a multiplatform and Multimodal Annotated Dataset for Societal Impact of Hurricane (MASH) to facilitate the understanding of societal impact of hurricanes. Specifically, the MASH dataset includes 59,607 relevant social media posts and is multimodal annotated in the three dimensions of Humanitarian Class, Bias Class, and Information Integrity Class. By leveraging annotations in three dimensions, this dataset provides new opportunities for research on the social impact of hurricanes such as the spread of false claims during disasters, the expression of public sentiment and bias, and the evolution of humanitarian needs. We implement systematic analyses to the MASH dataset including correlation analysis, temporal analysis, spatial analysis, baseline model evaluation, and sentiment analysis. Furthermore, we develop an online platform that supports data exploration, provides preliminary analysis, and enables stakeholders to share insights. The data collection, annotation framework, and online platform introduced in this paper are generalizable to other disaster contexts across diverse social media platforms. We will continue to explore methods to collect rich data on disaster impacts.

Acknowledgments

This research is supported in part by the National Science Foundation under Grant No. CNS-2427070, IIS-2331069, IIS-2202481, IIS-2130263, CNS-2131622. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- Abouelenin, A.; Ashfaq, A.; Atkinson, A.; Awadalla, H.; Bach, N.; Bao, J.; Benhaim, A.; Cai, M.; Chaudhary, V.; Chen, C.; et al. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*.
- Alam, F.; Alam, T.; Hasan, M. A.; Hasnat, A.; Imran, M.; and Offi, F. 2022. MEDIC: a multi-task learning dataset for disaster image classification. *Neural Comput. Appl.*, 35(3): 2609–2632.
- Alam, F.; Offi, F.; and Imran, M. 2018. CrisisMMD: Multimodal Twitter Datasets from Natural Disasters. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*.
- Alam, F.; Offi, F.; Imran, M.; and Aupetit, M. 2018. A Twitter Tale of Three Hurricanes: Harvey, Irma, and Maria. *Proc. of ISCRAM, Rochester, USA*.
- Alam, F.; Qazi, U.; Imran, M.; and Offi, F. 2021. Humaid: Human-annotated disaster incidents data from twitter with deep learning benchmarks. In *Proceedings of the International AAAI Conference on Web and social media*, volume 15, 933–942.
- Ataallah, K.; Shen, X.; Abdelrahman, E.; Sleiman, E.; Zhu, D.; Ding, J.; and Elhoseiny, M. 2024a. MiniGPT4-Video: Advancing Multimodal LLMs for Video Understanding with Interleaved Visual-Textual Tokens. *arXiv:2404.03413*.
- Ataallah, K.; Shen, X.; Abdelrahman, E.; Sleiman, E.; Zhuge, M.; Ding, J.; Zhu, D.; Schmidhuber, J.; and Elhoseiny, M. 2024b. Goldfish: Vision-Language Understanding of Arbitrarily Long Videos. *arXiv:2407.12679*.
- Bisiani, S.; Gulyas, A.; Wihbey, J.; and Heravi, B. 2025. UKTwtNewsCor: A Dataset of Online Local News Articles for the Study of Local News Provision. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, 2371–2384.
- Clark, K.; Luong, M.-T.; Le, Q. V.; and Manning, C. D. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In *ICLR*.
- Cohen, J. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1): 37–46.
- Dai, W.; Li, J.; Li, D.; Tiong, A.; Zhao, J.; Wang, W.; Li, B.; Fung, P. N.; and Hoi, S. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in neural information processing systems*, 36: 49250–49267.
- Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5): 378.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>.
- Gebbru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Gu, M.; Guo, H.; and Zhuang, J. 2021. Social media behavior and emotional evolution during emergency events. In *Healthcare*, volume 9, 1109. MDPI.
- Hirosawa, T.; Harada, Y.; Tokumasu, K.; Ito, T.; Suzuki, T.; and Shimizu, T. 2024. Comparative study to evaluate the accuracy of differential diagnosis lists generated by gemini advanced, gemini, and bard for a case report series analysis: cross-sectional study. *JMIR Medical Informatics*, 12: e63010.
- Imran, M.; Elbassuoni, S.; Castillo, C.; Diaz, F.; and Meier, P. 2013. Practical extraction of disaster-relevant information from social media. In *Proceedings of the 22nd international conference on world wide web*, 1021–1024.
- Jiang, Z.-H.; Yu, W.; Zhou, D.; Chen, Y.; Feng, J.; and Yan, S. 2020. Convbert: Improving bert with span-based dynamic convolution. *Advances in Neural Information Processing Systems*, 33: 12837–12848.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, 7871–7880.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Li, Y.; Liu, Z.; and Li, C. 2024. LLaVA-OneVision: Easy Visual Task Transfer. *arXiv:2408.03326*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, Y.; Li, Y.; and Wang, D. 2024. Intertwined Biases Across Social Media Spheres: Unpacking Correlations in Media Bias Dimensions. *arXiv preprint arXiv:2408.15406*.
- Liu, Y.; Liu, Y.; Li, Z.; Yao, R.; Zhang, Y.; and Wang, D. 2025a. Modality interactive mixture-of-experts for fake news detection. In *Proceedings of the ACM on Web Conference 2025*, 5139–5150.
- Liu, Y.; Liu, Y.; Mbuvi, P.; Li, Z.; Yao, R.; Lim, G.; and Wang, D. 2026. Mind the Ambiguity: Aleatoric Uncertainty Quantification in LLMs for Safe Medical Question Answering. *arXiv preprint arXiv:2601.17284*.
- Liu, Y.; Maussymbayeva, A.; Murzakhmetov, A.; Nemerenco, A.; Li, Y.; and Wang, D. 2025b. Reasoning-based uncertainty estimation for scalable multidimensional media bias annotation: A benchmark across diverse media spaces. *Knowledge-Based Systems*, 115058.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.

- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-Refine: Iterative Refinement with Self-Feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Manakul, P.; Liusie, A.; and Gales, M. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, 9004–9017.
- Marshall, J.; and Wang, D. 2016. Mood-sensitive truth discovery for reliable recommendation systems in social sensing. In *Proceedings of the 10th ACM conference on recommender systems*, 167–174.
- Mashkooor Siddiqui, S.; Sheikh, M. A.; Aleem, M.; and Singh, K. R. 2025. Comparative Analysis of Efficient Adapter-Based Fine-Tuning of State-of-the-Art Transformer Models. *arXiv e-prints*, arXiv–2501.
- Meng, L.; and Dong, Z. S. 2020. Natural hazards Twitter dataset. *arXiv preprint arXiv:2004.14456*.
- Miech, A.; Zhukov, D.; Alayrac, J.-B.; Tapaswi, M.; Laptev, I.; and Sivic, J. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2630–2640.
- Muhammed T, S.; and Mathew, S. K. 2022. The disaster of misinformation: a review of research in social media. *International journal of data science and analytics*, 13(4): 271–285.
- Murtefeldt, R.; Paik, S.; Alterman, N.; Kahveci, I.; and West, J. D. 2024. RIP Twitter API: A eulogy to its vast research contributions. *arXiv preprint arXiv:2404.07340*.
- Paciorek, E.; Zinnbauer, J.; Sweeney, K.; Schaaf, B.; and Sartori, N. 2024. 2024 Atlantic Hurricane Season.
- Pehlivan, Z.; Park, S.; Abrahams, A. S.; Desblancs, M. J. P.; Steel, B. D.; and Bridgman, A. 2025. Can-PolNews: A Multi-Platform Dataset of Political Discourse in Canada. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, 2550–2559.
- Shang, L.; Chen, B.; Liu, S.; Zhang, Y.; Zong, R.; Vora, A.; Cai, X.; Wei, N.; and Wang, D. 2025. SIDE: Socially Informed Drought Estimation Toward Understanding Societal Impact Dynamics of Environmental Crisis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 28359–28367.
- Shang, L.; Chen, B.; Vora, A.; Zhang, Y.; Cai, X.; and Wang, D. 2024. SocialDrought: A Social and News Media Driven Dataset and Analytical Platform towards Understanding Societal Impact of Drought. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, 2051–2062.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K. R.; and Yao, S. 2023. Reflexion: language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Sihag, M.; Li, Z. S.; Dash, A.; Arony, N. N.; Devathasan, K.; Ernst, N.; Albu, A. B.; and Damian, D. 2023. A data-driven approach for finding requirements relevant feedback from tiktok and youtube. In *2023 IEEE 31st International Requirements Engineering Conference (RE)*, 111–122. IEEE.
- Spinde, T.; Hinterreiter, S.; Haak, F.; Ruas, T.; Giese, H.; Meuschke, N.; and Gipp, B. 2023. The media bias taxonomy: A systematic literature review on the forms and automated detection of media bias. *arXiv preprint arXiv:2312.16148*.
- Team, G.; Georgiev, P.; Lei, V. I.; Burnell, R.; Bai, L.; Gulati, A.; Tanzer, G.; Vincent, D.; Pan, Z.; Wang, S.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Team, Q. 2025. Qwen2.5-VL.
- Tekumalla, R.; and Banda, J. M. 2022. TweetDIS: A Large Twitter Dataset for Natural Disasters Built using Weak Supervision. In *2022 IEEE International Conference on Big Data (Big Data)*, 4816–4823.
- Wang, D.; Kaplan, L.; Abdelzaher, T.; and Aggarwal, C. C. 2012. On scalability and robustness limitations of real and asymptotic confidence bounds in social sensing. In *2012 9th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 506–514. IEEE.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q. V.; Chi, E. H.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations*.
- Wessel, M.; Horych, T.; Ruas, T.; Aizawa, A.; Gipp, B.; and Spinde, T. 2023. Introducing MBIB-the first media bias identification benchmark task and dataset collection. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2765–2774.
- Wilcox, C.; and Jacobs, P. 2024. Hurricane-battered researchers assess damage. *Science (New York, N.Y.)*, 386: 367.
- Xu, H.; Ghosh, G.; Huang, P.-Y.; Okhonko, D.; Aghajanyan, A.; Metze, F.; Zettlemoyer, L.; and Feichtenhofer, C. 2021. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*.
- Zahra, K.; Imran, M.; and Ostermann, F. O. 2020. Automatic identification of eyewitness messages on twitter during disasters. *Information processing & management*, 57(1): 102107.
- Zhang, W.; Deng, Y.; Liu, B.; Pan, S.; and Bing, L. 2024. Sentiment Analysis in the Era of Large Language Models: A Reality Check. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Findings of the Association for Computational Linguistics: NAACL 2024*, 3881–3906. Mexico City, Mexico: Association for Computational Linguistics.

Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**
 - (f) Did you discuss any potential negative societal impacts of your work? **Yes**
 - (g) Did you discuss any potential misuse of your work? **Yes**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
 - (b) Have you provided justifications for all theoretical results? **NA**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
 - (e) Did you address potential biases or limitations in your theoretical framework? **NA**
 - (f) Have you related your theoretical results to the existing literature in social science? **NA**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **NA**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **NA**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? **Yes**
 - (b) Did you mention the license of the assets? **Yes**
 - (c) Did you include any new assets in the supplemental material or as a URL? **Yes**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **Yes**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? **Yes**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? **Yes**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? **NA**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **NA**
 - (d) Did you discuss how data is stored, shared, and de-identified? **NA**

Spatial Analysis

To examine the geographic distribution of posts in MASH dataset, we conduct a spatial analysis based on posts' location information. We employ two methods to extract geographic locations: (1) if a post explicitly discloses its location upon publication, we directly use this metadata; (2) if explicit location data is unavailable, we extract state-level information based on textual mentions of U.S. states within the post content. Posts for which neither method yields valid




Figure 4: Spatial Distribution of Posts Across U.S. States




Figure 5: Sentiment Spectrum of Social Media Platform

location information are excluded from this analysis. Figure 4 shows the state-level distribution of geolocated posts, reported as the proportion of posts located in each U.S. state. Darker shades indicate a higher proportion of posts originating from each state. We observe that a substantial majority of geolocated posts are concentrated in Florida and North Carolina, which aligns with the primary trajectory of Hurricane Helene and Hurricane Milton captured in the dataset. This pattern highlights the spatial relevance of the social media posts collected in our dataset.

Sentiment Analysis

We implement sentiment analysis on relevant posts from three social media platforms to examine the kinds of emotions expressed in the content. Following Zhang et al. (2024), we adopt their prompting strategy to guide the MLLM in analyzing the sentiment of each post in our dataset. Zhang et al. (2024) demonstrate that the MLLM achieves competitive performance in sentiment classification tasks, especially under few-shot or zero-shot settings, through a comprehensive test across 26 datasets. Therefore, we adopt the method because of the similarity of our task to the tested sentiment analysis task. Specifically, we let the MLLM consider the multimodal content of the post and classify the sentiment of the post into three categories: *Positive*, *Negative*, and *Neutral*. Figure 5 illustrates the distribution of sentiments across all platforms. We observe that the majority of posts showed negative or neutral sentiment, with only a small portion reflecting positive sentiment. This distribution highlights the prevalent feelings of tension and pessimism among the public when facing disasters like hurricanes. On the other hand, posts on TikTok showed the highest proportion of positive emotions, indicating that although hurricane-related content tends to be negative, TikTok users are more likely to share uplifting or supportive data samples using the platform’s visual and interactive features.

Temporal Analysis

We conduct a temporal analysis of the social media posts, focusing on the distribution of annotation labels over time. Figure 6 illustrates the time series distributions of Humanitarian Classes, Bias Classes, and Information Integrity Classes respectively. We observe a significant increase in the number of posts during the periods when Hurricane Helene and Hurricane Milton occurred, reflecting an increase in user activity in response to these events.

For the Humanitarian Classes, the categories Request and Advice are most prominent during the early stages of the hurricane, reflecting how users turn to social media to seek and offer help in anticipation of the disaster. As the hurricane strikes and its impact becomes visible, posts increasingly shift toward the Damage category, capturing reports of destruction and loss. In the aftermath, the Recovery category becomes more dominant, indicating a collective focus on rebuilding efforts.

For the Bias categories analysis, linguistic bias and political bias account for the majority of biased content. Notably, the distribution of both biases increases significantly during hurricanes. The increase in linguistic bias reflects that public discourse becomes more intense during disaster events, as individuals express stronger emotions in response to the crisis. Similarly, the increase in political bias is often driven by criticism of government relief efforts as public scrutiny of relief measures grows. These findings suggest that the impacts of hurricanes are not limited to physical dimensions, such as individual casualties and damage to infrastructure. Hurricanes also affect public discourse, highlighting the broad societal impact.

For Information Integrity Classes, the amount of both true information and false information increases significantly during the hurricanes. We observe that during the period of Hurricane Helene and Hurricane Milton, the proportion of false claims was notably high, peaking at around 40%. In the aftermath, this proportion dropped and stabilized at around 20%. This increase shows that as the level of information dissemination increases, the amount of true and false claims also increases. This increase also highlights the urgent need to verify facts during disasters, because the urgency of sharing information during disasters often leads to the spread of false claims. The spread of false claims could further cause public tension about natural disasters.

Text-based Baseline Model Evaluation

Finally, we evaluate the supervised performance of four text-based baseline models: RoBERTa (Liu et al. 2019), BART (Lewis et al. 2020), ConvBERT (Jiang et al. 2020), and ELECTRA (Clark et al. 2020) on the annotated labels. These models are selected due to their strong performance on a wide range of tasks, as well as their popularity and adoption as standard baselines in prior work (Mashkoor Siddiqui et al. 2025). Due to the substantial GPU resources and training time required for supervised multimodal models, we simplify the input to the textual content of social media posts in this evaluation. Similar to the work of Shang et al. (2024), we remove all URL links from the content. Additionally,




Figure 6: Temporal Trend of Humanitarian, Bias, and Information Integrity Classes

Category	RoBERTa		BART		ConvBERT		ELECTRA	
	Accuracy	Macro F1	Accuracy	Macro F1	Accuracy	Macro F1	Accuracy	Macro F1
Casualty	0.9306	0.7439	0.9258	0.7207	0.9303	0.7144	0.9315	0.7431
Evacuation	0.8009	0.6989	0.8116	0.7110	0.8032	0.6839	0.8040	0.6975
Damage	0.7732	0.7586	0.7664	0.7575	0.7732	0.7585	0.7710	0.7553
Advice	0.8344	0.6938	0.8382	0.7115	0.8365	0.6952	0.8326	0.6818
Request	0.8612	0.7345	0.8599	0.7441	0.8657	0.7463	0.8637	0.7474
Assistance	0.7884	0.7764	0.7888	0.7884	0.7871	0.7767	0.7874	0.7788
Recovery	0.8056	0.8035	0.7980	0.7924	0.8020	0.8001	0.8054	0.8026
Linguistic Bias	0.8395	0.7693	0.8314	0.7552	0.8423	0.7720	0.8417	0.7648
Political Bias	0.9195	0.8600	0.9216	0.8627	0.9232	0.8653	0.9208	0.8598
Gender Bias	0.9631	0.5712	0.9623	0.5776	0.9620	0.4903	0.9620	0.4903
Hate Speech	0.9769	0.6898	0.9783	0.6590	0.9785	0.6764	0.9746	0.4936
Racial Bias	0.9766	0.6947	0.9767	0.6826	0.9780	0.7085	0.9782	0.7014
True/False Information	0.8608	0.6934	0.8628	0.6801	0.8593	0.7156	0.8568	0.7201

Table 8: Accuracy and Macro F1 for Different Baseline Models

to protect user privacy, all tagged usernames, starting with “@”, are replaced with “@user”. Subsequently, we combine the posts from all four platforms into a single dataset and divide them into 70% training, 15% validation, and 15% testing sets. We train every annotated category using NVIDIA L40S GPU and the evaluation results are presented in Table 8. We employ two metrics for evaluation: accuracy and macro F1 score, one focuses on the overall performance and the other focuses on the performance of the minority group. From Tables 3 to 5, we observe that most of the binary classification tasks exhibit data imbalance, with one class accounting for less than 40% of the labeled samples. To address the class imbalance problem, we use macro F1 score as one of the evaluation metrics for the annotations because it can focus on underrepresented classes.

For most categories, the four baseline models achieve test accuracy over 0.8 and macro F1 score over 0.75, showing strong model performance. However, for some categories, such as gender bias, hate speech, and racial bias, the macro F1 scores are significantly lower. This is mainly due to the small proportion of these categories in the dataset, which is only 2-3% of the total posts. The small proportion of these categories in the dataset limited the model’s ability to receive sufficient training on these specific labels, thus affect-

ing the macro F1 performance. These findings highlight the challenges of training bias detection models in imbalanced datasets.

Prompt

Read the post, considering text, image, and video together. Determine whether the post is related to an actual hurricane disaster in North America (especially Hurricane Helene and Hurricane Milton) and give your reason.

Accepted examples:

- Hurricane disaster, even if it's not Helene or Milton

Some typical counter examples:

- Miami Hurricanes football team, Carolina Hurricanes hockey team, or other sports names

- WWE wrestler called 'Hurricane' or other people called 'Hurricane'

- Advertisement

- Use hurricane as a metaphor to describe other things

- The post contains hurricane-related hashtag, but the actual content is not related to the hurricane disaster

- Cyclone or Typhoon that does not happen in North America

Respond with:

{ "Judgment": "Judgment": True or False,

"Reason": "your detailed explanation based on the post, the examples, and the comparative reasoning between the two viewpoints." }

}

Figure 7: Prompt for Data Cleaning.

Prompt

You are reviewing a social media post related to a hurricane disaster. The post may include text, images, and videos. Carefully examine all available content together.

{ *Category Definition here* }

Respond in strict JSON format: { "Judgment": "Judgment": True or False,

"Reason": "Your detailed explanation based on the content of the post." }

Figure 8: Prompt for Humanitarian and Bias Class Annotation.

Prompt

You are reviewing a social media post related to a hurricane disaster. The post may include text, images, videos, and external links. Carefully examine all available content together. Your task is to determine whether the post contains false claims, describes truthful information, or is hard to identify due to a lack of verifiable evidence. To make your judgment, search online for reliable sources (e.g., government agencies, major news outlets, or verified fact-checking websites). Use these sources to verify any factual claims made in the post.

Label the post as one of the following:

- False Claims (-1): The post contains or promotes false or misleading information, even if only part of the content is inaccurate.

- True Claims (1): The post is accurate and supported by reliable sources.

- Hard to Identify (0): The post's truthfulness is difficult to determine (e.g., it describes personal experiences or no trustworthy sources can be found to verify it).

Please include the source(s) used to support your reasoning. Provide:

- A brief description of the source and how it supports your judgment.

- A URL link to the source (e.g., official website, fact-checking report).

Respond in strict JSON format: { "Judgment": one of the following integer values: 1 = True Claims, 0 = Hard to Identify, -1 = False Claims

"Reason": Your explanation, including supported sources and their descriptions + url links in string format }

Figure 9: Prompt for Information Integrity Annotation.

Prompt

You are reviewing a social media post related to a hurricane disaster. The post may include text, images, and videos. Carefully examine all available content together.

{Category Definition here}

For this specific post, here are two sets of arguments:

- Arguments in favor (why the post may be related to a real hurricane disaster): {true reasons here}
- Arguments against (why the post may NOT be related to a real hurricane disaster): {false reasons here}

Carefully compare and contrast both sides of the arguments above, and critically assess which side presents a more credible and reasonable explanation based on the content of the post.

Consider the relevance, specificity, and plausibility of each argument in the context of the post, and use that analysis to form your final judgment.

Respond in strict JSON format: {"Judgment": "Judgment": True or False, "Reason": "Your detailed explanation based on the content of the post." }

Figure 10: Prompt for Second Round Annotation Discussion.

Category	Definition
Casualty	The post reports people or animals who are killed, injured, or missing during the hurricane. For example: <ul style="list-style-type: none"> - Deaths resulting from the hurricane - Individuals or groups reported as missing - People or animals who are injured due to storm impact
Evacuation	The post describes the evacuation, relocation, rescue, or displacement of individuals or animals due to the hurricane. For example: <ul style="list-style-type: none"> - People leaving their homes or being urged to evacuate from at-risk areas to safer locations (e.g., temporary shelters, higher ground, or other towns) - Displacement of communities due to destruction of homes, rising water, or uninhabitable conditions - Rescue operations carried out by emergency personnel, neighbors, or volunteers
Damage	The post reports damage to infrastructure or public utilities caused by the hurricane. For example: <ul style="list-style-type: none"> - Destruction or severe damage to buildings, homes, or roads - Power outages, downed power lines, or transformers exploding - Cell towers, internet lines, or other communication infrastructure being down
Advice	The post provides advice, guidance, or suggestions related to hurricanes, including how to stay safe, protect property, or prepare for the disaster. For example: <ul style="list-style-type: none"> - Safety tips for individuals, families, or pets during a hurricane (e.g., “stay indoors,” “avoid floodwaters”) - Instructions on how to prepare emergency kits, store food/water, or stock up on supplies - Advice on how to evacuate safely, choose a shelter, or avoid traffic routes
Request	The post contains request for help, support, or resources due to the hurricane. For example: <ul style="list-style-type: none"> - People asking for rescue or evacuation assistance - Requests for food, water, medical supplies, or basic necessities - Posts seeking temporary shelter, housing, or safe relocation options
Assistance	The post contains assistance and support to victims. This includes both physical aid and emotional or psychological support provided by individuals, communities, or organizations. <p>Physical and Material Assistance:</p> <ul style="list-style-type: none"> - Distribution of relief supplies (e.g., food, water, clothing, medical kits) - Volunteers helping with debris removal, house repairs, or deliveries - Government or NGO programs providing direct support (e.g., FEMA aid, Red Cross response) <p>Emotional or Psychological Support:</p> <ul style="list-style-type: none"> - Expressions of solidarity, empathy, or moral support (e.g., “Praying for Florida,” “Stay strong New Orleans”) - Messages offering comfort, hope, or encouragement to victims and survivors - Posts acknowledging the suffering of affected communities and standing in support
Recovery	The post describes efforts or activities related to the recovery and rebuilding process after the hurricane. For example: <ul style="list-style-type: none"> - Rebuilding homes, businesses, infrastructure, or public spaces after the storm - Updates on clearing debris, restoring electricity, or repairing roads - Reopening of schools, government services, stores, or public transport

Table 9: Definition for each Humanitarian Category

Category	Definition
Linguistic Bias	The post contains biased, inappropriate, or offensive language, with a focus on word choice, tone, or expression. For example: <ul style="list-style-type: none"> - Use profanity, vulgarities, or swear words (e.g., “f**k,” “sh*t”) in a way that is excessive, aggressive, or unnecessary - Show emotionally charged language meant to belittle or ridicule victims, responders, or institutions - Contain exaggerated, alarmist, or inflammatory wording (e.g., “this hellstorm wiped out everything!” when not supported by evidence)
Political Bias	The post expresses political ideology, showing favor or disapproval toward specific political actors, parties, or policies. For example: <ul style="list-style-type: none"> - Promote or criticize political leaders (e.g., mayors, governors, presidents) in relation to the hurricane response or preparedness - Express support for or opposition to political parties or ideologies - Frame the hurricane or its response through a politicized lens
Gender Bias	The post contains biased, stereotypical, or discriminatory language or viewpoints related to gender. For example: <ul style="list-style-type: none"> - Contain sexist language, slurs, or dismissive terms toward certain genders - Reinforce traditional or harmful gender stereotypes - Make generalizations about a specific gender
Hate Speech	The post contains language that expresses hatred, hostility, or dehumanization toward a specific group or individual, especially those belonging to minority or marginalized communities. For example: <ul style="list-style-type: none"> - Racial or ethnic minorities - Immigrants or refugees - Religious minorities
Racial Bias	The post contains biased, discriminatory, or stereotypical statements directed toward one or more racial or ethnic groups. For example: <ul style="list-style-type: none"> - Black or African American communities - Latinx/Hispanic populations - Asian or Pacific Islander groups

Table 10: Definition for each Bias Category