Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Narratives, Emotions and the Logic of Disinformation During Natural Disasters: Evidence from X in the 2024 DANA in Spain

Maria Alló
Universidade da Coruña
Paolo Rungo
Universidade da Coruña
Maria L. Loureiro
maria.loureiro@usc.es

Universidade de Santiago de Compostela, Spain

Article

Keywords:

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10474639/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Narratives, Emotions and the Logic of Disinformation During Natural Disasters: Evidence from X in the 2024 DANA in Spain

Abstract

Natural disasters generate intense online communication, where messages providing practical information and mutual support coexist with political conflict and mis/disinformation. However, relatively little is known about how competing narratives evolve over time and how they shape user engagement and information diffusion. This paper examines more than 168,000 original posts published on X to study the interplay between narratives, emotions and engagement throughout the 2024 DANA floods in Spain. Tweets were classified into narrative categories using a supervised multi-label classifier based on multilingual sentence embeddings and logistic regression. We combined this classification with sentiment and emotion measures to estimate a series of regression models aimed at analyzing both user engagement and content diffusion. In the first days after the floods, most conversations focused on help, solidarity and practical information. As the crisis unfolded, however, political blame attracted disproportionately high levels of engagement, while mis/disinformation gained traction over time. Anger and hate were the emotions most strongly associated with amplification. Optimism, however, was also linked to higher engagement. Climate-related narratives followed a different trajectory. They peaked immediately after the floods, when discussion centered more heavily on possible explanations for the event, and then declined as attention shifted towards institutional responsibility. Overall, posts combining strong emotional content with political conflict received more engagement, especially when they expressed anger or assigned blame. These patterns matter for how public institutions communicate during crises and for how platforms respond to misleading and highly polarizing content.

Keywords: DANA, extreme events, misinformation, disinformation, social media, crisis communication, emotions, engagement

1. 1. Introduction

Fast-evolving disasters like flash floods shrink the public's response window to just a few hours, right when uncertainty is at its peak and verified information is scarce. In this context, digital platforms have become the go-to emergency communication infrastructure (Reynolds & Seeger, 2005; UNDRR, 2015). It is important to note that what makes them fast also makes them vulnerable to the quick spread of emotionally charged misinformation. In this sense, natural disasters create environments where information, emotions, and stories move at an extraordinary speed. After extreme events, digital platforms become key spaces for coordination, solidarity, and public communication, but also fertile ground for rumors, political blame, and emotionally charged falsehoods. A substantial body of research shows that emotionally and conflict-driven content tends to get more visibility in contexts marked by high uncertainty and strong emotional reactions, where people rely more on emotions than verification when judging the credibility of messages (Martel, Pennycook, & Rand, 2020; Pröllochs et al., 2021). Therefore, these crises offer a unique opportunity to see how informational disorders arise (Wardle & Derakhshan, 2017) and how moral and emotional signals shape the visibility of content on online networks (Brady et al., 2017).

The extreme weather event that hit eastern Spain at the end of October 2024, known locally as the DANA (Isolated Depression at High Levels), has been described as one of the deadliest floods in recent European history. Torrential rains and flash floods caused more than 230 deaths and estimated economic losses of over €17 billion (Government of Spain, 2025; Ivie, 2025). The disaster sparked massive online activity: millions of posts circulated in just a few days, mixing expressions of solidarity and grief with political accusations, misinformation, and debates about climate responsibility. In this sense, the DANA acts like a real-time observational setting: a large-scale, real-time scenario to observe how emotions and narratives influence both engagement and sharing during crises.

Although research on crisis communication and misinformation has documented the presence of false or misleading content after natural disasters and pandemics (Charquero-Ballester et al., 2021; Méndez-Muros, Alonso-González, & Pérez-Curiel, 2024), less is

known about the dynamic mechanisms through which different narratives generate interaction and gain amplification over time. Disaster research has extensively documented the beneficial uses of social media for early warning, situational awareness, and aid coordination (Vieweg et al., 2010; Kryvasheyeu et al., 2016; Simon et al., 2015; Ogie et al., 2022). It has also examined the spread of rumors during crises (Mendoza et al., 2010; Starbird et al., 2014; Oh et al., 2013) but, comparatively, little is known about misinformation as a dynamic process that evolves throughout a disaster and competes with prosocial communication for attention. In this regard, much of the existing literature is based on static snapshots or descriptive typologies, without modeling how engagement changes over time or how emotions interact with narrative frameworks. At the same time, research on moral-emotional contagion (Brady et al., 2017) suggests that anger and outrage are especially powerful drivers of visibility on social media, although empirical evidence from real disasters is still scarce. This aligns with studies indicating that proximity and emotional intensity drive the spread of misinformation during disasters (Huang et al., 2015), suggesting that outrage may act as the affective vehicle that platform amplification selectively rewards.

Engagement with content does not necessarily imply its dissemination through the network. While some narratives may attract reactions, others are more likely to be retransmitted and amplified through retweet dynamics.

The 2024 DANA crisis in Spain is particularly relevant because it has already been identified as a paradigmatic case of climate disinformation and diagonal political contestation in the Spanish public sphere (López-Carrión & Llorca-Abad, 2025). This article extends that line of work by focusing on engagement and diffusion on Twitter/X, emotional dynamics, and the temporal trajectory of competing narratives. Twitter/X has experienced a decline in users, particularly in some markets. However, it remained one of the largest online platforms in Europe during the study period, with approximately 105 million monthly active users in the European Union according to its Digital Services Act transparency reports.

Methodologically, we analyze more than 168,000 tweets posted during and after the 2024 DANA. We combine a supervised multi-label classifier of narratives, sentiment and emotion measures, and dynamic regression models to analyze how different narratives

and emotions shape two complementary dimensions of online communication: interaction and amplification. These models explain how engagement and diffusion unfold across distinct storylines: help and solidarity, political blame, mis/disinformation, climate-related explanations, useful or factual information and residual content. The main findings can be summarized as follows: (1) political blame, solidarity and mis/disinformation narratives generate significantly higher engagement and propagation than practical information, although amplification dynamics increasingly favor conflict-driven narratives over time; (2) strong emotions, such as anger, optimism and hate are associated with substantially higher engagement and transmission than other emotions; and (3) engagement and amplification follow a non-linear moral–emotional trajectory whereby prosocial content dominates initially but politicization and mis/disinformation intensify over time and sustain attention.

The article contributes to existing literature on three levels. First, it shows with large-scale data how emotionally intense and conflict-driven narratives progressively displace prosocial communication in terms of engagement and amplification. Second, it connects the framework of information disorder (Wardle & Derakhshan, 2017) with the literature on moral–emotional contagion (Brady et al., 2017). Third, it introduces dynamic models of engagement and propagation integrating narrative, emotional and temporal dimensions. The remainder of the paper is structured as follows: Section 2 reviews the relevant literature; Section 3 describes data and methods; Section 4 reports the results; Section 5 discusses their implications; and Section 6 concludes.

2. Literature Review

Extreme events (for example, hurricanes, wildfires, terrorist attacks, and epidemics) create an information environment with high uncertainty, where the demand for practical guidance exceeds the capacity of official channels to confirm and disseminate updates. Research shows that in these early stages, social media users engage in rapid sense-making: they share observations, interpret fragments of evidence, and coordinate help, while at the same time amplifying unverified claims that can turn into rumors and misinformation when information gaps persist (Palen et al., 2007). This helps explain why misinformation isn't simply an “extra” during disasters, but an inherent feature of digitally mediated crisis communication.

Within this context, empirical studies show that there are systematic asymmetries between misinformation and its correction. For example, after the Boston Marathon bombing in 2013, misinformation spread in much larger volumes than corrective messages, and the corrections often arrived late and failed to penetrate the same conversation channels in a significant way (Starbird et al., 2014). These crisis-specific dynamics align with broader evidence that falsehoods can spread faster and more widely than true information on social platforms, partly due to novelty and emotional engagement, conditions that are often intensified during emergencies (Vosoughi et al., 2018). Another line of literature emphasizes that misinformation during crises is multimodal and often visual. During Hurricane Sandy, much of the widely shared 'evidence' consisted of fabricated or out-of-context images, and its spread was largely driven by a small fraction of very visible accounts and retweet dynamics, rather than just the follower network structure (Gupta et al., 2013). This highlights a recurring operational problem for emergency management: manipulated visuals or images with incorrect captions can serve as high-credibility anchors for false narratives, complicating verification and speeding up behavioral responses (e.g., unnecessary travel, panic buying, or misallocation of resources).

During public health emergencies, the World Health Organization (WHO) has defined the problem as an "infodemic," that is, an overload of information of varying quality that makes it hard for people to find reliable guidance and can weaken response measures (World Health Organization, n.d.; World Health Organization, 2021). Discourse analyses on COVID-19 across different platforms show that the volumes of misinformation and its amplification vary by platform, suggesting that governance regimes, recommendation systems, and audience composition jointly shape the risk of an infodemic (Cinelli et al., 2020). As a result, research on mitigation increasingly tests scalable behavioral interventions: “accuracy nudges” that redirect attention to truthfulness can measurably reduce the willingness to share low-quality content, an effect consistent with explanations for sharing misinformation under cognitive load and limited attention (Pennycook et al., 2021).

3. Data and methods

3.1. Data collection and scope

The corpus used in this study contains all tweets that explicitly mentioned the term "DANA" (the Spanish acronym for Depresión Aislada en Niveles Altos) during the period corresponding to the Valencia floods of late 2024 (October 2024–April 2025). After removing duplicates, replies, and retweets, we obtained 172,184 original tweets, of which 168,670 contained sufficient text (minimum four tokens) to be encoded using semantic sentence embeddings. This dataset captures the full spectrum of discourse on X during a natural disaster with meteorological, informational and political dimensions.

Because the acronym "DANA" tends to dominate the lexical space and can induce semantic collapse in the embedding space, we created a dedicated text field where all occurrences of the term and its variants (for example, "post-DANA", "DANA_Valencia") were removed. This allowed us to distinguish thematic content beyond the mere mention of the event. And hence, it allows the analysis to focus on what users say about the event—such as causes, responsibilities or emotions—rather than on the mere mention of the event itself. This step improves the ability of semantic models to distinguish between different narratives within the same event.

We produced two parallel text representations. The first is a minimal-clean text, converted to lowercase and stripped of URLs, mentions, hashtags (while keeping their lexical core), emojis and punctuation; this version was used for embeddings. The second is a normalized text with more extensive lexical standardization and removal of numeric clutter; it was used exclusively for TF–IDF analysis and diagnostics. The cleaning strategy sought to preserve the semantic integrity of the tweets while reducing noise, arising from the orthographic variability of the Spanish language used in social media.

3.2. Annotation process and narrative definition

To capture the main discursive frames, we constructed a multi-label taxonomy in which each tweet could belong to several categories. The annotation process involved three

steps: (1) pilot definition, in which two researchers reviewed a random sample of 500 tweets to propose initial categories; (2) dual manual coding of an expanded sample (n = 800), followed by reconciliation to create a gold standard; and (3) assessment of inter-coder reliability using Cohen's kappa, which ranged from 0.47 (*Practical information*) to 0.87 (*Climate-related cause* and Bonaire hoax[1]), with an overall macro-kappa of approximately 0.62.

The resulting categories were:

•Climate-related cause: attribution of the DANA to climate change or other environmental drivers.
•Political blame: attribution of responsibility or blame to political or institutional actors.
•Mis/disinformation: false rumors and deceptive content, aggregating subtypes such as hoaxes about emergency services (112), dams, conspiracies (HAARP, chemtrails) and local hoaxes.
•Help/solidarity: expressions of help, empathy or calls for donations.
•Practical information: practical information, warnings and service messages.
• Other/not classified: residual or ambiguous messages.

The manually annotated gold standard served as the reference dataset for supervised learning. The supervised machine-learning models that automatically assign each tweet to one of the narrative categories are called classifiers. These models take the vector representation of each tweet as input and learn statistical decision boundaries that distinguish one narrative from the others. In our case, we implement a set of one-vs-rest logistic regression classifiers, each specialized in detecting a specific narrative. Once trained on the gold-standard sample, these classifiers can generalize the patterns learned from the manually annotated tweets to the full dataset. Each classifier was trained to recognize one narrative category (e.g. *political blame, mis/disinformation, climate-related cause, help/solidarity, practical information*) under a one-vs-rest scheme. The model learned the linguistic and semantic patterns associated with each label from the gold standard examples and was subsequently applied to the entire corpus to predict

[1] The Bonaire hoax was a massive viral disinformation campaign claiming that hundred of bodies were buried inside a parking garage at the Bonaire shopping center in Aldaia, Valencia.

narrative probabilities. The annotated data were subsequently divided into training and held-out validation subsets, with the latter used to evaluate model performance and calibrate label-specific decision thresholds.

Manual inspection of out-of-sample predictions suggested acceptable accuracy but also room for improvement in borderline cases. We therefore drew a second stratified validation sample, manually reviewed misclassified tweets, and merged them with the original annotations, resulting in an expanded gold standard (n = 1,274). For reasons of model stability, the *Other/not classified* category was excluded from training—due to its heterogeneity—and defined as a residual class at inference (assigned when no other label exceeded its threshold).

3.3. Semantic representation and classifier performance

All tweets were encoded using the multilingual BGE-M3 model (Chen et al., 2024), which generates 1,024-dimensional sentence embeddings, optimised for capturing semantic similarity across languages. Each embedding vector was L2-normalized before training, where "L2" refers to the Euclidean norm of the vector. In practice, this means that all embedding vectors were rescaled to have unit length, so that similarity comparisons depend on the direction of the vector (semantic content) rather than on differences in magnitude, preventing longer or more information-dense texts from exerting disproportionate influence on the classifier.

We evaluated two supervised multi-label approaches, both based on one-vs-rest logistic regression classifiers: a TF–IDF baseline (word bigrams and character 3–5 n-grams) and an embedding-based model using BGE-M3 representations. The embedding model achieved clearly superior predictive performance across standard evaluation metrics, particularly in its ability to capture semantically nuanced narratives. For this reason, we adopted it as our main classification strategy.

The selected model was then retrained on the expanded gold-standard dataset (80/20 train–validation split). The final classifier achieved strong out-of-sample performance

(macro-F1[2] = 0.82; macro-AUC[3] = 0.96), indicating a high capacity to distinguish narrative categories even in a noisy social-media environment. Performance was consistently high across labels, confirming the model's ability to disentangle overlapping frames with a balanced combination of precision and recall.

Because logistic regression produces continuous decision scores, we calibrated label-specific classification thresholds on the validation set. Tweets not exceeding any threshold were assigned to an "*Other/not classified*" category.

To assess out-of-sample reliability, we conducted a balanced human validation exercise on a 500-tweet sample (approximately 80 per main label plus 100 "Other/not classified"). Two annotators reviewed the model's predictions and compared them with human revisions. Agreement levels were high: overall micro-agreement reached 93.3%, with a macro-kappa[4] of 0.79 and a macro-F1 of 0.84. Performance was particularly strong for climate-related, help/solidarity, and informational tweets. For the mis/disinformation category, we prioritized recall over precision to minimise false negatives and ensure that potentially harmful content was not under-detected.

3.4. Engagement, emotions and statistical model

For each tweet, we computed a standardized engagement rate defined as follows:

$$engage_rate = (retweets\ count + replies\ count + likes\ count + quotes\ count) / views\ count \quad (1)$$

Summary statistics and full definitions of the variables in (1) can be found in Table 1. This approach follows prior work in social media analytics, which defines engagement rate based on the number of observed user interactions relative to the number of times content is seen by users, rather than relative to follower count alone (Tomblinson et al., 2019). In addition, we complement the engagement analysis with retweet counts to capture content diffusion and dissemination across the network. While engagement

[2] Macro-F1 reflects the balanced classification performance across categories.
[3] Macro-AUC measures the model's overall discriminative capacity.
[4] Macro-kappa reflects agreement with human annotations corrected for chance.

reflects user interaction with contents, retweets provide a direct measure of diffusion dynamics.

Our X metadata included the number of followers and accounts followed, as well as counts of likes, quotes, replies, retweets, and views. We also constructed a temporal variable measuring the number of days since the event, centered on the day of maximum flood intensity.

Emotional variables were derived from two sources: (1) an overall sentiment score using Valence Aware Dictionary and sEntiment Reasoner (VADER), and (2) emotion-specific scores for anger, hate, joy, optimism, and sadness, capturing the relative intensity of each emotional dimension in the text. Regarding VADER (Hutto and Gilbert, 2014), it uses a combination of linguistic rules and heuristics to find the intensity of the sentiment in the text. The first step was to automatically translate the entire corpus using the Helsinki-NLP opus-mt-es-en model, a neural translation model (MarianMT) specialized in the Spanish-English domain. VADER is especially suited for the language of social media, and it consists of a list of words that are labeled with a positive or negative valence score, as well as a list of booster words that can change the intensity of sentiment. To calculate the overall sentiment, the formula considers the valence scores of each individual word in the text, as well as their grammatical relationships and the existence of booster words or negations. The result provided is a normalized score between -1 and +1, where -1 represents an extremely negative sentiment and +1, an extremely positive sentiment, with 0 indicating a neutral sentiment. The transformer-based framework TweetNLP was used to obtain the emotion scores. TweetNLP gives us probabilistic scores for different emotion categories, capturing the relative intensity of each emotional dimension expressed in tweets. These emotion scores correspond to the probabilities predicted by the model and they are associated with each emotion label, reflecting the likelihood and relative prominence of each emotional state in the text. The descriptive statistics reported in Table 1 summarize the key variables and their means. Overall, the mean engagement rate is 0.0177; the average number of retweets is 16.76, the average sentiment is slightly negative (VADER ≈ -0.0314); and levels of sadness and anger are relatively high, in line with the tragic nature of the event.

On this basis, we estimated a series of models: Ordinary Least Squares (OLS), Tobit and Ordered Logit (OLogit) models with engagement rate as the dependent variable. In the case of retweet count, we estimate a Negative Binomial due to the count nature and strong over-dispersion of the retweet distribution, a common characteristic of social media data in which a small number of posts receive disproportionate visibility and amplification. In addition, we include the logarithm of view counts as an offset. This allows us to distinguish between content that receives many retweets simply because it is widely viewed and content that is disproportionately amplified relative to its level of visibility. All models include followers count and a temporal variable as controls. Successive specifications introduce: (1) narrative dummies; (2) emotional scores with sadness as the omitted category; (3) a combined specification of narratives and emotions; and (4) a dynamic model with interactions between time and narrative type[5]. Summary statistics are presented in Table 1.

Table 1: Summary statistics

Variable	Description	Mean
Engagement rate	(Retweets + Replies + Likes + Quotes) / Views	0.017
Twitter metadata variables		
Followers count	Number of followers.	53708.721
Following count	Number of accounts followed.	2001.332
Like count	Number of likes.	40.573
Quote count	Number of quote tweets.	0.835
Reply count	Number of replies.	2.493
Retweet count	Number of retweets received.	16.769
View count	Number of times the tweet was viewed.	3422.064
Emotional variables		
Anger score	Intensity of anger-related expressions in the text.	0.269
Hate score	Intensity of hate-related expressions in the text.	0.031

[5] Although VADER provides a general measure of sentiment polarity, it is moderately correlated with the emotion-specific variables included in the analysis. To avoid multicollinearity and improve interpretability, we focus on discrete emotional dimensions.

Joy score	Intensity of joy-related expressions in the text.	0.176
Optimism score	Intensity of optimism-related expressions in the text.	0.202
Sadness score	Intensity of sadness-related expressions in the text.	0.351
Narratives		
Help/solidarity	Indicates whether the tweet belongs to the "help and solidarity" narrative.	0.247
Political blame	Indicates whether the tweet frames responsibility or blame on political actors.	0.212
Misinformation /disinformation[6]	Identify tweets that reproduce, endorse, or discuss false, misleading, manipulated, or unverified claims related to the disaster, its causes, consequences, institutional response, or affected population.	0.105
Climate-related cause	Classification of tweets linking the event to climate change or environmental drivers.	0.037
Practical information	Indicates whether the tweet provides practical or operational information.	0.154
Other / not classified	Category for tweets that do not fit any of the main narratives.	0.396
Time variables		
Days since the event	Number of days since the event.	20.020

4. Results

4.1. Narrative prevalence

Applying the calibrated multi-label model to the full corpus (n = 168,670) yields the following prevalence of narratives, which serve as inputs for subsequent engagement models (See Table 2):

[6] We group under the same label "misinformation/disinformation" to identify content that contains false, misleading, or unfounded claims. The reason is that distinguishing between them requires evidence about the author's goal, which cannot be reliably inferred from the content of a tweet alone. Both phenomena involve the circulation of false, misleading, or insufficiently substantiated claims and may produce similar effects on public understanding and online debate.

Table 2: Narrative prevalence

Narrative	Percentage
Help/solidarity	24.76
Political blame	21.20
Mis/disinformation	10.56
Climate-related cause	3.07
Practical information	15.37
Other / not classified	39.64

A substantial proportion of tweets (39.64%) is classified as "*Other / not classified*". This category captures heterogeneous content that does not clearly fit into predefined narrative frames, including general reactions, emotional expressions, vague comments, humor or irony, and context-dependent messages.

Such a proportion is consistent with previous analyses of social media discourse related to crises, where a large part of the content consists of contextual, expressive messages, or loosely framed posts rather than clearly defined thematic narratives (Palen & Anderson, 2016). The next most prevalent categories are *help/solidar*ity (24.76), and *political blame* (21.20 %). *Mis/disinformation* accounts for just over 10 % of the corpus, *practical information* around 15.37%, while explicit *climate-related* references are comparatively rare (3.07 %).

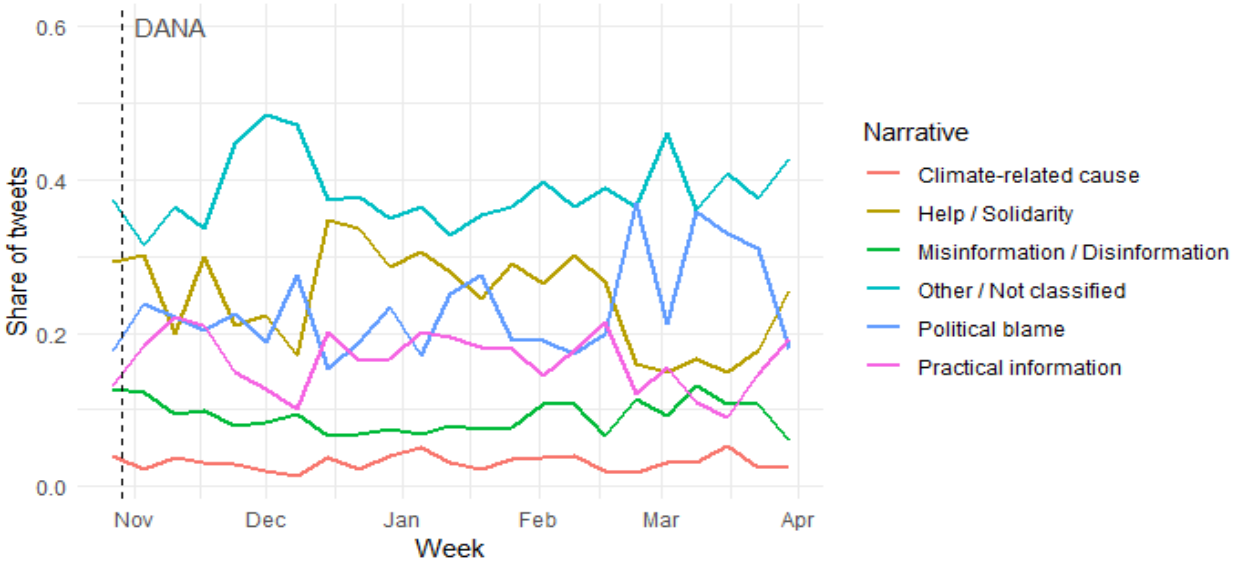
Figure 1 shows the temporal evolution of the narratives, suggesting that the trend has not been a gradual decline, but rather a sequence of reactivations driven by events. The first peak, during the week of the DANA, reflects the immediate emergency phase, in which the narratives *help/solidarity*, *practical information* and *climate-related cause* were the most prominent. Help/solidarity and useful information dominated the first days of the crisis, reflecting the immediate need for assistance and coordination. By contrast, climate-related tweets reflected an early effort to explain why the DANA had occurred, aligning with scientists who linked rainfall to climate change in public interviews.

In the following week, political blame remained high, in line with the public controversy surrounding the visit of national and regional authorities to Paiporta, a municipality at the epicenter of the flash floods, on 3 November. During the same period, the World Weather Attribution initiative published a report, which concluded that climate change played a role in increasing the likelihood and intensity of the event. The first protests over the perceived failures in emergency management started in the week beginning 10 November. Political blame was especially high in that period, when thousands of people gathered in Valencia, many calling for the resignation of Carlos Mazón, the president of the Valencian Government.

The Figure also shows a second phase of reactivation around late November and early December, coinciding with the one-month anniversary of the disaster, renewed protests, debates over aid and reconstruction, and media attention to institutions. For example, on November 27, Prime Minister Pedro Sánchez announced an additional aid package, defended the link between climate change and the floods, and criticized the regional management. On November 30, a second protest was called in Valencia against the handling of the DANA. Regarding the climate narrative, its modest resurgence is consistent with a technical report by AEMET, the Spanish Agency of Meteorology, and the broader public debate on climate preparedness.

Another large peak is observed in the week of 8 December, which seems to be less connected to the physical emergency itself than to the political reconstruction of responsibility. At the same time, public explanations and media coverage reopened questions about warning decisions and institutional coordination. The slight surge in political blame during the week of 29 December aligns with the third major protest at the end of the year. The Spanish outlet El País reported 80,000 demonstrators in Valencia demanding the resignation of the regional president and faster aid delivery. From January, conversations enter a low-intensity “long-tail”, where political blame remains as the most persistent narrative, suggesting that accountability dynamics outlasted the immediate humanitarian and informational phase.

Figure 1: Temporal evolution of narratives.



Considering the large number of tweets in the category "*Other / not classified*", we explicitly analyze whether this residual group may contain latent misinformation content. Thus, we conduct a comparative analysis between this group of tweets and those classified as explicit *mis/disinformation*. We focus our attention on engagement metrics, sentiments and emotions. Table 3 shows the main results.

Table 3: Category "mis/disinformation" vs. "other/not classified"

	Mis/disinformation	**Other/not classified**
Average RT count	22.16	14.6
Average Like count	52.09	37.2
Average Reply count	2.62	1.90
Average View count	4013	3426
VADER sentiment (average)		
Negative	0.14	0.09
Positive	0.05	0.07
Average sentiment	-0.29	-0.06
NLP (average)		
Negative	0.58	0.31
Positive	0.04	0.16

Neutral	0.36	0.52
VADER sentiment (% of tweets)		
Negative	62.88%	30.93%
Neutral	17.80%	17.72%
Positive	19.52%	22.82%
NLP sentiment (% of tweets)		
Negative	67.15%	23.05%
Neutral	30.93%	39.38%
Positive	1.93%	9.03%
NLP emotions (% of tweets)		
Anger	48.14%	12.46%
Joy	8.67%	19.26%
Optimism	4.89%	10.16%
Sadness	38.3%	29.58%

Table 3 shows that with respect to engagement patterns, tweets included in the category *mis/disinformation* show higher levels of user interaction than those included in the residual category (*other/not classified*). On average, *misinformation* tweets receive more retweets (22.16 vs. 14.6), likes (52.09 vs. 37.2), replies (2.62 vs. 1.9), and views (4013 vs. 3426). Thus, these differences suggest that, despite representing a smaller share of total tweets*, mis/disinformation*-related content achieves greater visibility and generates stronger audience engagement.

Considering the sentiment analysis (VADER scores), *mis/disinformation* tweets are substantially more negative, with an average compound sentiment of −0.288, compared to −0.067 for tweets in the "*Other/not classified*" category, which are close to neutral. Using the NLP-based sentiment classifier, results show that *mis/disinformation* tweets have a much higher proportion of negative sentiment (67.1%) and a very limited presence of positive sentiment (2%). In contrast, tweets classified as "*Other/not classified*" show lower negativity (23%), higher positivity (9%), and a higher proportion of neutral content (39.4%). If we look at the distribution of sentiment, we can see that negative sentiment

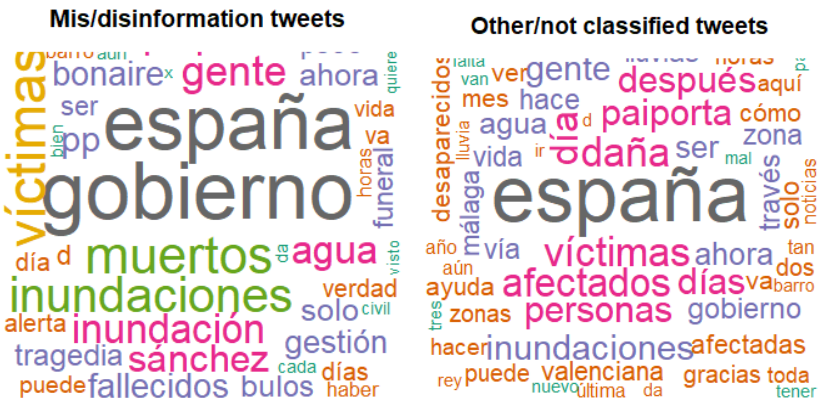
in *mis/disinformation* tweets is more than double that observed in the residual category (*other/not classified*).

With the analysis of emotions, we observe pronounced differences between the two groups: *Mis/disinformation* tweets display an emotional profile strongly driven by anger (anger accounts for 48.14% of emotional intensity), along with high levels of sadness (38.3%). Positive emotions such as joy (8.67%) and optimism (4.89%) are scarce in comparison. In contrast, tweets in the "*Other / not classified*" category show significantly lower levels of anger (12.46%) and a more balanced emotional composition, with a higher proportion of joy (19.26%) and optimism (10.16%), while sadness remains present but less dominant (29.58%).

These results support a conservative interpretation of the classification strategy, indicating that the "*Other/not classified*" category primarily captures heterogeneous and emotionally diffuse content rather than latent *mis/disinformation*.

Finally, we conducted a word cloud analysis (Figure 2) in order to provide additional qualitative support for the distinction between the two categories. *Mis/disinformation* tweets are characterized by highly politicized and alarmist language, with frequent references to government responsibility, fatalities, and crisis management (government, deaths, deceased, management, flood, alert, tragedy, hoaxes, truth, Sánchez). By contrast, the residual category (*other/not classified*) is dominated by descriptive and contextual terms related to the event itself and its general consequences, lacking the explicit claims or accusatory framing (Spain, DANA, victims, affected, people, days, areas, water, floods) observed in *mis/disinformation tweets*. Thus, we reinforce the interpretation of our previous results.

Figure 2: Word clouds mis/information and other/not classified



4.2. Narrative effects on engagement and diffusion

Once the narrative categories are presented, the next step is to assess whether the identified categories reflect distinct patterns of user engagement. To evaluate the robustness of the estimated effects, we employ different modeling strategies. For the engagement rate, we first estimate an OLS model using the continuous engagement rate as the dependent variable. Second, we estimate a Tobit model to explicitly account for the censoring at zero inherent to engagement data. Third, we estimate an ordered logit model based on an ordinal categorization of engagement levels. The results from the baseline OLS specification are presented in Table 4 (first three columns). For the retweet count we estimate a Negative Binomial and results are in Table 5.

Table 4. Engagement rate: Combined results for narrative, emotional, sentiment and temporal dynamics models

	Narrative effects			Emotional effects			Narrative and emotional effects			Temporal dynamics		
Variable	**OLS**	**Tobit**	**OLogit**	**OLS**	**Tobit**	**OLogit**	**OLS**	**Tobit**	**Ologit**	**OLS**	**Tobit**	**OLogit**
(Intercept)	0.0083***	-0.0036***		0.0109***	0.0017***		0.0057***	-0.0045***		0.0069***	-0.0027***	
	0.0003	0.0005		0.0002	0.0003		0.0003	0.0005		0.0004	0.0005	
Days since the event	0.0001***	0.0001***	0.0040***	0.0001***	0.0001***	0.0041***	0.0001***	0.0001***	0.0043***	0.0001***	0.0000***	0.0013**
	0.0000	0.0000	0.0002	0.0000	0.0000	0.0002	0.0000	0.000	0.0001	0.0000	0.0000	0.0007
Help/solidarity	0.0085***	0.0124***	0.7790***				0.0083***	0.0111***	0.7454***	0.0081***	0.0107***	0.7826***
	0.0004	0.0005	0.0208				0.0003	0.0004	0.0217	0.0004	0.0005	0.0266
Climate-related cause	0.0059***	0.0071***	0.4175***				0.0052***	0.0059***	0.3766***	0.0047***	0.0047***	0.3177***
	0.0006	0.0008	0.0341				0.0005	0.0007	0.0356	0.0006	0.0009	0.0424
Political blame	0.0115***	0.0109***	0.5187***				0.0090***	0.0089***	0.4603***	0.0060***	0.0044***	0.3497***
	0.0004	0.0005	0.0221				0.0003	0.0005	0.0241	0.0004	0.0006	0.0294
Mis/ disinformation	0.0069***	0.0071***	0.2838***				0.0047***	0.0053***	0.2495***	0.0021***	0.0018**	0.1239***
	0.0004	0.0006	0.0238				0.0004	0.0005	0.0255	0.0004	0.0006	0.0306
Other/ not classified	0.0045***	0.0037***	0.2260***				0.0038***	0.0042***	0.2574***	0.0022***	0.0018***	0.1637***
	0.0003	0.0005	0.0201				0.0003	0.0004	0.0214	0.0004	0.0005	0.0261
Anger score				0.0065***	0.0036***	0.1949***	0.0050***	0.0032***	0.2544***	0.0052***	0.0036***	0.2731***
				0.0004	0.0005	0.0226	0.0003	0.0005	0.0246	0.0003	0.0005	0.0247
Joy score				0.0002	-0.0001	0.0525***	0.0003	0.0002	0.0769***	0.0007*	0.0009	0.1038**
				0.0004	0.0005	0.022	0.0003	0.0004	0.0223	0.0003	0.0004	0.0224
Optimism score				0.0079***	0.0113***	0.8505***	0.0057***	0.0077***	0.5861***	0.0062***	0.0084***	0.6112***
				0.0004	0.0005	0.025	0.0004	0.0005	0.0261	0.0004	0.0005	0.0262
Hate score				0.0054**	0.0083**	0.2477**	0.0077***	0.0126***	0.6112***	0.0084***	0.0136***	0.6522***
				0.002	0.0027	0.1252	0.0020	0.0027	0.1267	0.0020	0.0027	0.1269
Days*help/solidarity										0.0000	0.00001	-0.0024***

										0.0000	0.0000	0.0008
Days*Climate-related cause										0.0000	0.00006	0.0003**
										0.0000	0.0000	0.0013
Days*political blame										0.0001***	0.0002***	0.0052***
										0.0000	0.0000	0.0008
Days*mis/ disinformation										0.0001***	0.0002***	0.0071***
										0.0000	0.0000	0.0009
Days*other/not classified										0.0001***	0.0001***	0.0048***
										0.0000	0.0000	0.0007
Log(scale)		-3.1179***			-3.2243***			-3.2286***			-3.23026***	
		0.0023			0.0024			0.0024			0.00241	
R-squared	0.021			0.021			0.028			0.030		
F	596.2			613			413.7			299.5		
p-value	0			0			0.000			0.000		
Wald-statistic		1991			1398			2325			2700	
p-value		0			0			0.000			0.000	
Log-likelihood		139764.4	-157335.8			-132642.6			-131681.9			-131517.9
AIC	-666229.7	-279512.7	314687.6	-582872.9	-269923.8	265299.18	-583908.4	-270836.4	263387.8	-584244.1	-271185.9	263069.7

Notes: Coefficients are shown with standard errors in parentheses. Sadness score is the omitted emotion category. Practical information is the reference category for narrative variables. Significance levels: * p<0.10, ** p<0.05, ***p<0.01.

Table 5. Diffusion (Retweet count): Combined results for narrative, emotional, sentiment and temporal dynamics models

	Narrative effects	Emotional effects	Narrative and emotional	Temporal dynamics
Intercept	-5.7369***	-5.5808***	-5.8258***	-5.7765**
	0.0166	0.0108	0.0187	0.0215
Days since the evento	0.0090***	0.0094***	0.0094***	0.0057**
	0.0002	0.0002	0.0002	0.0006
Help/Solidarity	0.4689***		0.4717***	0.4980***
	0.0183		0.0185	0.0223
Climate-related cause	0.2555***		0.2082***	0.1926**
	0.0308		0.0310	0.0364
Political blame	0.6358***		0.4929***	0.3778***
	0.0196		0.0207	0.0249
Mis/disinformation	0.3315***		0.1975***	0.0066
	0.0210		0.0217	0.0256
Other/not classified	0.0354*		0.0959***	0.0117
	0.0178		0.0182	0.0219
Anger score		0.3978***	0.3113***	0.3200***
		0.0198	0.0215	0.0215
Joy score		-0.151***	-0.1159***	-0.088***
		0.0196	0.0195	0.0196
Optimism score		0.3838***	0.2094***	0.2364***
		0.0220	0.0229	0.0229
Hate score		0.0583	0.3527**	0.3964***
		0.1106	0.1102	0.1100
Days*help/Solidarity				-0.0021**
				0.0007
Days*climate-related cause				0.0006
				0.0012
Days*political blame				0.0059**
				0.0007
Days*mis/disinformation				0.0112**
				0.0008
Days*other/not classified				0.0048**
				0.0007
AIC	555127	503517	501776	501289
Theta	0.5224	0.5425	0.5581	0.5618
Std. Err.	0.0031	0.0033	0.0034	0.0035

The first three models show that engagement during the 2024 DANA follows a systematic pattern shaped by the narrative features of the content. First, days since the event are strongly and positively associated with engagement and diffusion: each additional day

yields a clear increase in interaction rates, suggesting that conversation intensifies as the crisis unfolds and competing narratives consolidate (Tables 4 and 5).

Regarding the content, all narrative categories generate positive and statistically significant increases in engagement and diffusion relative to the reference group (*practical information*). *Help and solidarity* messages as well as *political blame* conversations show a substantial effect, consistent with the strong prosocial and emotional mobilization typical of the immediate aftermath of disasters. Also, these results align with previous research documenting the structural advantage of moralized and conflictual content in algorithmic environments (Brady et al., 2017).

Mis/disinformation also produces a significant increase in engagement and amplification although its effects remain below those associated with *political blame* and *help/solidarity narratives*, exceeding *climate-related* narratives (except for the OLogit model) and the residual category. This pattern is consistent with emotional amplification mechanisms: misleading content often triggers emotions such as outrage or fear, which facilitates its spread. Explanations related to climate and non-classified content have positive but more moderate effects, reflecting a lower capacity to provoke reactions compared to morally and politically charged narratives. Overall, although the R^2 for the OLS is modest (0.021), the results reveal a robust pattern: engagement is not random, instead it evolves over time, with solidarity dominating initially, while politicized and conflict-oriented narratives increasingly capture attention as the event progresses.

4.3. The effects of emotions on engagement and diffusion

Narrative categories capture the thematic framing of tweets, but emotional tone may constitute a key underlying mechanism that shapes user engagement. Therefore, we extend the analysis by examining the effects of emotional content on engagement. The models reported in Table 4 (columns 4 to 6) and Table 5 examine how specific emotions affect engagement and dissemination during the DANA[7].

[7] Sadness is the omitted emotional category, so coefficients for anger, joy, optimism and hate are interpreted relative to sad tweets.

As in the previous specification, *days since the event* continue to strongly predict engagement in all models, showing that interactions increase as the crisis unfolds. With *sadness* as the omitted emotion, *optimism* seems to be the strongest predictor of engagement, suggesting that hopeful or forward-looking messages provoke more reactions than sadness. This may be because they signal agency and collective efficacy. Tweets which express *anger* and *hate* also receive more interactions, reinforcing evidence that both are the most mobilizing emotions in crisis communication. In fact, *anger* is the dominant emotion when explaining amplification. *Joy* shows no statistically significant differences from sadness in the OLS and Tobit models, although the OLogit estimates detect a small but positive effect. In the case of diffusion, *joy* is negative and statistically significant, showing that this kind of emotions do not generate diffusion compared with *sadness.* Overall, emotional predictors reveal a clear hierarchy: *optimism*, *hate* and *anger* generate the strongest engagement, and *joy* is largely indistinguishable from sadness.

Thus, we can see that emotional dynamics differ between engagement and diffusion. While optimism and to a lesser extent joy, contribute positively to engagement, retweet dynamics are much more strongly driven by anger and hate, suggesting that amplification mechanisms selectively favor emotionally intense and conflictual content.

4.4. Emotions and narratives effects on engagement and diffusion

Until now, our analysis has considered narrative framing and emotional content separately. However, emotional expressions are inherently embedded within narrative contexts. Thus, now we explore how emotions, narratives, and sentiment jointly shape the tone and structure of crisis-related discourses. This third set of models (Table 4 columns 7 to 9 and Table 5) considers this point. Sadness remains the omitted emotional category. Relative to emotions, *hate* and *optimism* display strong positive effects on engagement, while *anger* also increases interactions in a more moderate manner but the power of diffusion is higher. For most models, *Joy* remains statistically indistinguishable from *sadness*, confirming that joyful expressions do not meaningfully alter engagement outcomes. Nevertheless, they negatively impact on the amplification of information.

Narrative categories remain powerful predictors: *help and solidarity* and *political blame* messages exhibit the largest increases in engagement and diffusion, while *climate-*

related explanations, *mis/disinformation* and residual content (*other/not classified*) produce smaller but still significant effects. Overall, the results seem to indicate that hatred, optimism, and anger consistently outweigh sadness when it comes to generating engagement, while joy is practically neutral.

4.5. Temporal evolution of narratives

Crisis-related communication on social media evolves over time, with narrative frameworks and emotional responses evolving as events unfold and new information emerges. In this section, we examine the temporal dynamics of the major narrative categories, analyzing how their relative prominence changes over time and whether specific narratives dominate at different stages of the crisis.

Figure 3 shows the weekly number of main narratives and Figure 4 shows the same but for the "*Other/not classified"* narrative. We present two separate figures because the residual category accounts for a large and relatively stable share of tweets throughout the period, which visually dominates the scale when plotted jointly with the remaining narratives.

Figure 3 shows that *help and solidarity* narratives rise sharply in the immediate aftermath and remain prominent over time. *Political blame* increases more gradually and gains relevance in later weeks, while *practical information* peaks early and then stabilizes at lower levels. *Mis/disinformation* remains present at relatively low but persistent levels, and *climate-related* narratives consistently account for the smallest number of tweets.

Figure 3: Weekly number of main narratives during DANA

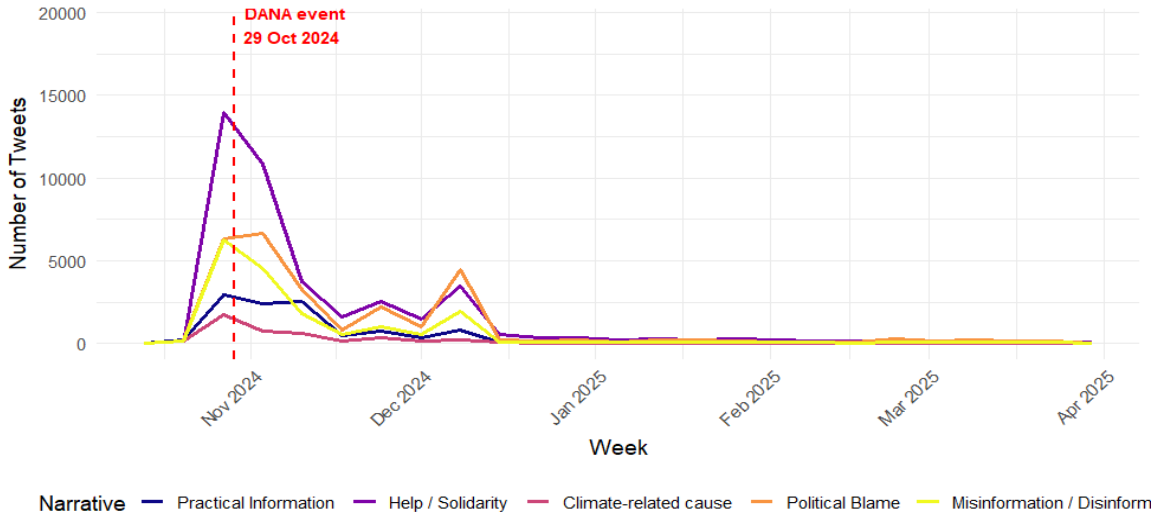
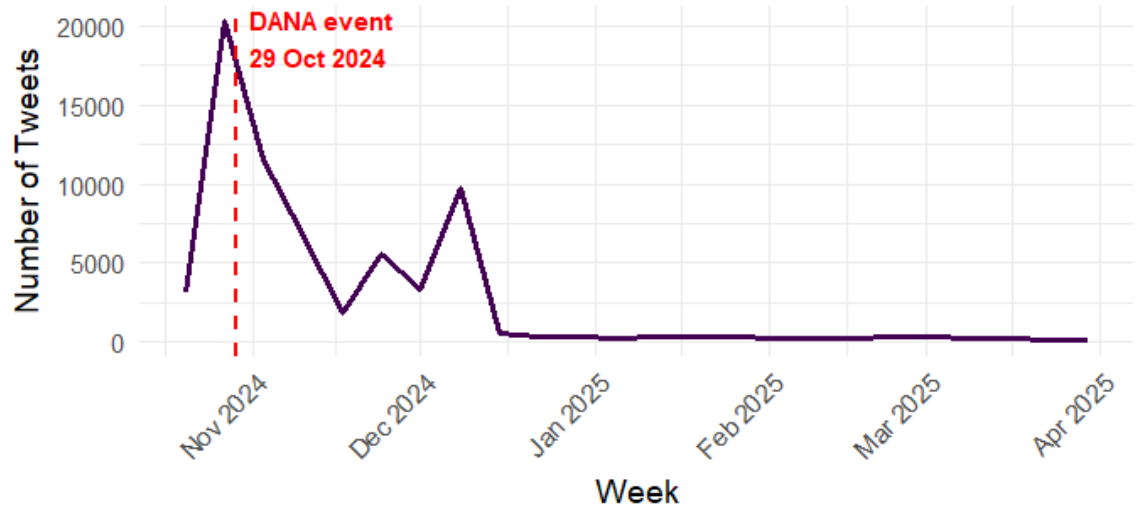


Figure 4 shows that, initially, the *other/not classified* number of tweets is very high and drops sharply after the event, before stabilizing at around 35–45% of tweets throughout the remaining period. This seems to indicate the persistence of heterogeneous and unstructured content over time.

Figure 4: Weekly number of "*other/not classified*" tweets during DANA



Figures 3 and 4 show that crisis-related discourse evolves from an initial phase dominated by diffuse and unstructured reactions toward more clearly articulated narrative frames. Immediately after the DANA event, stories of *help and solidarity* quickly emerge, while *political blame* takes center stage in later stages. *Practical information* plays a short-lived

role, and *misinformation* remains present at low but persistent levels. At the same time, the "*Other / not classified*" remains in a considerable part of the conversation.

To delve deeper into the temporal evolution, the fourth model considers time terms and interactions between time and narratives to capture different dynamic trajectories. Table 4 (columns 10 to 12) and Table 5 report the results. The results show that, in all specifications, emotional intensity stands as a key factor in engagement. *Anger*, *hate*, and *optimism* have strong positive effects on engagement and the spread of information.

The narrative framing further increases engagement, highlighting above all the basic effects of tweets related with help and solidarity, political responsibility, and climate-related issues. It is worth noting that the interaction terms show considerable temporal heterogeneity: while the engagement effect of *help and solidarity* decreases over time, *political blame* and *misinformation* become more attractive as the crisis unfolds. These results suggest a dynamic shift from prosocial narratives to more conflictive narratives after the event.

5. Discussion

5.1. From solidarity to blame and disinformation

The results show that the DANA crisis on Twitter/X was not just an informational emergency, but also a deeply emotional and political episode. First, the analysis confirms that political blame and emotionally charged narratives increasingly dominate the spread of information, although solidarity remains very active throughout the crisis. While help and support messages get high levels of engagement in the initial phase, in line with typical empathetic responses to disasters, politicized narratives and mis/disinformation increasingly concentrate attention. This finding is consistent with research that analyze the structural advantage of moralized and conflict-driven content in algorithmic environments (Brady et al., 2017; Arora et al., 2024). In the DANA case, this change is especially notable in the increase of rumors and hoaxes related to institutional responsibility (alleged failures of emergency services, dam management, governments blaming each other) and in conspiratorial narratives about climate manipulation. These findings support the qualitative analysis conducted by López-Carrión and Llorca-Abad

(2025), showing on a larger scale that misinformation and political narratives tend to generate more interaction than messages of solidarity as the crisis progresses.

5.2. Anger as a moral and emotional driver

Secondly, the models highlight the central role of anger as a spreading mechanism in crisis communication. In all specifications, anger clearly increases engagement compared to sadness, even after controlling for narratives. This pattern aligns with computational studies showing that anger is more contagious than joy and is more likely to cross community boundaries on social networks (Fan et al., 2016; Brady et al., 2017).

In DANA's data, expressions of anger are often present in *political blame* and *mis/disinformation* narratives, especially when the disaster is portrayed as preventable and attributed to political incompetence, institutional negligence, or hidden conspiracies. The positive, although minor, effect of hate indicates that openly antagonistic and dehumanizing language also contributes to maintain engagement. The fact that joy is indistinguishable from sadness suggests, however, that purely positive emotions are insufficient to generate visibility in a conflict-saturated environment.

5.3. The ambivalent power of optimism

A third remarkable finding concerns optimism. Contrary to the intuition that only negative emotions mobilize attention, our models indicate that optimism increases engagement to a degree comparable to, and sometimes greater than, anger. Hopeful messages oriented towards reconstruction or resilience generate substantial responses, likely because they provide a sense of agency in the face of devastation.

However, optimism does not replace politicization or mis/disinformation; it coexists with them. Optimistic approaches are combined with political criticism ("we will overcome this despite the incompetence of…") or are instrumentalized by disinforming actors to offer alternative "solutions" or "truths" in opposition to supposedly opaque institutions. Thus, optimism appears as an ambivalent emotion: it can sustain constructive mobilization, but it can also act as an affective glue to mix conspiracy, anti-elitism and self-empowerment.

5.4. Engagement as algorithmic governance in crises

Finally, dynamic models reinforce the idea that engagement is not merely a descriptive metric, but a mechanism of algorithmic governance. A possible interpretation is that privileging emotionally intense and conflictive content, recommendation algorithms seem to shape the trajectory of crisis narratives: prosocial and informational messages are prominent in the initial phase, but are progressively overshadowed by politicized and misleading frames that maintain a high level of engagement.

This has practical consequences for institutional communication and platform policy. If, as Méndez-Muros et al. (2024) show in other disasters, fact-checking and official messages tend to enter the information ecosystem later and with less emotional intensity, they are structurally disadvantaged in the competition for visibility. Our results suggest that something similar may have occurred during the DANA: institutional communications and debunking efforts struggled to compete with narratives driven by anger, hate, and moralized optimism.

6. Conclusions

The 2024 DANA in Spain triggered not only a large-scale physical disaster but also an informational and emotional crisis on X. Analysis of more than 168,000 tweets shows that engagement and diffusion followed partially distinct moral–emotional trajectories:

While a broad range of narratives generated interaction, amplification dynamics selectively privileged emotionally intense and conflict-driven content, particularly political blame and anger-driven communication.

1) Initial solidarity phase: in the first days, help and solidarity narratives dominate engagement, channeling grief into empathy and mutual support.

2) Politicization: over time, political blame frames gain prominence, mobilizing anger and introducing disputes over institutional responsibility and risk management.

3) Emotional disinformation: mis/disinformation narratives intensify, combining rumors, conspiracies and moral accusations. Their engagement increases as anger, hate and moralized optimism are reinforced.

4) Saturation and decline in the number of tweets: eventually, public attention wanes as the informational and emotional space saturates and the crisis loses centrality on the agenda. The predictability of this moral–emotional trajectory has a direct disaster-risk-reduction implication. Because emotional escalation—and particularly surges in outrage—appears to accompany the rise of mis/disinformation, real-time monitoring of affective intensity may serve as an early-warning signal, allowing responders to anticipate misinformation surges before they peak rather than reacting once falsehoods have already propagated (Starbird et al., 2014; Huang et al., 2015). Detection should therefore be timed to the early phase of the trajectory, when the window for effective intervention is still open.

First, institutional communication strategies in disaster contexts cannot assume that factual accuracy alone will guarantee visibility or trust. Our results indicate that official messages compete in an environment structurally biased towards moralized and conflictual content. To be effective, institutions should: a) intervene early, when solidarity and informational narratives still concentrate engagement; b)design messages that are not only accurate but also affectively resonant, acknowledging grief, anger and uncertainty without exploiting them; and c) anticipate recurrent patterns of mis/disinformation—especially those linking institutional responsibility, conspiracies and climate skepticism—and prepare rapid responses that combine factual rigour and emotional sensitivity.

From a platform governance perspective, our results support the need for crisis-sensitive algorithmic policies. These might include temporary mechanisms to dampen the diffusion of content that combines high emotional load with indicators of low veracity; priority amplification of verified coordination messages from high-trust sources; and greater transparency about how engagement signals are used to rank content during emergencies. Finally, this article offers a methodological contribution by proposing a dynamic engagement model that integrates narrative, emotional and temporal dimensions.

The current results, although valid, also have some shortcomings and limitations. First, we analyzed the social media information from a single platform (X), which, although predominant in Europe, would ideally be complemented with information from other new platforms. Second, our search criteria focused on conversations related to the DANA

episode, while there may be additional relevant data without considering this term of search. Third, social media users in general are not fully representative of the entire population, and hence, other feelings may have been present, in addition to those documented.

Future work could extend this approach by comparing different types of disasters (wildfires, heatwaves, pandemics), incorporating network structure (who interacts with whom) and combining data across platforms to provide a more comprehensive view of crisis communication. In addition, comparative analyses across countries with different levels of institutional trust and media systems could help identify the conditions under which mis/disinformation becomes dominant. Beyond the immediate response, the gradual shift from solidarity to politicized guilt can have lasting consequences for social cohesion and trust in public institutions, two essential elements of community resilience (UNDRR, 2015). Our findings, therefore, highlight the importance of strengthening informational resilience before disasters occur. Initiatives such as pre-bunking, media literacy, and trusted local communicators can help communities better manage rapidly changing information environments. In general, the DANA case suggests that misinformation during natural disasters should be seen as a moral and emotional issue that is part of larger problems related to trust and algorithm-driven systems. This implies that we need responses that combine fact-checking, emotionally intelligent communication by institutions, and regulation of platforms.

References

Arora, S., Agrawal, V., Kumar, D., Arora, S., & Banshal, S. K. (2024). Sentimental impact of fake news on social media using an integrated ensemble framework. Social Network Analysis and Mining, 14, 185.

Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. Proceedings of the National Academy of Sciences, 114(28), 7313–7318.

Charquero-Ballester, M., Walter, J. G., Nissen, I. A., & Bechmann, A. (2021). Different types of COVID-19 misinformation have different emotional valence on Twitter. Big Data & Society, 8(2).

Chen, J., Xiao, S., Zhang, P., Luo, K., Lian, D., & Liu, Z. (2024). BGE M3-Embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. arXiv preprint arXiv:2402.03216.

Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C. M., Brugnoli, E., Schmidt, A. L., Zola, P., Zollo, F., & Scala, A. (2020). The COVID-19 social media infodemic. Scientific Reports, 10, 16598. https://doi.org/10.1038/s41598-020-73510-5

Fan, R., Xu, K., & Zhao, J. (2016). Higher contagion and weaker ties mean anger spreads faster than joy in social media. arXiv preprint arXiv:1608.03656.

Government of Spain. (2025). [Report on the impacts of the 2024 DANA in the Valencian Community].

Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking Sandy: Characterizing and identifying fake images on Twitter during Hurricane Sandy. In Proceedings of the 22nd International Conference on World Wide Web (WWW '13 Companion) (pp. 729–736). ACM. https://doi.org/10.1145/2487788.2488033

Huang, Y. L., Starbird, K., Orand, M., Stanek, S. A., & Pedersen, H. T. (2015). Connected through crisis: Emotional proximity and the spread of misinformation online. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (pp. 969–980). ACM. https://doi.org/10.1145/2675133.2675202

Hutto, C. J., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. In Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media (pp. 216–225). AAAI Press.

Ivie. (2025). [Economic assessment of the impact of the 2024 DANA in the Valencian Community].

Klein, N. (2023). Doppelganger: A trip into the mirror world. Farrar, Straus and Giroux.

Kryvasheyeu, Y., Chen, H., Obradovich, N., Moro, E., Van Hentenryck, P., Fowler, J., & Cebrian, M. (2016). Rapid assessment of disaster damage using social media activity. Science Advances, 2(3), e1500779. https://doi.org/10.1126/sciadv.1500779

López-Carrión, A. E., & Llorca-Abad, G. (2025). Desinformación durante la crisis producida por la DANA de 2024 en España: Análisis, características, tipologías y desmentidos. Revista Mediterránea de Comunicación, 16(2), e29303.

Martel, C., Pennycook, G., & Rand, D. G. (2020). Reliance on emotion promotes belief in fake news. Cognitive Research: Principles and Implications, 5(1), 47.

Mendoza, M., Poblete, B., & Castillo, C. (2010). Twitter under crisis: Can we trust what we RT? In Proceedings of the First Workshop on Social Media Analytics (pp. 71–79). ACM. https://doi.org/10.1145/1964858.1964869

Méndez-Muros, S., Alonso-González, M., & Pérez-Curiel, C. (2024). Disinformation and fact-checking in the face of natural disasters: A case study on Turkey–Syria earthquakes. Societies, 14(4), 43.

Ogie, R. I., James, S., Moore, A., Dilworth, T., Amirghasemi, M., & Whittaker, J. (2022). Social media use in disaster recovery: A systematic literature review. International Journal of Disaster Risk Reduction, 70, 102783. https://doi.org/10.1016/j.ijdrr.2022.102783

Oh, O., Agrawal, M., & Rao, H. R. (2013). Community intelligence and social media services: A rumor theoretic analysis of tweets during social crises. MIS Quarterly, 37(2), 407–426. https://doi.org/10.25300/MISQ/2013/37.2.05

Palen, L., & Anderson, K. M. (2016). Crisis informatics—New data for extraordinary times: Focus on behaviors, not on fetishizing social media tools. Science, 353(6296), 224–225. https://doi.org/10.1126/science.aag2579

Palen, L., Hiltz, S. R., & Liu, S. B. (2007). Online forums supporting grassroots participation in emergency preparedness and response. Communications of the ACM, 50(3), 54–58. https://doi.org/10.1145/1226736.1226766

Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. Nature, 592(7855), 590–595. https://doi.org/10.1038/s41586-021-03344-2

Pröllochs, N., et al. (2021). Emotions explain differences in the diffusion of true and false social media rumours. Scientific Reports, 11, 23456.

Reynolds, B., & Seeger, M. W. (2005). Crisis and emergency risk communication as an integrative model. Journal of Health Communication, 10(1), 43–55. https://doi.org/10.1080/10810730590904571

Simon, T., Goldberg, A., & Adini, B. (2015). Socializing in emergencies—A review of the use of social media in emergency situations. International Journal of Information Management, 35(5), 609–619. https://doi.org/10.1016/j.ijinfomgt.2015.07.001

Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, R. M. (2014). Rumors, false flags, and digital vigilantes: Misinformation on Twitter after the 2013 Boston Marathon bombing. In iConference 2014 Proceedings (pp. 654–662). iSchools. https://doi.org/10.9776/14308

Tomblinson, C. M., Latimer, E., Glynn, C. H., McCarty, J. L., & Fitzgerald, R. T. (2019). Publicly available metrics underestimate AJNR Twitter impact and follower engagement. American Journal of Neuroradiology, 40(12), 1994–1997. https://doi.org/10.3174/ajnr.A6299

United Nations Office for Disaster Risk Reduction. (2015). Sendai Framework for Disaster Risk Reduction 2015–2030. UNDRR. https://www.undrr.org/publication/sendai-framework-disaster-risk-reduction-2015-2030

Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging during two natural hazards events: What Twitter may contribute to situational awareness. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1079–1088). ACM. https://doi.org/10.1145/1753326.1753486

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. Science, 359(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Wardle, C., & Derakhshan, H. (2017). Information disorder: Toward an interdisciplinary framework for research and policymaking. Council of Europe.

World Health Organization. (2021). WHO public health research agenda for managing infodemics. World Health Organization.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- DANAscript.r
- READMEII.txt
- CopiadeDANAsentimetiquetasautorX.csv