Safety Science 202 (2026) 107281

Contents lists available at ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/safety

Sensing digital risk discourse on algorithmically curated platforms: A case study of Typhoon Gaemi via LLM and explainable spatial analysis☆

Pu ZHANG

The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

ARTICLE INFO

Keywords:
Disaster Informatics
Digital risk discourse
Douyin (TikTok)
Large language models (LLM)
SHAP-GIS
Protective action decision model (PADM)
Public safety

ABSTRACT

Effective disaster management requires real-time sensing of public response, yet how risk discourse unfolds on algorithmically curated short-video platforms remains poorly understood. This case study analyzes 436,964 user comments on Douyin, China's leading short-video platform, distinguished from follower-based platforms by its algorithm-driven content recommendation, posted during Typhoon Gaemi (2024). Grounded in the Protective Action Decision Model (PADM) and the Social Amplification of Risk Framework (SARF), we classified discourse into ten thematic categories using a validated Large Language Model (DeepSeek R1 671B; 88.3% accuracy). Two key findings emerge. First, socio-cultural coping, including humor, regional identity narratives, and mutual aid (32.98%), achieved near-parity with functional information exchange (33.89%), challenging the information-centric paradigm in disaster communication. Second, spatial regression models identify robust associations between regional socio-economic characteristics and discourse patterns. Cropland Cover Proportion emerged as the most pervasive correlate, predicting both intensified situational reporting ($\beta = 0.764$, $q = 0.086$) and heightened scrutiny of information reliability ($\beta = 0.871$, $q = 0.009$); Natural Disaster Economic Loss per capita was independently associated with information trust discourse ($\beta = 0.636$, $q = 0.086$). All associations remained robust after correction for spatial autocorrelation and multiple testing. A novel SHAP-GIS workflow further decomposes these associations at the provincial level, revealing how macro-level socio-economic contexts contribute to local patterns of public response. This study demonstrates that online risk discourse is systematically associated with offline socio-economic realities, offering a scalable methodology for spatially targeted risk communication during climate-related emergencies.

1. Introduction

The 21st century is increasingly defined by the escalating frequency and intensity of climate-related extreme events, such as heatwaves, floods, and powerful tropical cyclones (Lange et al., 2020; Zhang et al., 2024c). These disasters pose profound threats to global health, economic stability, and social well-being (Thiery et al., 2021). In this landscape, understanding how the public perceives and responds to such hazards is a critical component of effective disaster management and societal resilience (Ma et al., 2024; Renn, 1998).

The rise of social media platforms has shifted public discourse from a top-down, one-way broadcast model to a real-time, many-to-many, interactive ecosystem (Alexander, 2013; Houston et al., 2014; Zhang et al., 2026c). Platforms like *Twitter*, *Facebook*, *Weibo*, and especially the short-video platform *Douyin* (抖音) in China, have become central arenas for public sense-making (Oh et al., 2015; Zhang et al., 2026d P. Zhang et al., 2024b). Unlike text-based platforms, short-video applications like *Douyin* offer immersive, visually rich, and algorithmically-driven experiences, potentially reshaping how risk is perceived and emotionally processed by the public (Kaye et al., 2021; Zhang et al., 2026a). During a disaster, these platforms are flooded with official warnings, citizen-generated ground-truth observations, expressions of emotion, and requests for help (Fauzi, 2023). This study analyzes this discourse through the integrated lenses of Risk Perception Theory (RPT), which explains the subjective, individual-level judgments of risk (Siegrist & Árvai, 2020; Slovic, 1987), and the Social Amplification of Risk Framework (SARF), which explains how these perceptions are socially transmitted, amplified, or attenuated (Kasperson et al., 1988; Wirz et al., 2018).

Despite the growing body of research on social media in disasters, significant gaps persist. First, much of the existing literature has focused on Western, text-based platforms like *Twitter* (Noori et al., 2025; Zou

☆ This article is part of a special issue entitled: 'Crisis and Emergency Management' published in Safety Science.

E-mail address: pzhang012@connect.hkust-gz.edu.cn.

https://doi.org/10.1016/j.ssci.2026.107281
Received 29 November 2025; Received in revised form 2 May 2026; Accepted 6 May 2026

et al., 2023). While research on *Douyin* is emerging, with recent systematic reviews beginning to map its role in disaster reporting and public engagement (Ilyas & Sharifi, 2025; Wahid et al., 2025), these studies often rely on descriptive content analysis of video metadata or relatively small-scale qualitative samples. Less is understood about the unique discourse on visually driven, algorithmically curated, and culturally specific platforms like *Douyin*, which has a massive user base and distinct communication norms (Dong & Xie, 2024; Kaye et al., 2021). Our work extends these pioneering efforts by capturing the large-scale digital risk discourse of 436,964 valid comments, providing a much more granular view of public sense-making.

Second, traditional computational methods (e.g., LDA or keyword-matching) often struggle to capture the full spectrum of nuanced, complex, and informal risk discourse (e.g., humor, regional metaphors, or risk dismissal) prevalent on visually driven platforms (Dong et al., 2022; P. Zhang et al., 2024a). We address this by employing a validated LLM-few-shot classification workflow (DeepSeek R1 671B), which allows for a more "human-like" understanding of socio-cultural nuances.

Third, a critical gap exists in connecting these "micro-level" online expressions of risk to the "macro-level" societal contexts from which they emerge. Few studies have systematically investigated how regional socio-economic characteristics correlate with the collective patterns of online risk discourse (Balog-Way et al., 2020; Cutter et al., 2003; Reid, 2013). While previous research has explored social media during disasters (P. Zhang et al., 2024b), no prior study has combined validated LLM thematic classification at this large scale with spatial econometric modeling and SHAP-GIS on a short-video platform. By situating our findings within the Protective Action Decision Model (PADM), this study provides an explainable spatial framework for sensing risk discourse in the TikTok/Douyin era.

This study addresses these thematic, methodological, and analytical gaps by conducting a multi-faceted investigation of public risk discourse on *Douyin* during Typhoon Gaemi—a recent event that generated over 740,000 raw comments. Our research is guided by two core questions:

RQ1: What was the aggregate thematic landscape of disaster-related risk discourse on Douyin during the cycle of Typhoon Gaemi, and how was this discourse spatially patterned across different provinces?

RQ2: What spatial associations exist between provincial-level socio-economic indicators and the proportional prevalence of these specific risk discourse categories, and how do these correlates characterize the regional disaster response?

To answer these questions, we employ a mixed-methods research design. Critically, to overcome the methodological challenges of analyzing this massive, informal dataset, we first develop and validate a robust classification workflow. We use an advanced Large Language Model (*DeepSeek R1 671B*) with a few-shot learning strategy, achieving an 88.3% accuracy against manual coding. This validated, high-accuracy classification allows us to confidently analyze the prevalence and spatial distribution of ten thematic categories using GIS. Finally, we apply Multiple Beta Regression, spatial autocorrelation tests, and SHAP analysis to model and interpret the relationship between topic prevalence and nine provincial-level socio-economic indicators.

This research makes several key contributions to the fields of disaster informatics and risk communication:

Empirical Contribution: It provides a rare large-scale (N = 436,964) characterization of public risk discourse on an algorithmically curated short-video platform. By capturing the collective "sensing" of Typhoon Gaemi on Douyin, this study moves beyond the video-metadata and creator-content analyses that dominate existing literature to reveal the granular, bottom-up sense-making patterns of the public.

Methodological Innovation: We demonstrate and validate a transparent, LLM-augmented SHAP-GIS workflow. By leveraging the few-shot reasoning capabilities of DeepSeek R1 671B, we achieve high-fidelity classification (88.3% accuracy) of nuanced social data. Furthermore, the integration of SHAP with spatial econometrics provides a scalable, open-source pipeline for interpreting how macro-level societal factors contribute to local digital response—a significant advancement in explainable AI (XAI) for social sensing.

Theoretical Advancement: This study bridges the Protective Action Decision Model (PADM) and the Social Amplification of Risk Framework (SARF). We provide empirical evidence for the "Information-Coping Equilibrium," showing that socio-cultural sense-making (e.g., humor, cultural identity) is as prevalent (32.98%) as functional information exchange (33.89%). This challenges the information-centric paradigm in risk theory and conceptualizes Douyin as a critical "amplification station" where interpretive signals are as vital as technical warnings.

Practical Implications: By identifying robust spatial associations between regional socio-economic indicators (e.g., agricultural exposure, historical disaster losses) and specific discourse themes, our findings offer actionable insights for emergency managers. This allows for targeted risk communication strategies that align with the specific vulnerabilities and discourse characteristics of different provincial populations, ultimately fostering societal resilience through data-driven situational awareness.

2. Conceptual framework

The conceptual framework of this study (see Fig. 1) transitions from a static data pipeline to an integrated risk sensing and response system. This system provides a process-oriented lens to understand how individual psychological appraisals are transformed into collective digital discourse through social and environmental mediation.

This framework draws on three complementary theoretical perspectives, each serving a distinct analytical function. Risk Perception Theory (RPT) provides the foundational understanding that individuals' judgments of threat severity and personal vulnerability are shaped not only by objective hazard characteristics but also by cognitive heuristics, affective responses, and socio-cultural contexts (Jenkins et al., 2024; Yin & Lui, 2024). RPT informs our identification of discourse categories that reflect these subjective appraisals, such as Worry/Fear and Risk Attenuation. The Social Amplification of Risk Framework (SARF) extends this individual-level lens to the societal level, theorizing that risk signals are filtered, amplified, or attenuated as they pass through social and institutional channels (Lechowska, 2021; Rohrmann & Renn, 2000)—a process that is particularly relevant to algorithmically curated platforms like Douyin, where content visibility is shaped by both user engagement and platform algorithms. SARF informs our analysis of how discourse patterns vary spatially and how digital platforms function as amplification stations. Finally, the Protective Action Decision Model (PADM) provides a process-oriented sequential structure that links these two perspectives, as detailed below.

At the core of the framework is the Protective Action Decision Model (PADM) (Lindell & Perry, 2011), which delineates the internal progression from information exposure and social observation, through risk and stakeholder appraisal, to action intentions. Unlike traditional models that treat perception as an isolated state (Cacioppo & Hawkley, 2009), our framework recognizes that these stages are socially mediated by the socio-economic context (e.g., agricultural exposure and historical disaster losses), which functions as a set of pre-existing predispositions associated with how individuals process environmental signals from Typhoon Gaemi.

We conceptualize social media discourse not merely as raw data, but as a dynamic observation window into the public's psychological stages. We explicitly distinguish between observable and unobserved components within this digital corpus: **Observable Digital Signals:** The corpus captures real-time information exchange, collective sense-making (e.g., humor and cultural narratives), and observational learning. These represent the communicative output of the cognitive appraisal and coping stages. **Unobserved Physical Actions:** We acknowledge that actual preparedness and physical protective behaviors (e.g., evacuation or property reinforcement) remain unobserved in this textual dataset. By situating our categories within the PADM, we

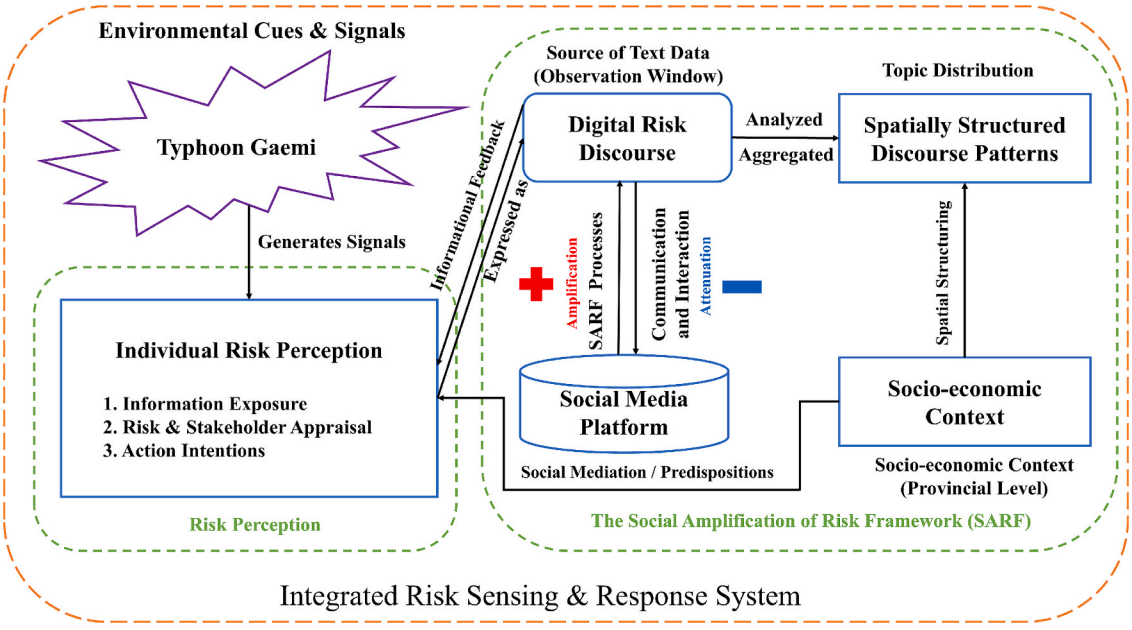

Fig. 1. The conceptual framework for the study.

interpret digital discourse as a proxy for psychological and social efficacy rather than a direct measure of behavioral outcomes.

The framework incorporates two critical systemic links: A direct analytical link exists between the socio-economic context and the resulting spatially structured discourse patterns. This represents the core of our spatial association analysis, where we examine how offline regional characteristics systematically correlate with the prevalence of specific thematic categories (Zhang et al., 2026e). A feedback loop is established where the prevailing digital risk discourse feeds back into individual risk perception. This reflects the process of social observation, where the collective wit, narratives, or reports encountered on platforms like Douyin continuously recalibrate individual appraisals and future responses.

The feedback loop described above draws on insights from SCT, particularly the concepts of observational learning and social efficacy (Ma et al., 2024). In the context of disaster-related social media discourse, users do not merely broadcast their own appraisals; they also observe and internalize the collective discourse of others—whether it amplifies fear, normalizes risk, or models coping behaviors. While the PADM provides the sequential decision-making structure for our

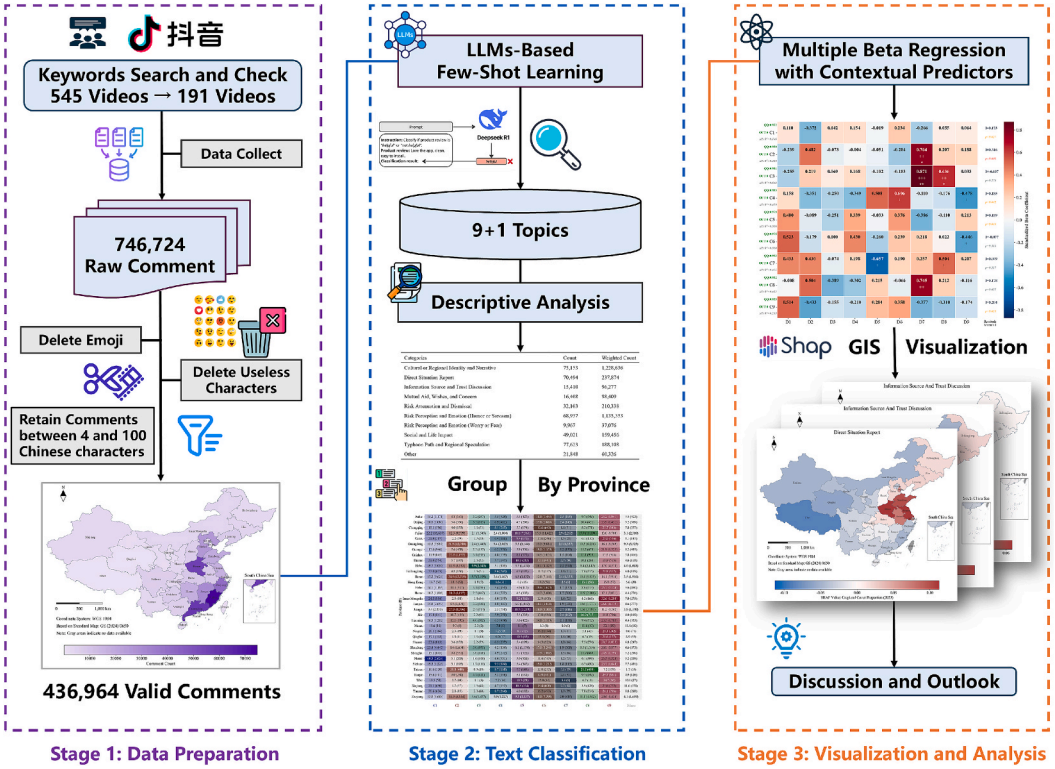

Fig. 2. The mixed methods research workflow.

analysis, SCT enriches our understanding of the social learning processes through which individual appraisals are continuously recalibrated by exposure to the digital discourse environment.

3. Research methodology

This chapter details the comprehensive, mixed-methods research design employed to address the research questions. As visually summarized in the research workflow (see Fig. 2), the methodology is structured in three primary stages: **Stage 1: Data Preparation:** This stage involved the collection, systematic preprocessing, and filtering of a large-scale comment dataset from the social media platform *Douyin* concerning Typhoon Gaemi, resulting in a final valid corpus for analysis. **Stage 2: Text Classification:** This computational stage utilized an advanced Large Language Model (LLM) with a few-shot learning approach to thematically categorize the entire comment corpus. This was followed by a descriptive analysis of the topic distribution and aggregation of the data by province. **Stage 3: Visualization and Analysis:** This final stage employed a suite of analytical techniques to interpret the classified data. It included GIS visualization for spatial patterns, Multiple Beta Regression to model the relationships between topic prevalence and socio-economic correlates, and SHAP analysis for in-depth model interpretation.

3.1. Data preparation

The foundation of this study rests on two distinct types of data. The first is the primary dataset of public discourse, collected from the *Douyin* social media platform, which captures real-time public risk discourse. The second is a set of provincial-level socio-economic indicators, which serve as the independent variables for the regression analyses. The preparation of each dataset is described below.

3.1.1. Social media data preparation

We collected public discourse data from Douyin (抖音), China's most popular social media platform and the leading short-video app. As the Chinese counterpart to TikTok, Douyin represents a transformative shift in social media architecture. Unlike traditional follower-based platforms such as Twitter or Weibo, Douyin employs a content-centric, algorithm-driven recommendation system that prioritizes engagement over social connections. Its immense popularity and unique short-video format make it particularly well-suited for this study, as it not only reflects authentic, real-time public reactions to events like Typhoon Gaemi but also effectively stimulates risk-related discussions among users. The primary unit of analysis is the user comments generated in response to disaster-related videos. We view the videos as the 'source signals' or 'visual stimuli' provided by the platform's algorithm, whereas the comments serve as the digital traces of public risk sensing and meaning making. By analyzing these comments, we aim to capture the reception and amplification of risk signals among the public, rather than the content production strategies of video creators. The data collection and preparation process, as illustrated in Stage 1 of the research workflow, followed several rigorous steps.

First, on October 30, 2024, a keyword search (i.e., "Typhoon Gaemi," "台风格美") was conducted on the Douyin platform to identify relevant videos. This initial query yielded 545 video. These video were then manually reviewed for their direct relevance to the disaster event (e.g., news reports, eyewitness accounts, or discussions about the typhoon). Videos that were irrelevant or purely entertainment-focused were discarded, resulting in a curated set of **191 relevant** videos for data collection.

Second, a systematic data collection process was implemented to extract all available comments from these 191 videos. For each video, the comment section was manually accessed and fully loaded by scrolling to trigger Douyin's dynamic content loading until no additional comments appeared. All visible comments (including comment text, timestamps, and engagement metrics) were exported using a browser-based data extraction tool. This approach accesses only publicly displayed content rendered in the user's browser and does not bypass any access controls, authentication mechanisms, or API rate limits, thereby maintaining compliance with the platform's Terms of Service. The process yielded a raw corpus of 746,724 comments collected.

Third, this raw data underwent a systematic preprocessing and cleaning pipeline to ensure data veracity and mitigate the influence of non-substantive or automated content. This pipeline included: **Deletion of Emojis and Useless Characters:** All emojis, special symbols, ASCII art, and excessive punctuation were removed to isolate the core textual intent. **Rigorous Anonymization:** To ensure user privacy and remove metadata noise, a pipeline stripped all personal identifiers, including nicknames and "@username" mentions. This also effectively filtered out tagging-heavy promotional content or bot-like social networking spam. **Length-Based Filtering (Bot Detection Proxy):** To ensure each comment contained a minimum level of analysable content while excluding overly long or off-topic posts (e.g., copypasta, advertisements), the dataset was filtered to retain only comments between **4 and 100 Chinese characters** in length. This threshold serves as a robust proxy for bot mitigation, as it prunes the high-frequency "one-word" echoes and repetitive automated scripts common in large-scale social media streams. **Geographic Contextualization:** As provincial-level IP locations are a mandatory public feature of the Douyin platform, their use in this study for aggregate-level spatial analysis complies with the platform's Terms of Service and ethical standards for public-domain research. This rigorous preparation process resulted in a final, valid dataset of 436,964 comments, representing a high-fidelity observation window into public risk discourse.

Fig. 3 presents a choropleth map illustrating the geographical distribution of this final valid comment corpus. The map shows that the volume of discussion was not evenly distributed; it was heavily concentrated in the coastal provinces directly in the typhoon's path and subsequent impact zones (such as *Fujian*, *Guangdong*, *Zhejiang*, and *Jiangxi*), as well as in populous inland provinces that were on high alert for heavy rainfall (such as *Henan*). This cleaned and geolocated corpus was then used for all subsequent thematic classification and statistical analysis.

3.1.2. Socioeconomic development data preparation

To investigate the spatial associations between macro-level factors and regional patterns of risk discourse (as outlined in RQ2), a secondary dataset of provincial-level socio-economic indicators was compiled as independent variables for subsequent regression analysis. The selection

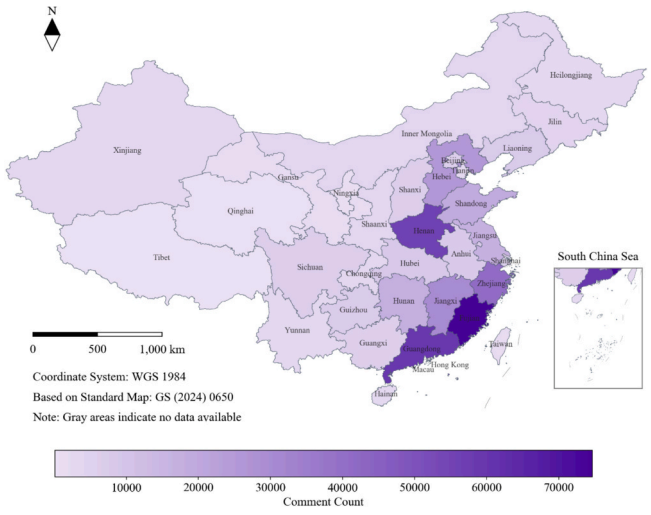

Fig. 3. Geographical distribution of the valid comment corpus (N = 436,964) by Province.

of indicators followed a theoretically informed, multi-dimensional framework commonly employed in disaster research, while being adapted to the specific context of online risk discourse in China.

The nine indicators were selected based on three key criteria: (1) theoretical relevance to disaster vulnerability and risk communication literature (AlQahtany & Abubakar, 2020; Askari et al., 2025; Li et al., 2023; Mallick et al., 2022), (2) data availability at the provincial level, and (3) representation of distinct dimensions of regional characteristics. These indicators represent both standard metrics in disaster studies (e.g., economic loss, affected area) and researcher-selected variables informed by social media behavior research (e.g., internet penetration, educational attainment).

The indicators encompass four analytical dimensions: **Economic Development: GDP per capita (CNY, 2023)** serves as a standard proxy for regional economic capacity and resource availability for disaster preparedness and response. **Demographic Structure: Gender Ratio (2020) and Population Growth Rate (2014–2023)** capture demographic composition and dynamics that may influence vulnerability patterns and information-seeking behaviors. **Human Capital & Information Access: Average Schooling Years (2020) and Internet Penetration Rate (2023)** measure education levels and digital connectivity, both critical factors shaping online discourse participation and risk comprehension. **Agricultural Exposure: Primary Industry GDP Proportion (2023) and Cropland Cover Proportion (2023)** indicate agricultural dependency and rural character, reflecting sector-specific vulnerabilities to natural hazards (Yang, 2024). **Disaster Impact Metrics: Natural Disaster Economic Loss per capita (10 K CNY, 2023) and Natural Disaster Affected Area per capita (Hectares, 2023)** represent direct disaster experience, following standard disaster statistics frameworks used by national authorities.

All data were collected from official publications, primarily the National Bureau of Statistics of China, ensuring reliability and consistency across provinces. Fig. 4 presents choropleth maps visualizing the spatial distribution of these indicators, illustrating baseline geographic variability in socio-economic and vulnerability contexts. This provincial-level dataset was then matched to the aggregated topic-proportion data derived from the Douyin comment corpus for spatial analysis.

3.2. Text classification

The second stage of this research involved the systematic classification of the 436,964 prepared comments into meaningful thematic categories. This computational phase was essential for transforming the unstructured text data into a structured dataset suitable for quantitative analysis. The process involved defining the analytical categories, classifying the data using a Large Language Model (LLM), and validating the classification performance.

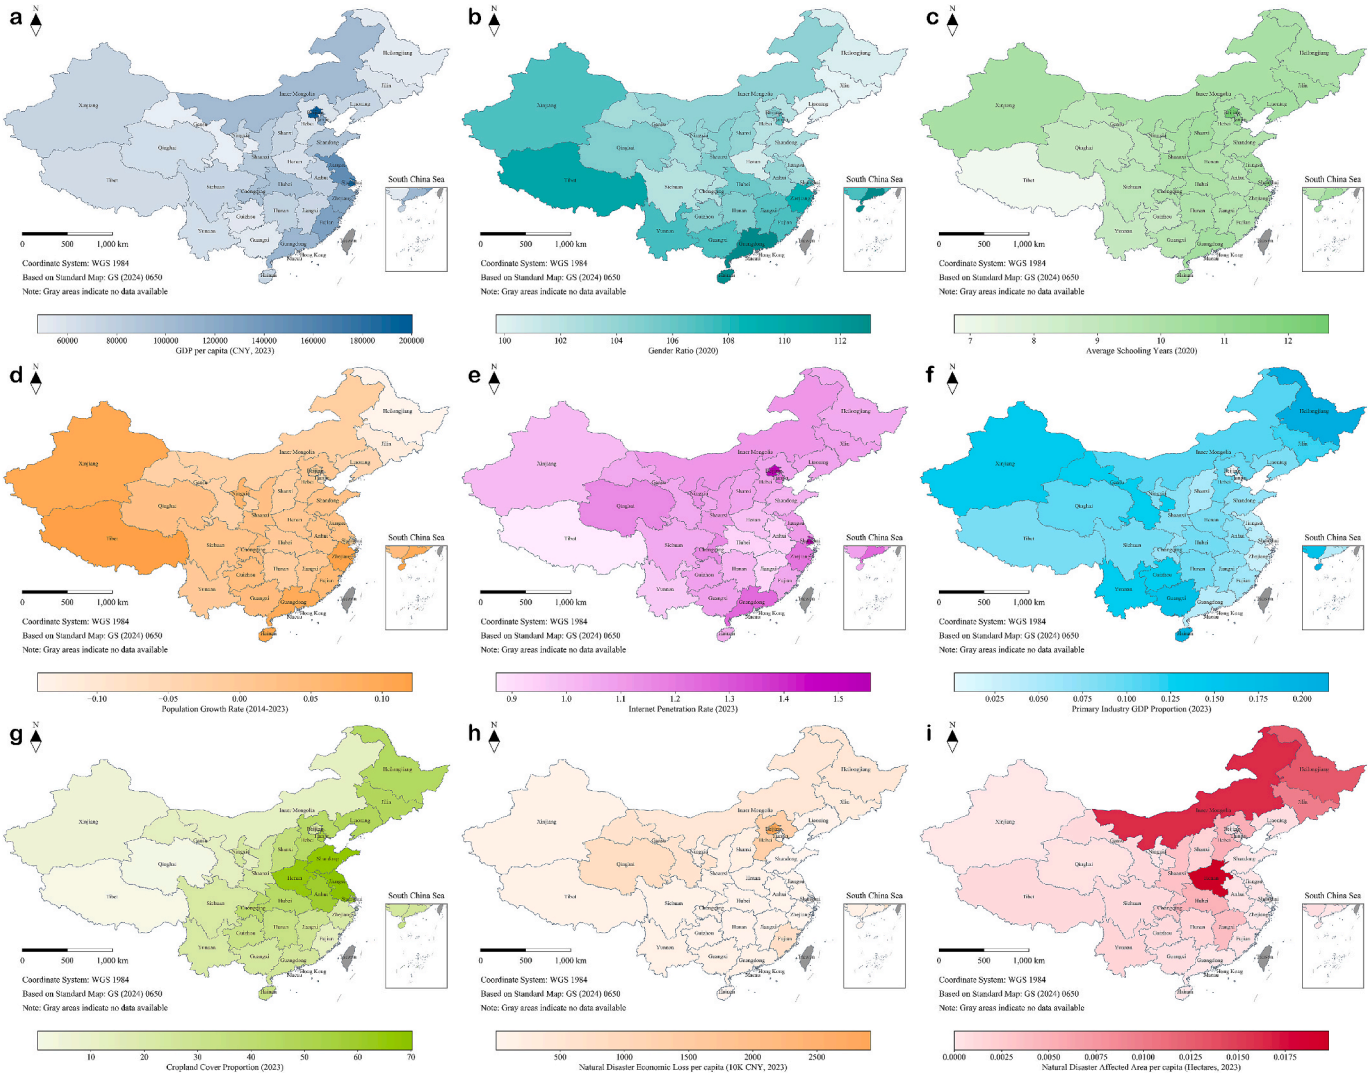

Fig. 4. Geographical distribution of the nine provincial-level socio-economic indicators.

3.2.1. Topic identification and definition

The development of the thematic classification scheme (the "codebook") followed a critical, iterative, and hybrid approach, combining top-down theoretical deduction with bottom-up empirical induction (Chen et al., 2018; Minaee et al., 2021). This process was rooted in the lead author's foundational qualitative work, involving an extensive and prolonged immersion in Douyin's disaster-related discourse.

First, a set of priori themes was deduced from established frameworks, primarily Risk Perception Theory (RPT) and the Social Amplification of Risk Framework (SARF), to capture core constructs such as emotional responses and information-seeking. Second, to ensure the codebook reflected the unique, algorithmically curated, and culturally specific environment of Douyin, a systematic, three-stage inductive open coding process was performed:

1) **Pilot Identification (n = 1,000):** An initial random sample of 1,000 comments was manually coded to identify preliminary thematic clusters and establish a baseline coding framework.
2) **Incremental Expansion (n = 1,000 + 2,000 = 3,000):** An additional 2,000 randomly sampled comments were introduced, bringing the cumulative total to 3,000. This stage revealed several new platform-specific thematic categories not captured in the pilot round, including localized memes, regional identity narratives, and diverse coping styles—demonstrating that the initial sample alone was insufficient to represent the full discursive landscape.
3) **Saturation Verification (n = 3,000 + 2,000 = 5,000):** A further 2,000 comments were added, bringing the final cumulative sample to 5,000. At this stage, no new thematic categories with a prevalence exceeding 0.5% emerged in the additional batches. The absence of novel themes confirmed that theoretical saturation had been reached, suggesting that the coding scheme successfully captured the full spectrum of dominant public discourse during the typhoon event.

The transition from open codes to the final ten-category scheme was governed by a principle of functional-theoretical parsimony: the codebook was refined to the level of granularity at which each retained category occupied a unique functional position within the PADM framework while maintaining sufficient empirical prevalence (>0.5%) to support meaningful spatial analysis. Approximately 15 initial open codes identified during Stages 1–2 were reduced to ten through theory-guided merging (e.g., "Humor" and "Sarcasm" consolidated under shared affective coping function) and retention of theoretically distinct pairs (e.g., "Direct Situation Reports" vs. "Path Speculation" preserved to capture different information-processing stages). The residual "Other" category absorbed all codes falling below the 0.5% prevalence threshold. We therefore characterize the final codebook as an analytically purposeful, judgment-based taxonomy whose structure is dictated by theoretical alignment and analytical utility rather than statistical optimization (Fereday & Muir-Cochrane, 2006).

Finally, categories derived from both phases were synthesized and refined using a functional-theoretical logic aligned with the Protective Action Decision Model (PADM) (Heath et al., 2017; Lindell & Perry, 2011). This step was crucial for ensuring that the categories were not merely descriptive but corresponded to specific social-cognitive stages of disaster response. For instance, "Risk Perception and Emotion (Humor or Sarcasm)" and "Cultural or Regional Identity and Narrative" were integrated as facets of Socio-cultural Coping, while "Direct Situation Reports" and "Typhoon Path and Regional Speculation" were maintained as distinct categories to differentiate between environmental cue observation and deliberate information seeking (Mukumbang, 2023). By utilizing the PADM as a taxonomical lens, we operationalize risk discourse as the observable communicative output of the public's internal cognitive appraisals. This approach bridges the gap between static content classification and the dynamic processes of disaster response (see Table 1).

3.2.2. LLM-based few-shot classification

To classify the entire corpus of 436,964 comments, this study employed a few-shot learning approach based on a Large Language Model (LLM). This method was chosen over traditional machine learning or dictionary-based methods due to its superior ability to understand the context, nuance, and informal language (e.g., slang, irony) prevalent in social media comments (Song et al., 2023; Wang et al.,

Table 1
Theoretical mapping of disaster risk discourse to the protective action decision model.

PADM stage	Discourse category	Theoretical mechanism and psychological impact	Description and example observations
I. Pre-decisional Information Processes	Typhoon path and regional speculation	Information Exposure and Attention. Publics use environmental cues and media to determine the risk's location.	Comments on predicted trajectory, speed, and landfall location, including inquiries about specific affected areas.
	Direct situation report	Information Exposure and Attention. Initial sensory awareness and localization of the threat.	First-person observations of current weather conditions, such as wind intensity, rain, or clear skies in a specific spot.
II. Risk and Threat Appraisal	Risk perception: worry and fear	Threat Perception. Individuals evaluate the severity of the disaster and their personal vulnerability.	Direct expressions of anxiety, fear, or dread regarding the typhoon's potential impact or current strength.
	Risk attenuation and dismissal	Threat Perception. Individuals downplay the threat or normalize the danger based on experience.	Comments expressing a lack of concern or normalizing the threat, such as "It's just a bit of rain" or "We are used to this."
	Social and life impact	Threat Perception. Individuals evaluate the tangible consequences of the disaster on their environment.	Discussions on how the typhoon affects daily life, including work, school, travel plans, infrastructure, and the economy.
III. Information and Stakeholder Evaluation	Information source and trust discussion	Information and Source Appraisal. A critical evaluation of the reliability of official and media communications.	Comments that evaluate, question, or express trust and distrust in weather forecasts, media reports, and official sources.
IV. Socio-cultural Coping	Risk perception: humor and sarcasm	Affective Response, Collective Sense-making, and Social Support. A coping mechanism is used to alleviate cognitive dissonance and collective fear.	Discussing the typhoon in a humorous, teasing, or sarcastic manner to manage stress through collective wit.
	Cultural or regional identity and narrative	Affective Response and Social Support. Using identity and history to process the event and strengthen resilience.	Linking the typhoon to local culture and regional history, such as references to provincial resilience or local myths.
V. Action Intentions and Efficacy	Mutual aid, wishes, and concern	Protective Action and Efficacy. Enhancing collective coping through pro-social behavior and support.	Expressions of empathy and prayers for safety, including well-wishes for affected areas and reminders for others to stay safe.
VI. Residual Category	Other	Data Exhaustiveness. Capturing noise data to ensure the coding scheme is complete.	Comments that are irrelevant, unclassifiable, or too generic to fit into a specific thematic category.

2020; Zhang et al., 2025). The specific model utilized was the ***DeepSeek R1 671B***,[1] chosen for its strong performance in complex language understanding and few-shot classification tasks (Guo et al., 2025). Instead of training a model on thousands of labeled examples, a carefully engineered prompt was provided to the LLM for each comment (Gilardi et al., 2023). This prompt included: (1) a system role definition, (2) the list of all 10 categories and their detailed definitions (from Table 1), (3) a curated set of representative examples (2–3 per category) illustrating the classification, and (4) the specific user comment to be classified. This few-shot approach allows for the efficient and scalable application of a nuanced, theoretically informed classification scheme across a massive dataset, which would be infeasible to achieve through manual coding. The core Python code and detailed prompt can be found in **Supplementary Materials**.

To assess the reliability and accuracy of the LLM-based classification, a rigorous validation was performed. A stratified random sample of 1,000 comments was extracted for validation. This sample size follows standard practices in text classification research (Artstein & Poesio, 2008; Hripcsak, 2005) and provides stable estimates of classification performance. These 1,000 comments were then independently and manually labeled by a human annotator (following the same 10-category scheme) to establish a "ground truth" dataset. The LLM's classifications for this same 1,000-comment subset were then compared against the human-annotated labels. The performance of the classifier is detailed in Table 2.

The few-shot classification workflow achieved a high overall accuracy of 88.3% on the 1,000-comment validation set. This robust performance, particularly in capturing nuanced discourse such as 90.34% for Risk Perception and Emotion (Humor or Sarcasm) and Risk Attenuation and Dismissal (77.63%), provides strong confidence in the reliability of the LLM-generated data for the full corpus. The comparatively lower accuracy for Cultural or Regional Identity and Narrative (73.68%) reflects the category's inherent subjectivity and its reliance on deeply localized cultural inferences, which are further addressed in the discussion.

Table 2
LLM classification validation performance (N = 1,000 sampled comments).

Category	Samples in Set	True	Accuracy Rate (%)
Cultural or regional identity and narrative	171	126	73.68
Direct situation report	150	143	95.33
Information source and trust discussion	28	28	100
Mutual aid, wishes, and concern	29	25	86.21
Risk attenuation and dismissal	76	59	77.63
Risk perception and emotion (humor or sarcasm)	176	159	90.34
Risk perception and emotion (worry or fear)	22	22	100
Social and life impact	112	105	93.75
Typhoon path and regional speculation	183	177	96.72
Other	53	45	84.91
Total	1,000	883	88.3

[1] This study utilized the official open-source release of DeepSeek-R1-671B, deployed as the vanilla version (without quantization or fine-tuning) on our institutional high-performance computing (HPC) cluster. The only configurable inference parameter was Temperature, which was set to 0 to ensure deterministic and reproducible outputs. All other model specifications, including the 128K-token context window and architectural details, are publicly available at: https://huggingface.co/deepseek-ai/DeepSeek-R1 . As the HPC infrastructure is provided to researchers at no direct cost, per-query inference costs are not reported; however, the use of the publicly available model weights ensures full reproducibility on any compatible infrastructure.

To further verify the methodological necessity of our approach, we conducted an ablation study comparing the few-shot workflow against a zero-shot baseline. The results demonstrated that the few-shot strategy significantly outperformed the zero-shot approach (which achieved 82.1% accuracy, see Appendix Table A1), particularly in disambiguating informal and context-heavy categories.

To validate the consistency of the classification framework in this sole-authored study, an intra-coder reliability test was performed to assess the stability of the annotation logic over time. A random validation set of 1,000 comments was subjected to two rounds of blind coding with a 10-month interval (T_1: May 2025; T_2: March 2026). Between these periods, the author further refined the Dominant-Intent Priority Hierarchy to resolve semantic overlaps.

The consistency between the two longitudinal sessions was evaluated using Cohen's Kappa (Cohen, 1960). The analysis revealed a kappa value of 0.89, indicating strong agreement under the specific coding scheme and test conditions (Landis & Koch, 1977). Of the 1,000 samples, 97.9% remained consistent, with minor discrepancies (2.1%) primarily occurring in early-stage labels that were subsequently corrected by stricter adherence to the Fact-First and Information Scrutiny rules in the final iteration. This high level of longitudinal stability confirms the robustness of the analytical pipeline and the reliability of the 88.3% LLM classification accuracy.

3.2.3. *Descriptive analysis and provincial aggregation*

Following the validated classification of all 436,964 comments, two initial descriptive analyses were performed, as shown in the Stage 2 workflow (Fig. 2). **Overall Topic Distribution:** The total count for each of the 10 themes were calculated to understand the overall landscape of the discourse. **Provincial Aggregation:** Using the IP-based location data for each comment, the classified data was grouped by province. For each province, the proportional prevalence of each topic (i.e., the percentage of that province's total comments belonging to a specific theme) was calculated. This proportional data serves as the primary dependent variable for the analyses in Stage 3.

3.3. *Visualization and analysis*

To investigate the relationship between provincial socio-economic contexts (the nine independent variables) and the prevalence of nine substantive discussion themes (excluding the residual 'Other' category), we conducted a statistical analysis on the data from 31 data available provinces. First, all independent variables were tested for multicollinearity using the Variance Inflation Factor (VIF) (O'brien, 2007). The results (Fig. 5) confirmed that all indicators were within acceptable limits for regression analysis (all VIFs < 5.3, with eight of the nine indicators below VIF 4.0).

Second, because the dependent variables are proportions (bounded between 0 and 1), standard OLS regression is inappropriate. We therefore employed nine **separate and independent** Multiple Beta Regression models (Cribari-Neto & Zeileis, 2010), one for each topic. This independent modeling structure ensures that the statistical estimation for each discourse type is decoupled, preventing potential classification noise in one category from propagating to or biasing the results of others. Given the 81 simultaneous regression tests (9 topics $\times$ 9 indicators), we applied the Benjamini-Hochberg (FDR) procedure (alpha $= 0.1$) to all resulting p-values to control the false discovery rate (Thissen et al., 2002). Model residuals were also checked for spatial autocorrelation using Moran's I (Zhang et al., 2008) to ensure that the identified socio-economic associations were not artifacts of spatial clustering.

Finally, to interpret the model's results at a local level, we used SHAP (SHapley Additive exPlanations)(Lundberg & Lee, 2017). While regression coefficients show the *average* effect of a factor, SHAP analysis allows us to decompose each prediction and quantify the specific contribution of each socio-economic indicator to the prevalence of a

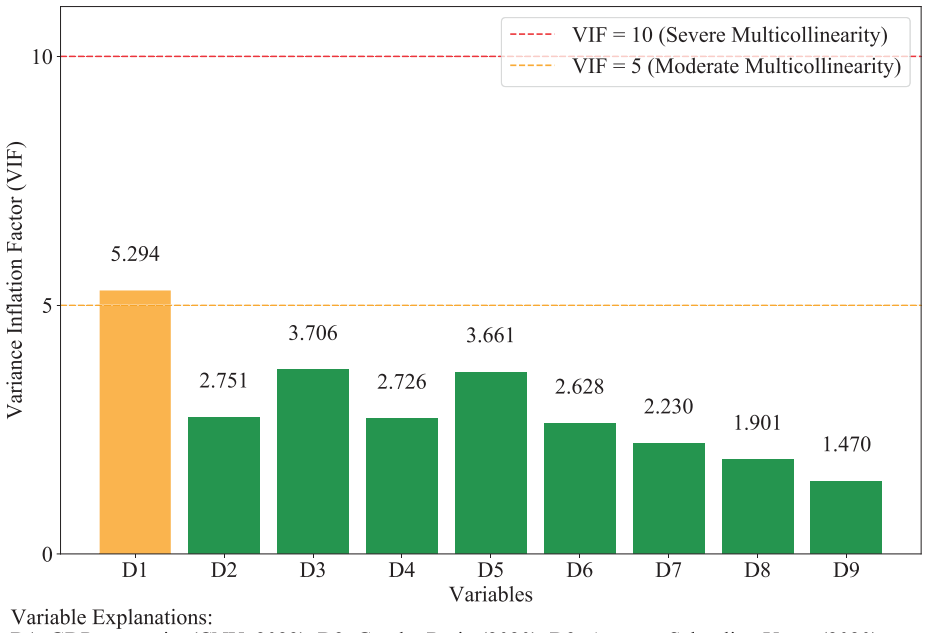

Fig. 5. Variance Inflation Factor (VIF) results for independent variables.

topic *within each individual province* (Zhang et al., 2023). SHAP values represent the marginal contribution of a feature to the model's prediction for a specific instance. While they provide a powerful lens for local interpretability, SHAP values reflect the statistical contribution to model prediction and do not necessarily imply a direct causal effect. In our spatial analysis, a negative SHAP value for a province indicates that its socio-economic profile contributes to a lower predicted prevalence of a specific theme compared to the global average of the dataset (Zhang et al., 2026b). As a key feature of our study, these individual SHAP values were then mapped using GIS, enabling a spatial interpretation of the model's findings. This pipeline constitutes a generalizable framework designed for application to visually driven, algorithmically curated platforms in collectivist contexts, extending beyond the specific case of Typhoon Gaemi on Douyin.

4. Results

This chapter presents the empirical findings derived from the multi-stage research methodology. First, the overall results of the thematic classification of the *Douyin* comment corpus are presented, establishing the general landscape of public discourse. Following this, the chapter delves into a detailed spatial analysis, using GIS visualizations to explore the geographical distribution of each discussion theme. Finally, the results of the Beta Regression and SHAP analyses are presented to examine the associations between socio-economic factors and these observed patterns.

4.1. Text classification results

As detailed in the methodology (Section 3.2), the entire corpus of 496,964 valid *Douyin* comments was processed and classified using the validated LLM-based few-shot learning approach. This classification produced a structured overview of the public's primary topics of discussion during the Typhoon Gaemi event. Table 3 presents the final distribution of comments across the ten defined thematic categories, showing the raw comment count and the corresponding percentage of the total discourse.

As Table 3 demonstrates, the public discourse surrounding Typhoon Gaemi was diverse and multifaceted. The four most dominant themes were *Typhoon Path and Regional Speculation* (17.76%), *Cultural or Regional Identity and Narrative* (17.20%), *Direct Situation Report* (16.13%), and *Risk Perception and Emotion (Humor or Sarcasm)* (15.78%). These four categories alone accounted for over two-thirds (66.87%) of the entire conversation.

This indicates that the public's focus was almost evenly split between: Practical Information-Seeking: Users sought and shared information about the storm's trajectory and their real-time weather conditions. Socio-Cultural Sense-Making: Users processed the event through the lenses of cultural identity and humor, engaging in collective coping and social bonding.

Discussions of tangible *Social and Life Impact* also constituted a significant portion of the discourse (11.22%). The remaining themes, while less frequent, highlight important facets of the response: *Risk Attenuation and Dismissal* (7.35%) was notably more common than overt expressions of *Risk Perception and Emotion (Worry or Fear)* (2.28%), and discussions on *Information Source and Trust* (3.53%) and *Mutual Aid, Wishes, and Concern* (3.76%) represented smaller, more specific conversational niches.

Table 3
Overall distribution of thematic categories (N = 436,964).

Categories	Count	Percentage (%)
Cultural or regional identity and narrative	75,153	17.20
Direct situation report	70,494	16.13
Information source and trust discussion	15,410	3.53
Mutual aid, wishes, and concern	16,408	3.76
Risk attenuation and dismissal	32,103	7.35
Risk perception and emotion (humor or sarcasm)	68,937	15.78
Risk perception and emotion (worry or fear)	9,967	2.28
Social and life impact	49,021	11.22
Typhoon path and regional speculation	77,623	17.76
Other	21,848	5.00

4.2. Regional variations in discussion topic proportions

While the aggregate distribution (Table 3) provides a national overview, it obscures crucial regional differences in public focus. To investigate how the discourse varied geographically, the 436,964 classified comments were geolocated to their province of origin (as shown in Fig. 3). For each province, the proportional prevalence of each theme (i.e., the percentage of that province's total comments belonging to a specific theme) was calculated. Table 4 presents a comprehensive overview of this regional distribution, showing both the raw comment count and the percentage proportion for every topic across all provinces. The heatmap visualization within the table highlights the relative intensity of each theme.

This spatial analysis reveals that the public response to Typhoon Gaemi was not monolithic across China. Instead, distinct regional patterns emerged, often linked to the physical proximity of the threat and local context. The geographical distributions of four key themes, presented in Fig. 6 through Fig. 9, illustrate these dynamics.

A clear pattern emerges related to uncertainty and physical distance. As shown in Fig. 6, the theme of *Typhoon Path and Regional Speculation* was most prevalent in inland and northern provinces, far from the typhoon's most direct impacts. The highest concentrations were found in *Inner Mongolia* (32.59%), *Gansu* (29.61%), *Ningxia* (29.33%), and *Shaanxi* (28.31%). For these distant populations, the primary mode of engagement was monitoring and speculating on the storm's trajectory. Spatial autocorrelation analysis (Moran's I = 0.3168, p = 0.004) shows that the geographical distribution of this topic is not random but exhibits a strong spatial positive correlation and clustering effect, discussions about path prediction are more prevalent in northern and inland provinces.

In direct contrast, discourse focused on immediate, tangible reality dominated the regions most affected by the storm. As seen in Fig. 7, *Direct Situation Report* was highly concentrated in the central and southern provinces that were on high alert for heavy rainfall. Spatial autocorrelation analysis confirms this pattern, yielding a Moran's I value of 0.3557 (p = 0.0050), which indicates a statistically significant positive spatial clustering. The highest proportions were found in *Hunan* (28.73%), *Jiangxi* (27.87%), *Henan* (24.48%), and the coastal province of *Guangdong* (21.88%). For these residents, the risk was not an abstract speculation but a present reality to be documented and shared.

Beyond immediate risk, themes of cultural sense-making also showed unique spatial patterns. Fig. 8 illustrates that *Cultural or Regional Identity and Narrative* had an exceptionally strong concentration in *Shanxi* (40.30%), an inland province not directly in the storm's path. This suggests a powerful tendency in some regions to frame a distant national event through local cultural lenses, history, and identity. High proportions were also seen in other northern provinces like *Shandong* (23.36%) and *Shaanxi* (23.61%), as well as the directly impacted *Fujian* (22.25%), indicating that cultural framing is a versatile tool used by both those directly and indirectly affected. A statistically significant clustering pattern was confirmed for regional identity discourse, with a Moran's I = 0.3244 (p = 0.002). This suggests a strong spatial dependency, where regions with high identity-related discussions tend to be geographically proximal, forming distinct narrative hubs across the map.

Finally, the theme of *Risk Attenuation and Dismissal* (Fig. 9) reveals a notable "bimodal", or "U-shaped," distribution. This discourse was highly prevalent in two distinct types of regions: Experienced Coastal Provinces: such as *Hainan* (10.90%) and *Fujian* (10.13%), where high prevalence may be driven by habituation and a normalization of the threat ("typhoon fatigue"). Remote Inland Provinces: such as *Xinjiang* (10.84%) and *Qinghai* (10.52%), where high prevalence was likely driven by perceived physical distance and a sense of invulnerability or irrelevance. There is a clear spatial clustering effect in the discourse of risk dismissal, as evidenced by a Moran's I value of 0.3461 (p = 0.003). This indicates that the tendency to downplay the storm's risk is not randomly distributed but is geographically concentrated among neighboring provinces in the non-affected zones.

To provide a comprehensive overview of the digital risk landscape, Fig. 10 visualizes the spatial distribution of the remaining themes that exhibited significant spatial autocorrelation. Specifically, 'Information Source and Trust Discussion' (Fig. 10a) shows notable concentrations across several northern provinces. The socio-economic correlations of this spatial pattern are formally examined in Section 4.3. 'Mutual Aid, Wishes, and Concern' (Fig. 10b) exhibits high prevalence in Sichuan and the Northeast, potentially reflecting strong regional community bonds or prior disaster experience. Interestingly, 'Risk Perception and Emotion (Humor or Sarcasm)' (Fig. 10c) displays a western and southwestern clustering (e.g., Xinjiang, Chongqing), suggesting regional variations in affective coping styles. 'Social and Life Impact' (Fig. 10d) is understandably concentrated in the eastern coastal provinces directly intersected by Typhoon Gaemi's path.

Notably, 'Risk Perception and Emotion (Worry or Fear)' was the only category that did not exhibit significant spatial clustering (Moran's I = 0.0316, p = 0.2770, see Fig. 11). This suggests that while cognitive appraisals (like information scrutiny) and coping mechanisms (like humor) are spatially associated with by socio-economic contexts, acute emotional distress may act as a more universal and diffuse response that transcends regional boundaries during a national-level crisis.

4.3. Socio-economic correlates of topic prevalence

While Section 4.2 revealed notable provincial-level variation in topic prevalence, for instance, the prominent role of Direct Situation Report in agricultural provinces like Henan, these patterns alone do not explain why such differences exist. This section therefore moves from description to explanation by modeling the relationship between topic prevalence and provincial-level socio-economic factors. Crucially, Moran's I tests on the initial model residuals (summarized in the right-most column of Fig. 12) revealed that 5 of the 9 models exhibited significant positive spatial autocorrelation. This means that standard regression results would be unreliable. Therefore, to account for these spatial effects, Spatial Beta Regression models (specifically, Spatial Error Models (SEM), based on model diagnostics) were employed (Rüttenauer, 2019). The final, robust results of these spatial regression models are presented in Fig. 12.

The regression models demonstrated a good overall fit, with a mean Pseudo R^2 of 0.427 across all nine models, indicating that the selected socio-economic indicators could explain a substantial portion of the variance in topic prevalence. After applying the strict Benjamini-Hochberg (FDR) correction (alpha = 0.1) across all 81 (9x9) tests, three relationships remained statistically significant. The detailed statistical parameters for all 81 regression tests, including standard errors, t-statistics, and FDR-corrected q-values, are provided in the **Supplementary Materials** (Excel File).

First, a strong positive relationship was found between *Direct Situation Report* and 'Cropland Cover Proportion' (β = 0.764, q = 0.086). This indicates that provinces with a higher percentage of land dedicated to agriculture had a significantly higher proportion of real-time situational reports. This pattern is consistent with regions with extensive cropland (such as *Henan* and *Shandong*) which are highly vulnerable to flooding and waterlogging from typhoon-related rainfall, posing a direct threat to agricultural livelihoods. This vulnerability context co-occurs with more residents observing and reporting on immediate weather conditions. This statistical finding provides a socio-economic explanation for the provincial pattern observed in Section 4.2, where Henan and Shandong—provinces with among the highest cropland proportions nationally—were visually prominent in the prevalence of *Direct Situation Report*.

Second, the strongest and most significant relationship identified across all 81 tests was between *Information Source and Trust Discussion* and 'Cropland Cover Proportion' (β = 0.871, q = 0.009). This result,

Table 4
Detailed distribution of thematic categories by province.

Province (IP)	C1	C2	C3	C4	C5	C6	C7	C8	C9	Others
Anhui	15.2 (1,171)	8.8 (681)	3.2 (247)	5.5 (424)	5.5 (427)	18.8 (1,445)	2.5 (195)	9.7 (749)	25.2 (1,941)	5.5 (423)
Beijing	18.0 (1,080)	5.6 (336)	5.3 (319)	6.9 (411)	4.7 (280)	17.8 (1,065)	2.4 (143)	10.4 (621)	23.9 (1,431)	5.2 (309)
Chongqing	18.1 (596)	4.0 (133)	1.6 (53)	8.1 (265)	5.3 (174)	19.6 (643)	1.6 (51)	8.2 (271)	25.7 (843)	7.8 (257)
Fujian	22.2 (16,607)	12.9 (9,599)	2.1 (1,545)	2.4 (1,804)	10.1 (7,563)	15.3 (11,422)	2.9 (2,202)	15.9 (11,894)	13.0 (9,700)	3.1 (2,310)
Gansu	21.4 (387)	2.2 (39)	1.7 (30)	6.4 (116)	9.4 (171)	13.6 (246)	1.5 (28)	6.1 (110)	29.6 (536)	8.1 (147)
Guangdong	13.5 (7,987)	21.9 (12,980)	2.4 (1,445)	3.4 (2,023)	5.3 (3,144)	15.3 (9,081)	2.3 (1,337)	10.5 (6,236)	16.1 (9,569)	9.3 (5,523)
Guangxi	15.6 (946)	7.6 (458)	2.3 (137)	6.2 (376)	5.9 (356)	19.0 (1,153)	2.2 (135)	11.2 (678)	21.9 (1,325)	8.2 (495)
Guizhou	11.9 (845)	23.8 (1,692)	3.1 (218)	4.8 (339)	8.1 (573)	14.5 (1,035)	1.5 (104)	13.1 (929)	13.5 (963)	5.8 (416)
Hainan	21.4 (554)	5.7 (147)	2.0 (52)	5.7 (147)	10.9 (282)	17.1 (441)	2.7 (70)	10.1 (260)	19.9 (515)	4.6 (118)
Hebei	19.7 (5,050)	13.9 (3,583)	9.9 (2,545)	3.1 (805)	5.7 (1,474)	16.1 (4,125)	2.2 (574)	8.4 (2,147)	17.0 (4,378)	4.0 (1,018)
Heilongjiang	22.0 (623)	4.8 (136)	1.8 (51)	8.6 (243)	6.9 (195)	11.7 (332)	1.8 (51)	7.6 (215)	28.8 (817)	6.0 (169)
Henan	17.2 (9,624)	24.5 (13,715)	5.7 (3,195)	2.6 (1,467)	6.8 (3,822)	12.8 (7,148)	2.4 (1,339)	10.4 (5,835)	14.1 (7,914)	3.5 (1,966)
Hong Kong	16.7 (60)	14.4 (52)	1.9 (7)	8.6 (31)	4.2 (15)	14.2 (51)	2.5 (9)	16.1 (58)	15.8 (57)	5.6 (20)
Hubei	14.1 (1,165)	11.3 (937)	3.1 (256)	5.4 (450)	5.0 (412)	19.6 (1,623)	1.7 (137)	8.6 (713)	25.7 (2,124)	5.6 (460)
Hunan	10.2 (1,808)	28.7 (5,097)	2.5 (442)	4.4 (772)	4.5 (802)	14.7 (2,600)	1.7 (306)	11.9 (2,106)	17.1 (3,031)	4.4 (776)
Inner Mongolia	24.3 (954)	2.5 (98)	1.4 (54)	4.0 (159)	9.2 (361)	12.9 (508)	1.8 (72)	4.2 (166)	32.6 (1,280)	7.0 (275)
Jiangsu	15.8 (2,659)	8.9 (1,498)	3.2 (541)	4.8 (813)	6.0 (1,012)	18.1 (3,038)	1.9 (320)	10.6 (1,777)	26.0 (4,379)	4.6 (777)
Jiangxi	9.4 (2,830)	27.9 (8,396)	2.4 (714)	2.5 (758)	10.7 (3,219)	16.9 (5,102)	2.7 (816)	12.6 (3,791)	11.2 (3,362)	3.8 (1,140)
Jilin	19.8 (841)	10.7 (455)	2.2 (92)	5.9 (250)	5.5 (235)	13.8 (584)	2.4 (102)	16.9 (717)	18.8 (796)	4.0 (169)
Liaoning	16.9 (1,282)	13.1 (992)	4.0 (302)	6.3 (474)	5.6 (422)	14.8 (1,119)	2.1 (158)	9.4 (712)	23.5 (1,783)	4.4 (333)
Macau	18.6 (16)	9.3 (8)	2.3 (2)	7.0 (6)	8.1 (7)	9.3 (8)	0.0 (0)	11.6 (10)	22.1 (19)	11.6 (10)
Ningxia	20.1 (166)	2.3 (19)	1.1 (9)	7.2 (59)	8.7 (72)	16.2 (134)	1.3 (11)	5.1 (42)	29.3 (242)	8.6 (71)
Qinghai	19.1 (118)	1.8 (11)	1.6 (10)	6.0 (37)	10.5 (65)	15.5 (96)	1.6 (10)	6.3 (39)	28.6 (177)	8.9 (55)
Shaanxi	23.6 (863)	3.6 (133)	2.2 (79)	6.5 (237)	5.4 (199)	14.3 (523)	1.8 (66)	7.5 (274)	28.3 (1,035)	6.8 (247)
Shandong	23.4 (4,447)	8.4 (1,604)	5.0 (959)	4.2 (804)	6.1 (1,154)	17.0 (3,240)	1.9 (368)	8.5 (1,614)	20.9 (3,977)	4.6 (873)
Shanghai	15.1 (845)	9.8 (551)	3.0 (168)	4.9 (272)	5.7 (322)	17.9 (1,002)	1.5 (86)	11.9 (669)	24.9 (1,396)	5.2 (290)
Shanxi	40.3 (2,624)	3.1 (203)	1.6 (101)	4.6 (301)	5.0 (326)	11.4 (745)	1.2 (75)	6.1 (399)	23.5 (1,528)	3.2 (209)
Sichuan	19.3 (1,228)	3.1 (195)	1.7 (111)	9.5 (604)	5.4 (347)	19.0 (1,212)	1.8 (115)	6.7 (430)	25.7 (1,640)	7.7 (491)
Taiwan	16.6 (320)	20.8 (400)	0.9 (18)	8.7 (168)	7.7 (149)	12.0 (232)	3.1 (59)	21.2 (408)	7.2 (139)	1.7 (33)
Tianjin	15.6 (401)	9.0 (230)	4.7 (120)	5.3 (136)	5.1 (132)	17.9 (461)	2.1 (53)	9.7 (249)	25.7 (661)	4.9 (126)
Tibet	19.5 (52)	3.7 (10)	1.1 (3)	5.2 (14)	10.5 (28)	15.0 (40)	3.4 (9)	6.7 (18)	24.7 (66)	10.1 (27)
Xinjiang	21.1 (690)	1.7 (57)	1.5 (48)	4.7 (152)	10.8 (354)	18.4 (600)	2.7 (88)	3.9 (128)	23.6 (772)	11.6 (378)
Yunnan	20.6 (626)	2.8 (85)	2.0 (60)	8.7 (264)	6.0 (182)	16.2 (493)	1.9 (59)	7.0 (214)	26.1 (793)	8.8 (268)
Zhejiang	13.9 (5,688)	14.5 (5,964)	3.6 (1,477)	3.0 (1,227)	9.3 (3,827)	18.0 (7,390)	2.0 (819)	11.1 (4,542)	20.6 (8,434)	4.1 (1,669)

Percentage (%)
20 40 10 20 2.5 5.0 7.5 2.5 5.0 7.5 5 10 10 15 0 2 10 20 10 20 30

C1: Cultural or Regional Identity and Narrative C2: Direct Situation Report C3: Information Source and Trust Discussion C4: Mutual Aid, Wishes, and Concern
C5: Risk Attenuation and Dismissal C6: Risk Perception and Emotion (Humor or Sarcasm) C7: Risk Perception and Emotion (Worry or Fear) C8: Social and Life Impact
C9: Typhoon Path and Regional Speculation

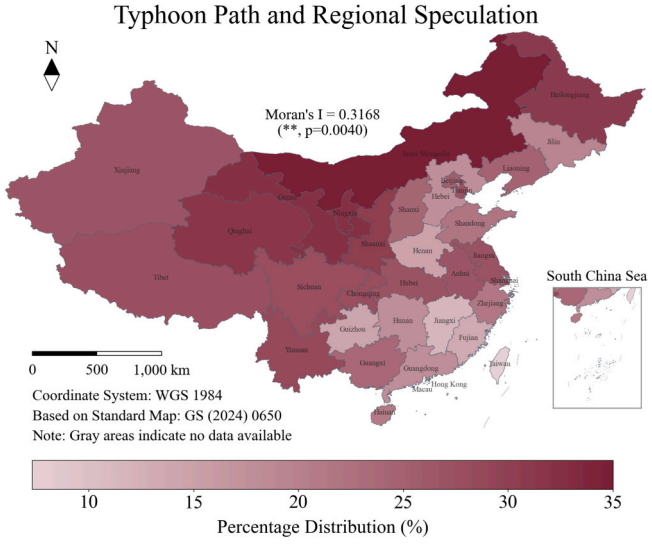

Fig. 6. Geographical distribution of "typhoon path and regional speculation".

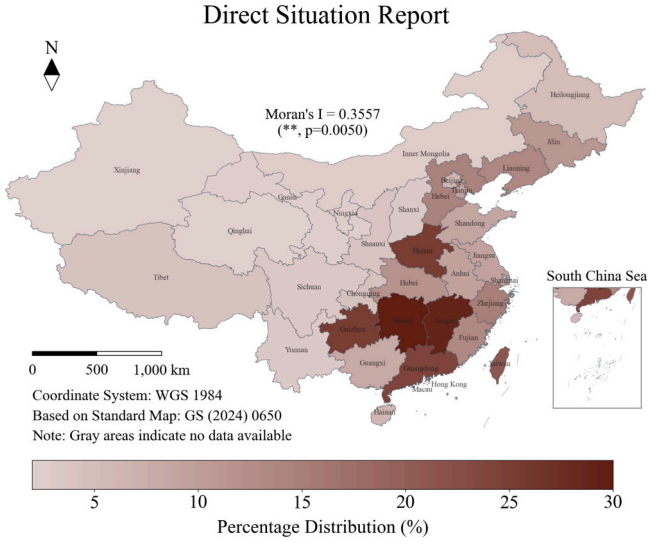

Fig. 7. Geographical distribution of "direct situation report".

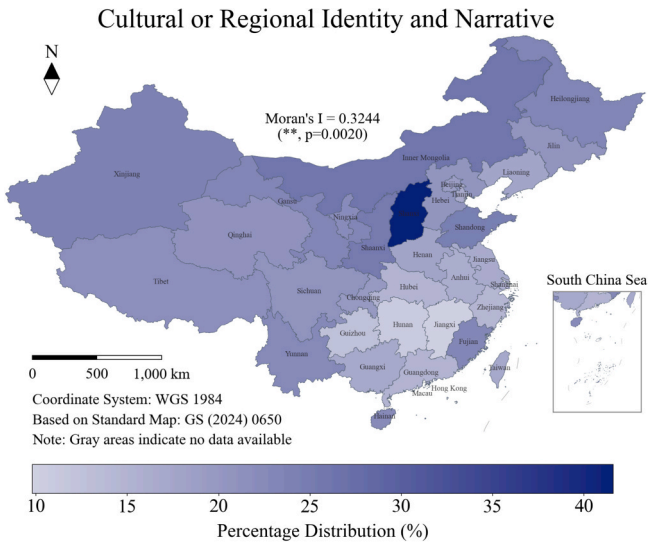

Fig. 8. Geographical distribution of "cultural or regional identity and narrative".

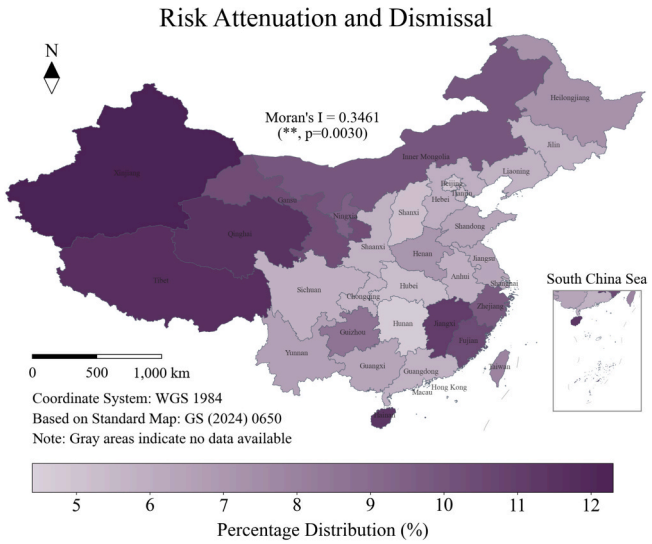

Fig. 9. Geographical distribution of "risk attenuation and dismissal".

significant even at the stringent $q < 0.01$ level, reveals that agricultural provinces are not only the primary sources of direct situational reports (as shown above), but are also characterized by significantly heightened critical discourse about the accuracy and reliability of information sources. This dual association suggests that in regions where agricultural livelihoods are directly threatened by typhoon-related hazards, the public's stake in accurate forecasting is exceptionally high—associated with both real-time environmental reporting and active scrutiny of whether official information sources can be trusted. The emergence of Cropland Cover Proportion as the most consistently significant correlation across two distinct discourse categories underscores the central role of agricultural vulnerability as a multifaceted contextual factor associated with the regional disaster discourse landscape.

Third, *Information Source and Trust Discussion* was also positively correlated with 'Natural Disaster Economic Loss per capita' ($\beta = 0.636$, $q = 0.086$). This result contextualizes the elevated prevalence of trust-related discourse observed in Section 4.2 for provinces like Hebei, whose recent histories of significant disaster-related economic losses co-occur with residents' tendency to scrutinize information quality more actively. This suggests that populations in provinces that have previously suffered greater per capita economic losses from disasters are more sensitized to the importance of accurate information, which is associated with more critical scrutiny of forecasts and reports during a new event.

In summary, the emergence of Cropland Cover Proportion as the most pervasive spatial correlate after FDR correction—significantly predicting both Direct Situation Report and Information Source and Trust Discussion—highlights that agricultural vulnerability functions as a multifaceted contextual factor, associated with both immediate environmental monitoring and a heightened demand for reliable disaster information. The additional independent association with Natural Disaster Economic Loss per capita reinforces the interpretation that accumulated disaster experience further compounds this critical orientation toward information reliability.

4.4. Model interpretation using SHAP and GIS

While the regression coefficients in Fig. 12 identify the average, global relationships between variables, they do not explain the model's feature importance patterns for specific provinces. To "unpack" these findings and understand *how* local contexts contribute to the model's predictions, we employed SHAP (SHapley Additive exPlanations) analysis.

First, to validate the alignment between the underlying regression and the SHAP analysis, we compared the SHAP values against the regression coefficients. As shown in Fig. 13a, there is a clear positive relationship between the absolute WLS coefficient and the SHAP importance, confirming that variables deemed significant in the regression (red diamonds) are also identified as highly important by SHAP. Fig. 13b further validates this by showing that the sign of the Regression Coefficient (positive or negative effect) aligns with the mean SHAP value.

With the model's interpretability validated, we then applied our novel SHAP-GIS visualization method to spatially decompose the three significant relationships identified in Section 4.3. This allows us to move beyond *what* the relationship is to *where* and *how* this relationship manifests. The Beta Regression identified a significant positive relationship between a province's 'Cropland Cover Proportion' and the prevalence of *Direct Situation Report*. Fig. 14 clearly shows that provinces with high cropland proportions, particularly major agricultural hubs like Henan and Shandong, exhibit strong positive SHAP values (deep red), indicating that the extensive agricultural land in these regions contributed positively to the model's prediction for direct situational reports. This reinforces the interpretation that direct vulnerability to agricultural damage during a typhoon is strongly associated with immediate, real-time reporting from affected populations. Conversely, regions with very low cropland cover, such as Tibet, show strong negative SHAP contributions (deep blue), demonstrating that low cropland cover has reduced predictive importance for *Direct Situation Report* in these regions.

The strongest positive relationship found was between 'Cropland Cover Proportion' and *Information Source and* Trust *Discussion*. Fig. 15 spatially unpacks this correlation, highlighting that most agricultural provinces, such as *Jiangsu*, *Shandong*, and *Henan*, exhibit substantial positive SHAP values. This indicates that their higher proportion of cropland is associated with a greater predicted prevalence of critical discussions surrounding information veracity. This underscores how regional agricultural exposure is associated with heightened public concern regarding information trust during disaster events. Conversely, provinces with minimal cropland cover, such as *Tibet*, *Qinghai*, and the highly urbanized *Beijing*, show negative SHAP contributions for this specific indicator.

Finally, 'Natural Disaster Economic Loss per capita' also showed a significant positive association with *Information Source and Trust Discussion*. Fig. 16 reveals that provinces with a history of significant per capita economic losses from natural disasters, such as *Beijing* and *Hebei*,

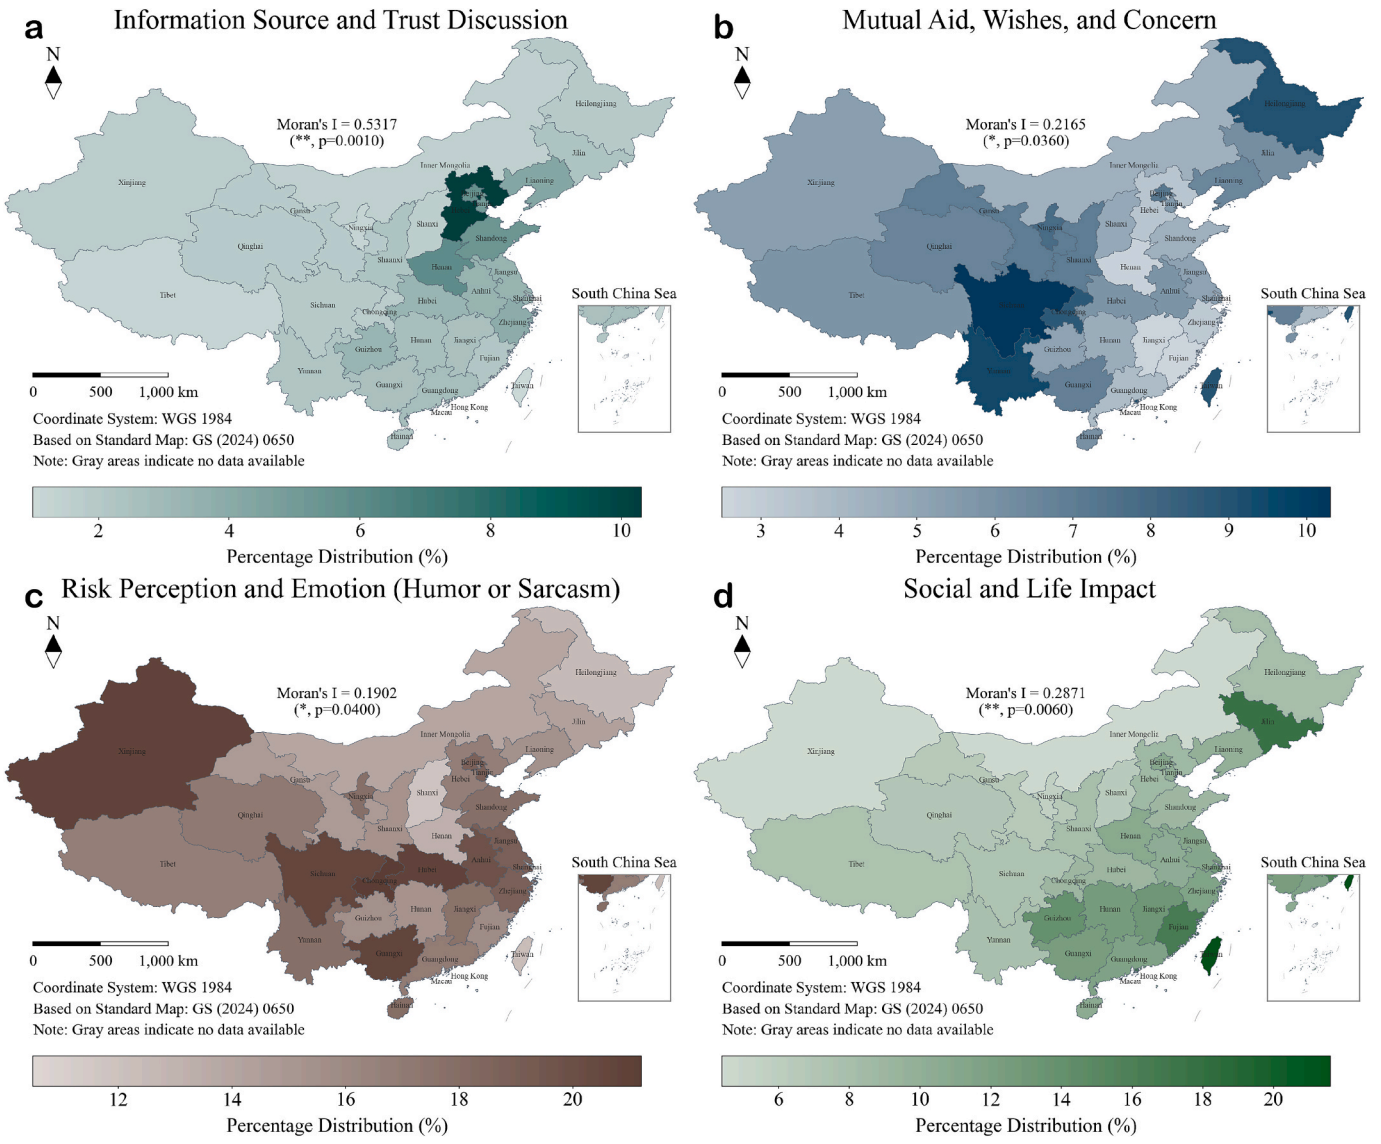

Fig. 10. Geographical distribution of 'information source and trust discussion', 'mutual aid, wishes, and concern', 'risk perception and emotion (humor or sarcasm)', and 'social and life impact'.

exhibit strong positive SHAP values. This suggests that prior adverse experiences are associated with heightened public sensitivity to the reliability of disaster-related information, which co-occurs with more prevalent discussions about trust and sources. Notably, some provinces with relatively high recent economic loss, like *Shandong*, show negative SHAP values, indicating that this feature's marginal contribution to the model's prediction was negative for these specific provinces, suggesting that other local contextual factors exert a stronger influence and the relationship between disaster loss and trust discourse is not spatially uniform. This nuance, revealed by the SHAP-GIS approach, enriches our understanding beyond the aggregated regression coefficient.

Taken together, the three SHAP-GIS maps demonstrate that the statistically significant socio-economic correlations identified in Section 4.3 do not operate uniformly across the country. Instead, their predictive contributions are geographically concentrated in provinces where the underlying socio-economic conditions are most pronounced: agricultural heartlands for both cropland-driven situational reporting and information trust discourse, and disaster-experienced provinces for trust-related scrutiny. This spatial specificity reinforces the value of combining regression analysis with local-level interpretability.

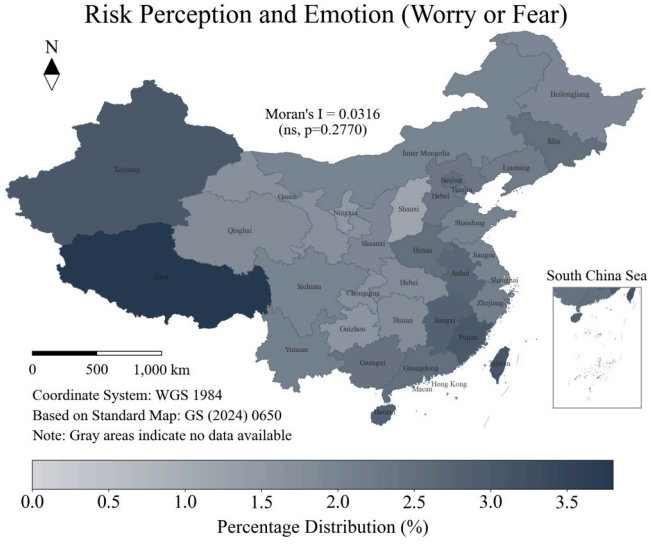

Fig. 11. Geographical distribution of "risk perception and emotion (worry or fear)".

5. Discussion

This study employed an advanced Large Language Model (LLM) and a spatial regression framework to analyze public risk discourse on *Douyin* during Typhoon Gaemi. Our analysis identified ten distinct thematic categories, revealed significant regional variations in their prevalence, and, most importantly, quantified the socio-economic factors associated with these patterns. This section interprets these core findings within our theoretical framework, highlighting the study's contributions to risk discourse and communication research.

5.1. The thematic landscape: information exchange vs. socio-cultural sense-making

The first research question sought to identify the principal themes and characteristics of public risk discourse on *Douyin*. The thematic classification of 436,964 comments (presented in Table 3) reveals a landscape far more complex than simple information dissemination. The four most dominant themes—*Typhoon Path and Regional Speculation* (17.76%), *Cultural or Regional Identity and Narrative* (17.20%), *Direct Situation Report* (16.13%), and *Risk Perception and Emotion (Humor or Sarcasm)* (15.78%)—exhibited a remarkably balanced distribution. By situating these categories within the Protective Action Decision Model (PADM), we observe that the discourse is not merely a reflection of a single cognitive state but a manifestation of a multi-stage social-cognitive progression, ranging from initial information exposure to complex affective appraisal (Terpstra & Lindell, 2012).

This finding presents a notable challenge to the information-centric paradigm that has long dominated social media risk communication literature (Alexander, 2013). Traditional studies, particularly those focused on text-centric platforms like Twitter or Weibo during hurricanes and floods, often report that functional information-seeking and situational reporting dominate over 50–60% of the total volume (Cao et al., 2021; Dong et al., 2024; Zhang & Cozma, 2022). In contrast, our data demonstrates a unique "coping-information parity" on *Douyin*: **Informational Exchange (***Typhoon Path and Regional Speculation* +

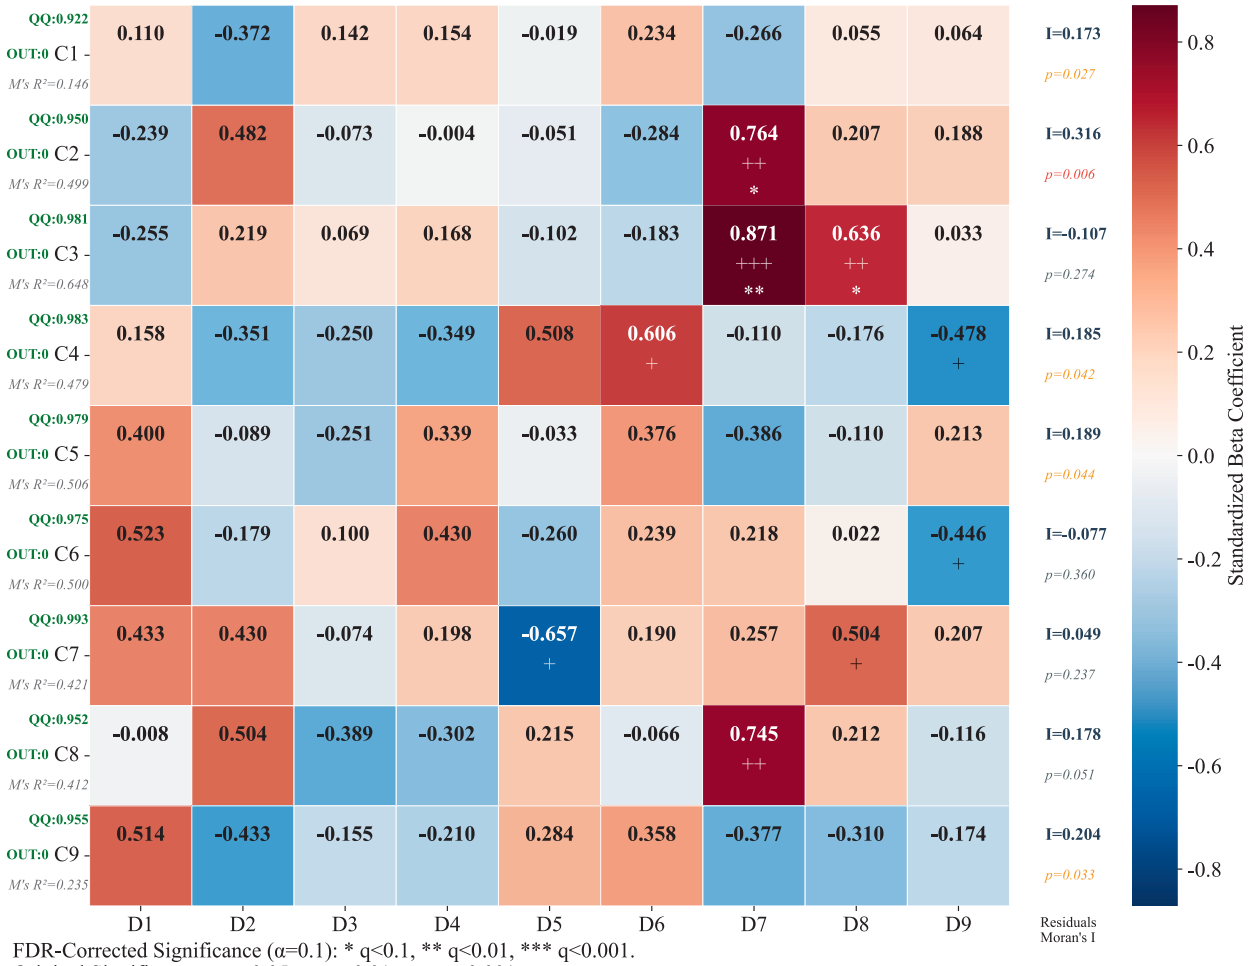

FDR-Corrected Significance (α=0.1): * q<0.1, ** q<0.01, *** q<0.001.
Original Significance: + p<0.05, ++ p<0.01, +++ p<0.001.
QQ: Q-Q correlation in logit space (>0.92 excellent, 0.85-0.92 acceptable, <0.85 poor).
OUT: Pearson residual outliers >3σ in original proportion space.
Socioeconomic Development Indicators:
D1: GDP per capita (CNY, 2023) D2: Gender Ratio (2020) D3: Average Schooling Years (2020)
D4: Population Growth Rate (2014-2023) D5: Internet Penetration Rate (2023) D6: Primary Industry GDP Proportion (2023)
D7: Cropland Cover Proportion (2023) D8: Natural Disaster Economic Loss per capita (10K CNY, 2023)
D9: Natural Disaster Affected Area per capita (Hectares, 2023)
Text Classification Category:
C1: Cultural or Regional Identity and Narrative C2: Direct Situation Report C3: Information Source and Trust Discussion
C4: Mutual Aid, Wishes, and Concern C5: Risk Attenuation and Dismissal C6: Risk Perception and Emotion (Humor or Sarcasm)
C7: Risk Perception and Emotion (Worry or Fear) C8: Social and Life Impact C9: Typhoon Path and Regional Speculation

Fig. 12. Heatmap of standardized beta coefficients from spatial beta regression models.

Direct Situation Report)**:** Accounted for 33.89% of comments. **Socio-Cultural Sense-Making (***Cultural or Regional Identity and Narrative + Risk Perception and Emotion (Humor or Sarcasm)***):** Accounted for 32.98% of comments.

This result provides direct empirical evidence for the Social Amplification of Risk Framework (SARF). It shows that social media platforms like *Douyin* are not merely passive conduits for expert information but are active and powerful "amplification stations" for *interpretive signals* (Kasperson et al., 1988; Wirz et al., 2018). In the collectivist and emerging-economy context of China (Fan, 2020), the high prevalence of cultural narratives suggests that a primary function of public discourse during Typhoon Gaemi was not merely to acquire factual knowledge about the threat, but to process and normalize the risk through shared identity and affect.

The spatial concentration of Cultural or Regional Identity, particularly in non-impacted or inland provinces, aligns with Risk Perception Theory (RPT), emphasizing that risk is filtered through existing cultural schemas (Rippl, 2002; Wilkinson, 2001). Similarly, the high prevalence of *Humor or Sarcasm* demonstrates a critical collective coping mechanism. This aligns with the affect heuristic (Finucane et al., 2000), showing a public effort to manage the dread of the risk through shared affect, a common behavior on social media platforms (Dong & Xie, 2024). These findings suggest that on algorithmically curated platforms, disaster discourse is fundamentally a dual-track process: while individuals use the platform for situational awareness (Stage I of PADM), they simultaneously utilize it for collective meaning-making and psychological resilience (Stage IV of PADM).

5.2. Socio-economic context associated with risk discourse patterns

The second research question, and a core focus of this work, explored the statistical association between provincial-level socio-economic factors and the prevalence of online risk discourse themes. The application of Spatial Beta Regression (Fig. 12) provided a robust model, and the SHAP-GIS analysis (Figs. 14–16) offered unprecedented, localized interpretability of feature importance. The results strongly support the hypothesis that macro-level societal contexts are systematically associated with collective online discourse. This finding empirically grounds the "amplification station" concept within SARF (Kasperson et al., 1988; Kasperson et al., 2022), demonstrating that regional socio-economic characteristics are significant predictors of the thematic texture of public risk processing. Two core findings emerged after stringent FDR correction across all 81 tests:

First, and most strikingly, Cropland Cover Proportion emerged as the most pervasive spatial correlate in the entire study, significantly predicting two distinct discourse categories: Direct Situation Report ($\beta = 0.764$, $q = 0.086$) and Information Source and Trust Discussion ($\beta = 0.871$, $q = 0.009$). This dual association reveals that agricultural vulnerability operates as a multifaceted contextual factor that co-occurs with multiple PADM stages simultaneously.

On one hand, the association with Direct Situation Report reflects a vulnerability-driven monitoring response. As illustrated by the SHAP-GIS map (Fig. 14), this relationship is most pronounced in major agricultural provinces such as Henan and Shandong. This aligns with social vulnerability literature, which identifies economic dependence on climate-sensitive sectors as a key factor in disaster risk (Cutter et al., 2003; Minea et al., 2025). For these populations, the typhoon is not an abstract event but a direct and immediate threat to their livelihood, which co-occurs with a higher volume of real-time observation and reporting of environmental conditions (PADM Stage I: Information Exposure).

On the other hand, the even stronger association with Information Source and Trust Discussion ($\beta = 0.871$, $q = 0.009$—the single most significant result across all 81 tests) suggests that agricultural vulnerability is associated not only with immediate environmental monitoring but also with a heightened demand for reliable disaster information. In provinces where farming livelihoods depend directly on accurate weather forecasting, the public tends to critically scrutinize whether official information sources can be trusted (PADM Stage III: Stakeholder Appraisal). This is consistent with RPT literature suggesting that populations with high material stakes in disaster outcomes develop a stronger demand for information accuracy (Balog-Way et al., 2020; Siegrist & Árvai, 2020).

Second, Information Source and Trust Discussion was also independently associated with Natural Disaster Economic Loss per capita ($\beta = 0.636$, $q = 0.086$), suggesting a complementary pattern of experience-based sensitization. This relates to the "trust" component of risk discourse (Siegrist & Cvetkovich, 2000). Populations in provinces that have previously suffered greater per capita economic losses from disasters (as shown in Fig. 16, where this factor had high feature importance in Hebei and Beijing) are likely more sensitized to the importance of accurate information, which co-occurs with higher scrutiny of official forecasts and media reports (Mase et al., 2015). This finding complements the cropland association: while agricultural exposure captures current structural vulnerability, disaster economic loss captures accumulated experiential vulnerability—and both independently predict heightened critical discourse on information reliability.

Collectively, these findings demonstrate that the public's online response to a disaster is not random. It is a structured phenomenon, predictable at the macro level by underlying social and economic realities. Agricultural vulnerability (Cropland Cover Proportion) emerged as the most pervasive spatial correlate—associated with both situational reporting and critical information scrutiny—while accumulated disaster experience (Natural Disaster Economic Loss per capita) independently reinforced the demand for information reliability.

5.3. Cross-platform implications: algorithmically curated discourse in collectivist contexts

A critical question arising from Section 5.1 is why socio-cultural coping mechanisms achieved near-parity with functional information exchange on Douyin, when studies on text-centric platforms consistently report information dominance exceeding 50% (Dong & Xie, 2024; Zhang & Cozma, 2022). We propose that Douyin's algorithmic architecture provides a structural explanation. Unlike chronological or follower-based feeds, Douyin's recommendation algorithm prioritizes content engagement—measured through watch time, shares, and comment activity—over informational utility or source authority. Humor, regional memes, and culturally resonant narratives are inherently high-engagement content formats that the algorithm is likely to amplify disproportionately, thereby elevating socio-cultural coping discourse to a visibility level comparable to factual reporting. This suggests that the Information-Coping Equilibrium observed in Section 5.1 may not solely reflect the public's communicative preferences but is also co-produced by the platform's algorithmic mediation, a finding consistent with SARF's recognition that "amplification stations" are not neutral conduits but actively shape the signal landscape (Kasperson et al., 1988).

Similarly, the dual role of Cropland Cover Proportion identified in Section 5.2—predicting both situational reporting and information trust scrutiny—may be partially platform-specific. On Douyin, the short-video format lowers the barrier for first-person environmental reporting: a farmer filming flooded fields requires no written literacy or journalistic framing, making agricultural provinces natural hubs of "ground-truth" content production. This same video-centric accessibility may also stimulate trust-related discourse, as users in agricultural regions can visually compare their own observations with official forecasts, creating a feedback loop between experiential evidence and institutional credibility assessment. Whether this dual association would replicate on text-only platforms, where situational reporting requires more deliberate textual composition, remains an open empirical question.

These platform-specific mechanisms caution against uncritical

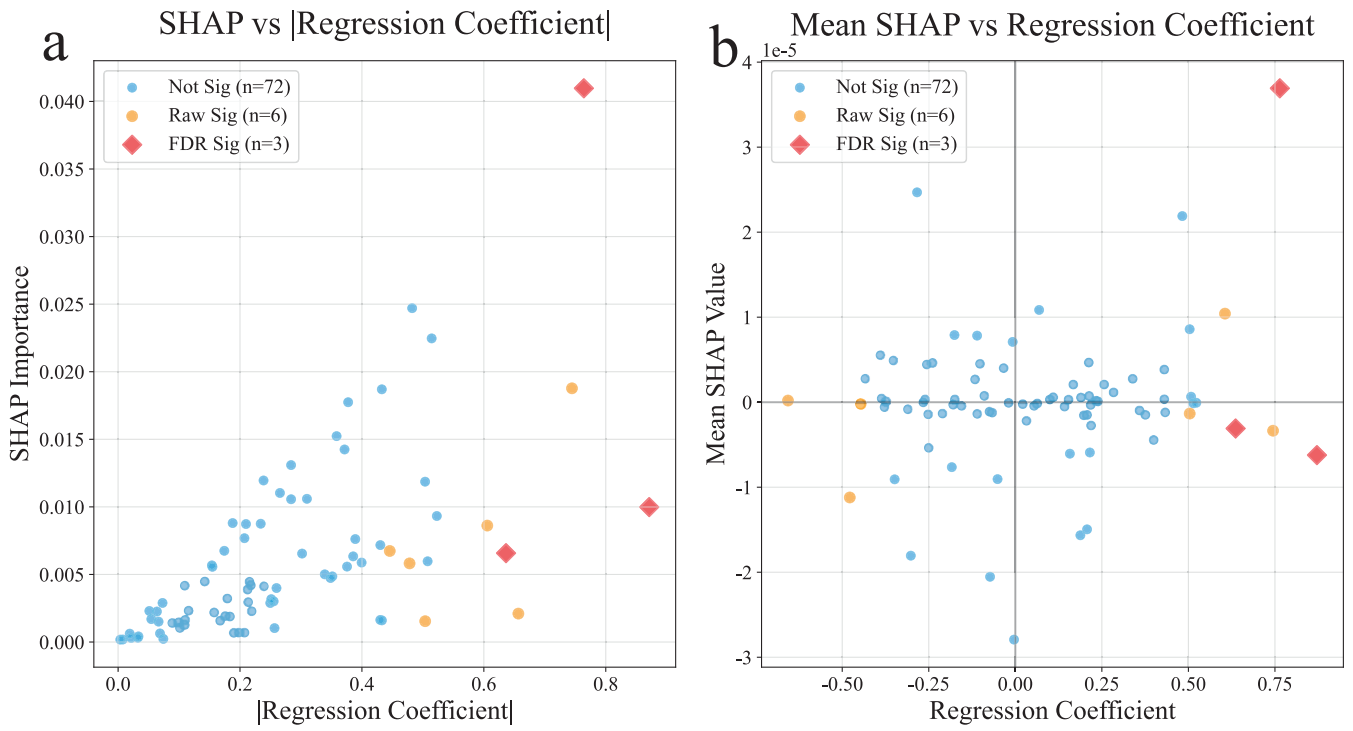

Fig. 13. SHAP validation plots.

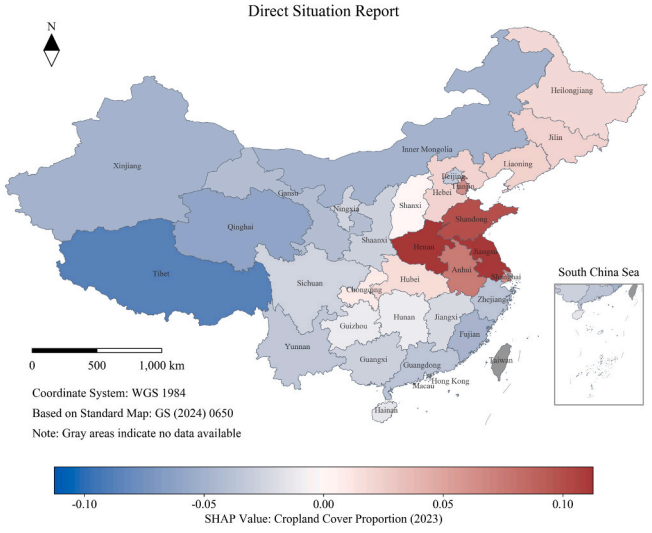

Fig. 14. SHAP-GIS Map: contribution of 'cropland cover proportion' to "direct situation report".

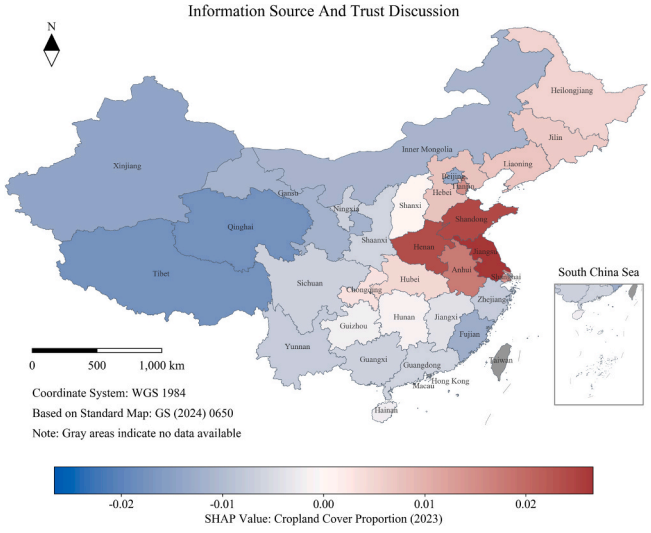

Fig. 15. SHAP-GIS Map: contribution of 'cropland cover proportion' to "information source and trust discussion".

generalization of our findings to other social media ecosystems. While the methodological pipeline (LLM classification → spatial regression → SHAP-GIS) is designed for cross-platform transferability, the substantive findings—particularly the Information-Coping Equilibrium and the elevated role of visual situational reporting—are likely moderated by platform architecture. On text-centric platforms like Weibo or Twitter, where algorithmic curation plays a smaller role and textual fluency is a prerequisite for participation, we would expect information-seeking themes to dominate and the cultural coping component to be comparatively attenuated. Conversely, on other algorithm-driven short-video platforms such as TikTok, YouTube Shorts, or Xiaohongshu, similar equilibria between information and coping discourse may emerge, particularly during climate-related events in collectivist cultural contexts. Future cross-platform comparative studies are needed to empirically test these hypothesized boundary conditions.

5.4. *Theoretical Implications*

5.4.1. *Challenging the "information-seeking" hegemony in RPT*

Traditional Risk Perception Theory (RPT) often views social media use during disasters through a utilitarian lens, prioritizing functional information-seeking (Siegrist & Árvai, 2020; Slovic, 1987). Our finding challenges this paradigm by demonstrating a near-perfect equilibrium between Functional Information Exchange (33.89%) and Socio-cultural Coping (32.98%). This suggests that in the era of algorithmically-curated short videos, risk perception is not merely a cognitive calculation of threat but a dual-track process of information processing and collective meaning-making (Liu et al., 2025; Wilkinson, 2001). We

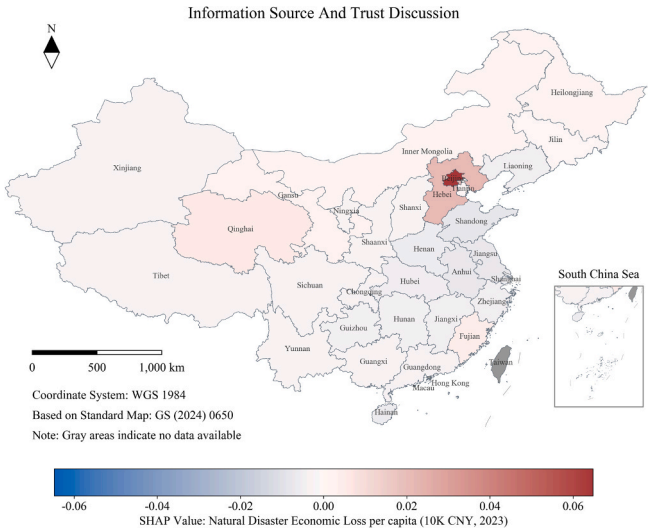

Fig. 16. SHAP-GIS Map: contribution of 'natural disaster economic loss per capita' to "information source and trust discussion".

empirically enrich RPT by showing that affective responses (humor) and cultural identity are not "noise", but core components of how the public sustains psychological resilience in digital environments.

5.4.2. *Spatially structured amplification in SARF*

SARF identifies "amplification stations (Kasperson et al., 2003)," yet it rarely explains why these stations behave differently across geographic space. By linking SARF with spatial econometrics, this study provides a "spatialized" validation of SARF. We demonstrate that amplification is not random but systematically structured by macro-level societal characteristics (Kasperson et al., 2022; Wirz et al., 2018). For instance, the positive association between Cropland Cover Proportion and 'Information Source and Trust Discussion' ($\beta = 0.871$, $q = 0.009$) illustrates how provincial socio-economic contexts—specifically agricultural vulnerability—act as 'pre-decisional filters' that calibrate how risk signals are amplified or attenuated within a regional digital ecosystem.

5.4.3. *Integrating PADM: discourse as a proxy for cognitive stages*

Responding to the need for a process-oriented understanding, this study makes a significant contribution by mapping digital discourse categories onto the Protective Action Decision Model (PADM) (Heath et al., 2017). We transition from a static thematic analysis to a process-based view, where categories like " Typhoon Path and Regional Speculation " and " Risk Perception and Emotion (Humor or Sarcasm) " are conceptualized as digital traces of Stage I (Pre-decisional Information Processes) and Stage IV (Socio-cultural Coping) of the PADM. This integration allows risk theory to move beyond "what is said" to "where the public stands" in the decision-making cycle, positioning digital discourse as a real-time observation window into the pre-decisional cognitive processes of a population.

5.5. *Methodological implications and transferability*

The success of this study hinges on a dual-layered methodological innovation: the high-fidelity classification via an advanced LLM and the explainable spatial grounding provided by the SHAP-GIS pipeline.

First, the LLM-Augmented Thematic Sensing: Our findings confirm that the *DeepSeek R1 671B* model, utilizing a few-shot learning approach, effectively addresses the challenges of scale and linguistic complexity inherent in disaster social media data (Rutsaert et al., 2013; Schroeder et al., 2013). The validation process yielded 88.3% overall accuracy, demonstrating robust performance and high transferability. Traditional NLP methods, such as dictionary-based approaches or simpler topic models, often struggle with the informal, context-dependent, and ironic language prevalent in online discourse (Minaee et al., 2021). This human-supervised approach, which effectively captures context-dependent nuances like humor and regional metaphors, can be efficiently adapted and deployed to other risk communication scenarios with minimal retraining requirements.

Second, SHAP-GIS Explainable Pipeline: The most significant methodological contribution lies in the integration of SHAP with GIS. While traditional spatial regressions identify global trends, the SHAP-GIS workflow provides local interpretability,[2] allowing researchers to "see" how specific provincial socio-economic profiles contribute to the model's predictions.

Third, Practical Replication for Emergency Management: To facilitate the transferability of this framework, we demonstrate that this pipeline can be replicated using widely accessible open-source tools. Emergency managers can implement this "sensing-to-explanation" workflow by: Sensing: Utilizing the LLMs API for rapid, low-cost thematic classification of social media streams. Mapping: Employing geopandas for spatial aggregation of discourse categories. Explaining: Applying the shap Python library to identify which regional indicators (e.g., agricultural exposure, disaster loss) are the strongest correlates of specific public concerns.

By providing the underlying logic and Python-based implementation (see Supplementary Material), this study offers an adaptable framework whose core pipeline (LLM classification → spatial regression → SHAP-GIS decomposition) can be systematically replicated with platform-specific codebook adjustments. In principle, the framework is adaptable to other geo-tagged social media platforms (e.g., Weibo, TikTok, YouTube Shorts) and potentially to other text-based data sources with geographic metadata, though platform-specific adaptations to the classification codebook would be required.

5.6. *Practical Implications*

The findings offer several actionable insights for governmental agencies, emergency managers, and risk communicators: **Move Beyond Informational Broadcasting:** Simply disseminating information is insufficient for public engagement, although it is necessary. Our data shows that the public spends just as much energy on cultural and emotional processing. Effective risk communication must also engage with these "heart-and-mind" issues, not just the "what-and-where" of the hazard. **Acknowledge and Leverage Local Discourse:** The high prevalence of *Cultural or Regional Identity'* and its unique spatial patterns (e.g., in *Shanxi*) is a critical insight. Risk communicators should move away from one-size-fits-all messaging and, where possible, tailor communications to resonate with local identities, historical narratives, and cultural values to build trust and increase message salience. **Use Socio-Economic Contexts for Targeted Communication:** Our regression findings provide a basis for data-driven communication strategies. For example: In regions with high agricultural vulnerability (high cropland cover), the public is simultaneously focused on real-time environmental 'ground truth' and actively scrutinizing the reliability of official information. Communication strategies for these regions should therefore prioritize both localized, real-time impact information and transparent, evidence-based sourcing to meet this dual demand for accuracy and accountability. In regions with high prior disaster economic losses, expect a more critical and information-scrutinizing public. Communication here must be highly transparent and explicitly address trust, acknowledging the population's accumulated experience with disaster

[2] Note for Interpretation: It must be emphasized that SHAP values reflect the marginal contribution to model prediction rather than a direct causal effect. These values act as a diagnostic tool for sensing spatial associations, characterizing how regional vulnerabilities align with digital discourse.

impacts. **Monitor Humor as a Barometer:** The high prevalence of *Humor or Sarcasm* should not be dismissed as frivolous. It serves as a key public coping mechanism (Dong & Xie, 2024). Monitoring the *type* of humor can provide valuable, real-time insights into underlying public anxieties, skepticism, or areas where official messaging is failing to resonate.

5.7. Limitations and future research

While this study provides significant insights, its limitations offer clear pathways for future inquiry. Our findings are specific to one event (Typhoon Gaemi) on one platform (*Douyin*). The unique communication norms of this short-video platform and the cultural context of this event mean the results may not be generalizable to all disasters or other social media ecosystems (e.g., *Weibo*). Therefore, future research should conduct comparative studies across different platforms, disaster types, and cultural contexts to test the generalizability of these thematic patterns.

Additionally, while our comment-level sample (N = 436,964) provides robust statistical power, the engagement-based selection of 191 videos prioritizes high-visibility content, which may over-represent mainstream disaster narratives while under-sampling less popular or alternative perspectives. This sampling strategy means our findings primarily characterize dominant patterns in public disaster discourse rather than the full diversity of risk framings present on the platform. Furthermore, our data cleaning pipeline relied on content-level heuristics (length filtering and noise symbol removal) rather than account-level bot detection methods (e.g., account age or posting frequency), as the Douyin public comment interface does not expose user-level metadata such as account creation date or historical posting activity to non-API data collection methods. While our filtering effectively removed non-substantive content, future research with direct platform API access could implement account-level screening to more rigorously exclude automated or bot-driven comments and report the percentage of content removed.

This study analyzes the typhoon-related discourse as a cumulative cross-sectional dataset. Consequently, it does not distinguish between the different temporal phases of the disaster (pre-event vs. post-event) or the varying levels of direct hazard experience among individual commenters. While the use of provincial-level geographic aggregation serves as a robust contextual proxy for regional vulnerability, it inherently masks the heterogeneity of individual-level responses within each province. Future studies should consider temporal segmentation and user-level metadata to investigate how risk discourse evolves as the threat progresses from speculative to immediate impact.

While our LLM methodology proved highly effective (88.3% accuracy), the thematic validation of the LLM classification was primarily performed by the lead author. To ensure the robustness and temporal stability of the coding logic, the lead author conducted a blind second-round validation on the 1,000-comment sample after a three-month interval, achieving an intra-rater reliability of 88.3%. While this longitudinal consistency mitigates concerns regarding individual caprice, the absence of independent coders remains a limitation. Future research should involve multiple independent participants to conduct cross-verification and report inter-annotator agreement metrics (e.g., Cohen's Kappa) to further eliminate potential subjective bias.

We acknowledge the inherent assumption that IP-derived locations perfectly capture lived experience. A potential robustness check would involve restricting the analysis to comments containing explicit geographic references in their text. However, on Douyin, most comments are brief (median length approximately 15 characters) and rarely include explicit place names, as users' geographic identity is already publicly displayed via their IP-derived provincial tag. Restricting the location-mentioned comments would drastically reduce the sample size and introduce a selection bias favoring geographically oriented categories (e.g., Path Speculation, Direct Situation Reports) while systematically excluding affect-driven discourse (e.g., Humor, Worry/Fear). We therefore maintain that IP-derived provincial aggregation provides a more representative geographic proxy for this specific platform context. While micro-level variations exist, IP data serves as a contextual proxy for regional exposure. Furthermore, our use of 2020–2023 socio-economic data assumes relative stability in provincial structures; future work should prioritize real-time individual-level controls (e.g., age, gender) via survey triangulation where API access permits. Relatedly, the single-event cross-sectional design precludes the use of lagged socio-economic variables or quasi-experimental approaches such as difference-in-differences analysis. Future multi-event studies—tracking discourse across successive disaster seasons—could employ such designs to more rigorously test the temporal stability and directionality of the identified associations.

Finally, our statistical analysis, while robust in identifying spatial and correlational patterns, does not establish causation. The use of aggregate provincial-level data means ecological fallacy must be avoided; we can only claim associations at the regional level, not for individuals. Specifically, the cross-sectional ecological design cannot rule out reverse causation (e.g., regions with more active online discourse may attract greater institutional investment in disaster preparedness) or omitted-variable bias (e.g., unmeasured factors such as local media ecology, provincial government communication strategies, or individual-level demographic characteristics may confound the observed associations). These inferential constraints are inherent to the study design and underscore the associational—rather than causal—nature of our findings. Future studies should aim to bridge this micro–macro gap, perhaps by integrating our macro-level findings with individual-level surveys or qualitative interviews to explore the causal mechanisms linking regional socio-economic characteristics, agricultural vulnerability, and risk discourse patterns.

6. Conclusion

This study conducted a large-scale, mixed-methods analysis of public risk discourse on Douyin during Typhoon Gaemi. We moved beyond traditional approaches by developing and validating a sophisticated LLM-based methodology to classify over 436,964 public comments into ten distinct thematic categories.

Our first key finding is that public discourse during a disaster is a dual process of information-seeking and socio-cultural sense-making. We demonstrated that themes of *Cultural or Regional Identity and Narrative* and *Risk Perception and Emotion (Humor or Sarcasm)* were just as prevalent as *Typhoon Path and Regional Speculation* and *Direct Situation Report*. This empirically confirms that on a platform like *Douyin*, the social and emotional processing of risk is as fundamental as the consumption of information, providing a clear illustration of the Social Amplification of Risk Framework (SARF) in a contemporary digital context.

Our second key finding is that this public discourse is not random but is systematically associated with regional socio-economic contexts. Using spatial Beta Regression and an innovative SHAP-GIS interpretation method, we identified robust, significant relationships that survived strict FDR correction. Most notably, Cropland Cover Proportion emerged as the single most pervasive correlation, significantly predicting two distinct discourse categories: it was associated with both intensified Direct Situation Reports ($\beta = 0.764$, $q = 0.086$) and heightened scrutiny of Information Source and Trust ($\beta = 0.871$, $q = 0.009$). This reveals that agricultural vulnerability operates as a dual contextual stimulus—provinces where farming livelihoods are directly threatened by typhoon hazards exhibit not only greater real-time environmental reporting but also a stronger demand for information reliability, reflecting the high economic stakes attached to forecast accuracy in these communities. Furthermore, Natural Disaster Economic Loss per capita was independently associated with Information Source and Trust Discussion ($\beta = 0.636$, $q = 0.086$), indicating that accumulated disaster

experience compounds this critical orientation toward information sources. Together, these findings suggest that the regional disaster discourse landscape is *systematically associated with* the intersection of current economic exposure (agricultural vulnerability) and historical disaster experience (accumulated economic losses), both of which are associated with more active and critical public engagement with disaster-related information.

This research makes a threefold contribution: Theoretically, it provides robust, large-scale evidence for RPT and SARF in the emerging but rapidly growing *Douyin* ecosystem. Methodologically, it introduces a transparent, validated workflow combining an advanced LLM with a novel SHAP-GIS approach to link complex models with interpretable spatial insights. Practically, it provides evidence-based guidance for risk communicators, demonstrating the need to move beyond simple information-sharing and to tailor messages to the specific, measurable socio-economic and cultural contexts of different regions. Ultimately, this study underscores that understanding contemporary disaster response requires listening to and interpreting the public's rich, diverse, and socially embedded online voices.

Our findings provide a roadmap for spatially targeted risk communication. Based on the identified spatial correlations, we offer the following actionable recommendations: In agricultural provinces with high Cropland Cover Proportion, public discourse simultaneously reflects both active situational reporting and intense scrutiny of information reliability. Emergency managers in these regions should adopt a dual strategy: (a) harness the population's propensity for real-time environmental reporting by establishing streamlined citizen-reporting channels to supplement official monitoring, treating these regions as reliable 'ground-truth' sensory networks; and (b) proactively address information trust by providing transparent, timely, and locally relevant forecast updates, recognizing that forecast accuracy in these regions carries direct consequences for agricultural livelihoods and economic survival. In provinces with high historical Natural Disaster Economic Loss per capita, the public exhibits heightened scrutiny of information sources and trust. Communication strategies for these regions should prioritize transparency and evidence-based messaging, explicitly acknowledge prior disaster experiences and demonstrate institutional learning from past events to rebuild and maintain public trust.

While this study is grounded in a specific typhoon event on Douyin, the validated LLM-few-shot workflow and the SHAP-GIS pipeline demonstrate significant transferability. These methodologies are well-suited for analyzing public response on other algorithm-driven, short-video platforms such as TikTok, YouTube Shorts, or Xiaohongshu during diverse climate-induced crises, including floods, wildfires, or heatwaves. By balancing socio-cultural sense-making with functional information exchanging, our framework offers a scalable tool for emergency managers to sense and navigate the spatially structured landscape of digital risk discourse in a rapidly evolving social media era.

7. Data ethics and compliance statement

This study utilizes publicly accessible user comments from Douyin. In accordance with ethical guidelines for internet-mediated research: All analyzed data is in the public domain; IP-based location data is automatically displayed by the platform at the provincial level. A custom script was used to strip all PII (Personally Identifiable Information), including usernames and account tags, before LLM processing. No individual-level tracking was performed. All socio-economic correlations were conducted at the provincial aggregate level to prevent re-identification.

CRediT authorship contribution statement

Pu ZHANG: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The author declares that he have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The author would like to express sincere gratitude to the Red Bird MPhil Program for its support throughout this research. This paper is based on the author's MPhil thesis completed at The Hong Kong University of Science and Technology (Guangzhou). The authors deeply appreciate the members of the thesis examination committee, namely Prof. Corey Kewei Xu, Prof. Jing Tang, Prof. Muzhi Zhou, and Dr. Arthur Kar Leung LIN, for their valuable guidance, rigorous review, and insightful suggestions during the thesis examination process. The author also thanks The Hong Kong University of Science and Technology (Guangzhou) for providing essential API computing resources that made this work possible.

Appendix 1. Thematic classification: development, decision rules, and validation

A1.1. Thematic development and saturation process

The thematic codebook was developed through an iterative, incremental sampling process rooted in the lead author's MPhil research. To ensure both theoretical robustness and empirical sensitivity to the Douyin ecosystem, the following stages were executed:

Stage 1: Pilot Coding (n = 1,000). A simple random subsample of 1,000 comments was drawn from the full corpus (N = 436,964). Stratification was not applied during sampling to preserve ecological validity. Given the highly skewed natural distribution of themes in disaster discourse, stratified sampling would artificially inflate the prevalence of rare categories, potentially overestimating classification accuracy under real-world deployment conditions. Simple random sampling provides a more conservative accuracy estimate that reflects the natural base rates of the Douyin ecosystem. Initial open coding identified preliminary clusters focusing on information-seeking and basic emotional expression (Worry/Fear).

Stage 2: Incremental **Expansion (n = 1,000 + 2,000 = 3,000).** The sample was expanded to 3,000 comments. During this phase, platform-specific discourses emerged, such as regional “rain god” metaphors and collective humor, necessitating the creation of the *Socio-cultural Coping* and *Cultural Identit* categories.

Stage 3: Saturation Verification (n = 1,000 + 2,000 + 2,000 = 5,000). With a total sample of 5,000 comments, theoretical saturation was reached. No new thematic categories with a prevalence exceeding 0.5% emerged in the final 2,000 comments. The final ten-category taxonomy was then mapped onto the Protective Action Decision Model (PADM) stages to ensure functional alignment with disaster psychology.

A1.2. Validation protocol: intra-coder reliability

Given the sole-authored nature of this study, an intra-coder reliability (stability) test was conducted to assess the reproducibility of the classification framework over time. Rather than measuring agreement between two independent coders (inter-rater reliability), this longitudinal design measures the stability of the same coder's judgments over a prolonged interval, thereby isolating the clarity and reproducibility of the decision rules from short-term memory effects.

Sampling for Validation. A validation set of 1,000 comments was randomly sampled from the full corpus, independent of the thematic development sample. Simple random sampling was again employed to preserve ecological validity.

Coding Procedure. The coder first familiarized themselves with the ten thematic categories, their operational definitions, and representative examples from the codebook (i.e., the same system prompt subsequently provided to the LLM). For each comment, the coder read the full text and assigned the single most appropriate category according to the Dominant-Intent Priority Hierarchy (see Section A1.3). Edge cases were logged with brief notes for consistency checking.

Longitudinal Design. The two coding sessions were separated by a 10-month interval (T1: May 2025; T2: March 2026). The temporal metadata of the research repository confirms rigorous tracking using identical comment IDs. This extended gap ensures that the second coding session was conducted without any reliance on short-term memory, strictly following the same Dominant-Intent Priority Hierarchy.

Reliability Results. Consistency between the two longitudinal sessions was evaluated using Cohen's Kappa (Cohen, 1960). The analysis revealed a kappa value of 0.89, indicating strong agreement under the specific coding scheme and test conditions (Landis & Koch, 1977). Of the 1,000 samples, 97.9% remained consistent, with minor discrepancies (2.1%) primarily occurring in early-stage labels that were subsequently corrected by stricter adherence to the Fact-First and Information Scrutiny rules in the final iteration. This high level of longitudinal stability confirms that the decision rules are reproducible and not dependent on contextual fluctuations or short-term memory. It further corroborates the reliability of the 88.3% LLM classification accuracy reported in the main manuscript.

A1.3. Dominant-Intent Priority Hierarchy (Decision Rules)

To address the challenge of semantic overlap (a single comment may contain multiple intents), we established a Dominant-Intent Priority Hierarchy. This protocol ensures statistical consistency for the spatial beta regression, which requires mutually exclusive proportional data. When a comment presented hybrid characteristics, the classification followed these priority rules:

1. **Fact-First Rule.** If a comment contained specific, first-person observations of environmental conditions (e.g., wind intensity, rainfall, water levels), it was prioritized as a *Direct Situation Report*, regardless of any accompanying emotional tone or humorous framing.
2. **Affective Dominance Rule.** If a comment used highly stylized language, regional memes, sarcasm, or culturally specific metaphors to frame the risk, it was categorized under *Humor or Sarcasm* or *Cultural Identity*, as these interpretive signals represent the primary socio-cultural coping mechanism rather than literal information exchange.
3. **Action Over Sentiment Rule.** Comments expressing empathy alongside specific well-wishes, safety reminders, or offers of assistance were classified as *Mutual Aid, Wishes, and Concern*, reflecting the "Social Efficacy" stage of the PADM. This prioritizes actionable social support over passive emotional expression.
4. **Information Scrutiny Rule.** Any discussion questioning the reliability of forecasts, official media reports, or information sources was prioritized as *Information Source and Trust Discussion*, given its critical role in the "Stakeholder Evaluation" stage of the PADM. This rule overrides neutral information-seeking classifications.

A1.4. Zero-Shot ablation experiment and classification challenges

To evaluate the marginal utility of the few-shot prompting strategy, a zero-shot ablation experiment was conducted using the same model (DeepSeek R1 671B). The zero-shot configuration relied solely on the category definitions without representative examples. The specific results are shown in Table A1 below. The comparison between Zero-Shot (82.1%) and Few-Shot (88.3%) performance reveals several critical insights into the nature of disaster-related discourse on short-video platforms:

Table A1
LLM Classification Validation Performance (Zero-Shot Baseline, N = 1,000).

Category	Samples in Set	True	Accuracy Rate (%)
Cultural or Regional Identity and Narrative	81	76	93.83
Direct Situation Report	152	142	93.42
Information Source and Trust Discussion	37	31	83.78
Mutual Aid, Wishes, and Concern	34	25	73.53
Risk Attenuation and Dismissal	100	62	62.00
Risk Perception and Emotion (Humor or Sarcasm)	192	143	74.48
Risk Perception and Emotion (Worry or Fear)	30	28	93.33
Social and Life Impact	116	99	85.34
Typhoon Path and Regional Speculation	193	169	87.56
Other	65	46	70.77
Total	1,000	821	82.10

Precision-Recall Trade-off in Subjective Categories: In the Zero-shot setting, categories like *Cultural or Regional Identity* showed high accuracy but low recall (identifying only 81 samples compared to 171 in the few-shot set). This suggests that without examples, the model adopted a "conservative" bias, only labeling comments that met a narrow, literal interpretation of the category. The few-shot approach "calibrated" the model to recognize broader, more implicit cultural cues and regional dialects.

The Challenge of Risk Attenuation: The significant performance gap in *Risk Attenuation and Dismissal* (62.0% in Zero-shot vs. 77.6% in Few-shot) underscores the difficulty of identifying "normalization" discourse (e.g., "just another rainy day"). Without contextual examples, the model frequently confused risk dismissal with neutral reports.

Semantic Overlap and Multi-intent Discourse: A qualitative review of misclassified instances indicates that many comments are inherently **multi-intent**. For example, a comment stating *"Stay safe, Fuzhou! Hopefully Mazu will keep the wind away"* simultaneously contains elements of **Mutual Aid**, **Cultural Identity**, and **Worry**.

In this study, to ensure statistical consistency for the spatial Beta Regression (where dependent variables must sum to a coherent proportion for each province) we implemented a **single-label, dominant-intent classification strategy**. This required the model to prioritize the primary discourse function. While this methodology effectively captures macro-level spatial patterns, it inevitably introduces a degree of "taxonomic overlap" where secondary intents are suppressed. Future work will explore **multi-label architectures** to more fully preserve the semantic richness of multi-intent disaster discourse.

Appendix B. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.ssci.2026.107281.

Data availability

The datasets generated and analyzed during the current study are not publicly available due to privacy concerns related to social media user data. The core code and the LLM prompts used for the analysis will be available in Supplementary materials.

References

Alexander, D.E., 2013. Social media in disaster risk reduction and crisis management. Sci. Eng. Ethics 20 (3), 717–733. https://doi.org/10.1007/s11948-013-9502-z.

AlQahtany, A.M., Abubakar, I.R., 2020. Public perception and attitudes to disaster risks in a coastal metropolis of Saudi Arabia. Int. J. Disaster Risk Reduct. 44, 101422. https://doi.org/10.1016/j.ijdrr.2019.101422.

Artstein, R., Poesio, M., 2008. Inter-coder agreement for computational linguistics. Comput. Linguist. 34 (4), 555–596. https://doi.org/10.1162/coli.07-034-R2.

Askari, E., Savari, M., Rezaei, M., 2025. Factors affecting the application of protective measures before flood occurrence among local communities in Iran. Clim. Risk Manag. 49, 100730. https://doi.org/10.1016/j.crm.2025.100730.

Balog-Way, D., McComas, K., Besley, J., 2020. The evolving field of risk communication. Risk Anal. 40 (S1), 2240–2262. https://doi.org/10.1111/risa.13615.

Cacioppo, J.T., Hawkley, L.C., 2009. Perceived social isolation and cognition. Trends Cogn. Sci. 13 (10), 447–454. https://doi.org/10.1016/j.tics.2009.06.005.

Cao, G., Shen, L., Evans, R., Zhang, Z., Bi, Q., Huang, W., Yao, R., Zhang, W., 2021. Analysis of social media data for public emotion on the Wuhan lockdown event during the COVID-19 pandemic. Comput. Methods Programs Biomed. 212, 106468. https://doi.org/10.1016/j.cmpb.2021.106468.

Chen, N.-C., Drouhard, M., Kocielnik, R., Suh, J., Aragon, C.R., 2018. Using machine learning to support qualitative coding in social science: shifting the focus to ambiguity. ACM Trans. Interact. Intell. Syst. 8 (2), 9. https://doi.org/10.1145/3185515.

Cohen, J., 1960. A coefficient of agreement for nominal scales. Educ. Psychol. Meas. 20 (1), 37–46. https://doi.org/10.1177/001316446002000104.

Cribari-Neto, F., Zeileis, A., 2010. Beta Regression in R. J. Stat. Softw. 34 (2), 1–24. https://doi.org/10.18637/jss.v034.i02.

Cutter, S.L., Boruff, B.J., Shirley, W.L., 2003. Social vulnerability to environmental hazards*. Soc. Sci. Q. 84 (2), 242–261. https://doi.org/10.1111/1540-6237.8402002.

Dong, Z., Xie, T., 2024. Why do people love short-form videos? The motivations for using Chinese TikTok (Douyin) and implications for well-being. Curr. Psychol. 43 (26), 22283–22296. https://doi.org/10.1007/s12144-024-05927-4.

Dong, B., Xia, J., Li, Q., Zhou, M., 2022. Risk assessment for people and vehicles in an extreme urban flood: case study of the "7.20" flood event in Zhengzhou China. Int. J. Disaster Risk Red. 80, 103205. https://doi.org/10.1016/j.ijdrr.2022.103205.

Dong, X., Ye, L., Lian, Y., 2024. Who has dominated information spreading on social media during the early stage of COVID-19 pandemic in China? A temporal network analysis. Int. J. Disaster Risk Reduct. 107, 104493. https://doi.org/10.1016/j.ijdrr.2024.104493.

Fan, Y., 2020. *The Social Amplification of risk on Weibo-take the COVID-19 Epidemic in China as an example*. In: Proceedings of the 2020 Conference on Artificial Intelligence and Healthcare. https://doi.org/10.1145/3433996.3434013.

Fauzi, M.A., 2023. Social media in disaster management: review of the literature and future trends through bibliometric analysis. Nat. Hazards 118 (2), 953–975. https://doi.org/10.1007/s11069-023-06079-7.

Fereday, J., Muir-Cochrane, E., 2006. Demonstrating rigor using thematic analysis: a hybrid approach of inductive and deductive coding and theme development. Int. J. Qual. Methods 5 (1), 80–92. https://doi.org/10.1177/160940690600500107.

Finucane, M.L., Alhakami, A., Slovic, P., Johnson, S.M., 2000. The affect heuristic in judgments of risks and benefits. J. Behav. Decis. Mak. 13 (1), 1–17. https://doi.org/10.1002/(SICI)1099-0771(200001/03)13:1<1::AID-BDM333>3.0.CO;2-S.

Gilardi, F., Alizadeh, M., Kubli, M., 2023. ChatGPT outperforms crowd workers for text-annotation tasks. Proc. Natl. Acad. Sci. 120 (30), e2305016120. https://doi.org/10.1073/pnas.2305016120.

Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z.F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Zhang, Z., 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. Nature 645 (8081), 633–638. https://doi.org/10.1038/s41586-025-09422-z.

Heath, R.L., Lee, J., Palenchar, M.J., Lemon, L.L., 2017. Risk communication emergency response preparedness: contextual assessment of the protective action decision model. Risk Anal. 38 (2), 333–344. https://doi.org/10.1111/risa.12845.

Houston, J.B., Hawthorne, J., Perreault, M.F., Park, E.H., Goldstein Hode, M., Halliwell, M.R., Turner McGowen, S.E., Davis, R., Vaid, S., McElderry, J.A., Griffith, S.A., 2014. Social media and disasters: a functional framework for social media use in disaster planning, response, and research. Disasters 39 (1), 1–22. https://doi.org/10.1111/disa.12092.

Hripcsak, G., 2005. Agreement, the F-measure, and reliability in information retrieval. J. Am. Med. Inform. Assoc. 12 (3), 296–298. https://doi.org/10.1197/jamia.M1733.

Ilyas, B., Sharifi, A., 2025. A systematic review of social media-based sentiment analysis in disaster risk management. Int. J. Disaster Risk Reduct. 123, 105487. https://doi.org/10.1016/j.ijdrr.2025.105487.

Jenkins, S.C., Lachlan, R.F., Osman, M., 2024. An integrative framework for mapping the psychological landscape of risk perception. Sci. Rep. 14 (1). https://doi.org/10.1038/s41598-024-59189-y.

Kasperson, R.E., Renn, O., Slovic, P., Brown, H.S., Emel, J., Goble, R., Kasperson, J.X., Ratick, S., 1988. The social amplification of risk: a conceptual framework. Risk Anal. 8 (2), 177–187. https://doi.org/10.1111/j.1539-6924.1988.tb01168.x.

Kasperson, J.X., Kasperson, R.E., Pidgeon, N., Slovic, P., 2003. The social amplification of risk: assessing fifteen years of research and theory. In: Pidgeon, N., Kasperson, R.E., Slovic, P. (Eds.), *The Social Amplification of Risk*. Cambridge University Press, pp. 13–46. https://doi.org/10.1017/CBO9780511550461.002.

Kasperson, R.E., Webler, T., Ram, B., Sutton, J., 2022. The social amplification of risk framework: New perspectives. Risk Anal. 42 (7), 1367–1380. https://doi.org/10.1111/risa.13926.

Kaye, D.B.V., Chen, X., Zeng, J., 2021. The co-evolution of two Chinese mobile short video apps: parallel platformization of Douyin and TikTok. Mobile Media & Communication 9 (2), 229–253. https://doi.org/10.1177/2050157920952120.

Landis, J.R., Koch, G.G., 1977. The measurement of observer agreement for categorical data. Biometrics 33 (1), 159. https://doi.org/10.2307/2529310.

Lange, S., Volkholz, J., Geiger, T., Zhao, F., Vega, I., Veldkamp, T., Reyer, C.P.O., Warszawski, L., Huber, V., Jägermeyr, J., Schewe, J., Bresch, D.N., Büchner, M., Chang, J., Ciais, P., Dury, M., Emanuel, K., Folberth, C., Gerten, D., Frieler, K., 2020. Projecting exposure to extreme climate impact events across six event categories and three spatial scales. Earth's Future 8 (12), e2020EF001616. https://doi.org/10.1029/2020EF001616.

Lechowska, E., 2021. Approaches in research on flood risk perception and their importance in flood risk management: a review. Nat. Hazards 111 (3), 2343–2378. https://doi.org/10.1007/s11069-021-05140-7.

Li, J., Huang, C., Yang, Y., Liu, J., Lin, X., Pan, J., 2023. How nursing students' risk perception affected their professional commitment during the COVID-19 pandemic: the mediating effects of negative emotions and moderating effects of psychological capital. Human. Soc. Sci. Commun. 10 (1). https://doi.org/10.1057/s41599-023-01719-6.

Lindell, M.K., Perry, R.W., 2011. The protective action decision model: theoretical modifications and additional evidence. Risk Anal. 32 (4), 616–632. https://doi.org/10.1111/j.1539-6924.2011.01647.x.

Liu, B., Wei, R., Tang, J., Hong, J., Lu, Q., Guo, C., Wu, H., 2025. Human-centric disaster resilience: uncovering social inequity in climate change. Risk Anal. https://doi.org/10.1111/risa.70140.

Lundberg, S. M., & Lee, S.-I. (2017). *A unified approach to interpreting model predictions* Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, California, USA.

Ma, C., Culhane, D.P., Bachman, S.S., 2024. Understanding the dynamic process of human behavior changes towards disaster preparedness: an application of the integrated TTM with SCT and PMT. Int. J. Disaster Risk Reduct. 110, 104606. https://doi.org/10.1016/j.ijdrr.2024.104606.

Mallick, J., Salam, R., Amin, R., Islam, A.R.M.T., Islam, A., Siddik, M.N.A., Alam, G.M. M., 2022. Assessing factors affecting drought, earthquake, and flood risk perception: empirical evidence from Bangladesh. Nat. Hazards 112 (2), 1633–1656. https://doi.org/10.1007/s11069-022-05242-w.

Mase, A.S., Cho, H., Prokopy, L.S., 2015. Enhancing the social amplification of risk framework (SARF) by exploring trust, the availability heuristic, and agricultural advisors' belief in climate change. J. Environ. Psychol. 41, 166–176. https://doi.org/10.1016/j.jenvp.2014.12.004.

Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., Gao, J., 2021. Deep learning–based text classification. ACM Comput. Surv. 54 (3), 1–40. https://doi.org/10.1145/3439726.

Minea, I., Chelariu, O.E., Boicu, D., Iosub, M., 2025. An integrated approach to social vulnerability assessment of groundwater resources in the Moldavian Plain under temperate climatic conditions. J. Hydrol.: Reg. Stud. 58, 102308. https://doi.org/10.1016/j.ejrh.2025.102308.

Mukumbang, F.C., 2023. Retroductive theorizing: a contribution of critical realism to mixed methods research. J. Mixed Methods Res. 17 (1), 93–114. https://doi.org/10.1177/15586898211049847.

Noori, M.A.R., Sharma, B., Mehra, R., 2025. Disaster tweet classification using enhanced salp swarm algorithm. Web Intelligence 23 (1), 23–38. https://doi.org/10.3233/web-230019.

Oh, O., Eom, C., Rao, H.R., 2015. Research note—role of social media in social change: an analysis of collective sense making during the 2011 Egypt revolution. Inf. Syst. Res. 26 (1), 210–223. https://doi.org/10.1287/isre.2015.0565.

O'brien, R.M., 2007. A caution regarding rules of thumb for variance inflation factors. Qual. Quant. 41 (5), 673–690. https://doi.org/10.1007/s11135-006-9018-6.

Reid, M., 2013. Disasters and social Inequalities. Sociol. Compass 7 (11), 984–997. https://doi.org/10.1111/soc4.12080.

Renn, O., 1998. The role of risk perception for risk management. Reliab. Eng. Syst. Saf. 59 (1), 49–62. https://doi.org/10.1016/S0951-8320(97)00119-1.

Rippl, S., 2002. Cultural theory and risk perception: a proposal for a better measurement. J. Risk Res. 5 (2), 147–165. https://doi.org/10.1080/13669870110042598.

Rohrmann, B., Renn, O., 2000. Risk perception Research. In: Renn, O., Rohrmann, B. (Eds.), *Cross-Cultural Risk Perception: A Survey of Empirical Studies*. Springer, US, pp. 11–53. https://doi.org/10.1007/978-1-4757-4891-8_1.

Rutsaert, P., Regan, Á., Pieniak, Z., McConnon, Á., Moss, A., Wall, P., Verbeke, W., 2013. The use of social media in food risk and benefit communication. Trends Food Sci. Technol. 30 (1), 84–91. https://doi.org/10.1016/j.tifs.2012.10.006.

Rüttenauer, T., 2019. Spatial regression models: a systematic comparison of different model specifications using Monte Carlo experiments. Sociol. Methods Res. 51 (2), 728–759. https://doi.org/10.1177/0049124119882467.

Schroeder, A., Pennington-Gray, L., Donohoe, H., Kiousis, S., 2013. Using social media in times of crisis. J. Travel Tour. Mark. 30 (1–2), 126–143. https://doi.org/10.1080/10548408.2013.751271.

Siegrist, M., Árvai, J., 2020. Risk perception: reflections on 40 years of research. Risk Anal. 40 (S1), 2191–2206. https://doi.org/10.1111/risa.13599.

Siegrist, M., Cvetkovich, G., 2000. Perception of hazards: the role of social trust and knowledge. Risk Anal. 20 (5), 713–720. https://doi.org/10.1111/0272-4332.205064.

Slovic, P. (1987). Perception of Risk. *Science*, *236*(4799), 280–285. https://doi.org/10.1126/science.3563507#core-collateral-purchase-access.

Song, Y., Wang, T., Cai, P., Mondal, S.K., Sahoo, J.P., 2023. A comprehensive survey of few-shot learning: evolution, applications, challenges, and opportunities. ACM Comput. Surv. 55 (13s), 271. https://doi.org/10.1145/3582688.

Terpstra, T., Lindell, M.K., 2012. Citizens' perceptions of flood hazard adjustments. Environ. Behav. 45 (8), 993–1018. https://doi.org/10.1177/0013916512452427.

Thiery, W., Lange, S., Rogelj, J., Schleussner, C.-F., Gudmundsson, L., Seneviratne, S.I., Andrijevic, M., Frieler, K., Emanuel, K., Geiger, T., Bresch, D.N., Zhao, F., Willner, S. N., Büchner, M., Volkholz, J., Bauer, N., Chang, J., Ciais, P., Dury, M., Wada, Y., 2021. Intergenerational inequities in exposure to climate extremes. Science 374 (6564), 158–160. https://doi.org/10.1126/science.abi7339.

Thissen, D., Steinberg, L., Kuang, D., 2002. Quick and easy implementation of the benjamini-hochberg procedure for controlling the false positive rate in multiple comparisons. J. Educ. Behav. Stat. 27 (1), 77–83. https://doi.org/10.3102/10769986027001077.

Wahid, J.A., Xu, M., Ayoub, M., Jiang, X., Lei, S., Gao, Y., Hussain, S., Yang, Y., 2025. AI-driven social media text analysis during crisis: a review for natural disasters and pandemics. Appl. Soft Comput. 171, 112774. https://doi.org/10.1016/j.asoc.2025.112774.

Wang, Y., Yao, Q., Kwok, J.T., Ni, L.M., 2020. Generalizing from a few examples: a survey on few-shot learning. ACM Comput. Surv. 53 (3), 63. https://doi.org/10.1145/3386252.

Wilkinson, I., 2001. Social theories of risk perception: at once indispensable and insufficient. Curr. Sociol. 49 (1), 1–22. https://doi.org/10.1177/0011392101049001002.

Wirz, C.D., Xenos, M.A., Brossard, D., Scheufele, D., Chung, J.H., Massarani, L., 2018. Rethinking social amplification of risk: social media and zika in three languages. Risk Anal. 38 (12), 2599–2624. https://doi.org/10.1111/risa.13228.

Yang Jie, & Huang Xin. (2024). The 30 m annual land cover datasets and its dynamics in China from 1985 to 2023 [Data set]. In Earth System Science Data (1.0.3, Vol. 13, Number 1, pp. 3907–3925). Zenodo. doi.org/10.5281/zenodo.12779975.

Yin, J.-D.-C., Lui, J.-N.-M., 2024. Factors influencing risk perception during public health emergencies of international concern (PHEIC): a scoping review. BMC Public Health 24 (1). https://doi.org/10.1186/s12889-024-18832-z.

Zhang, X.A., Cozma, R., 2022. Risk sharing on Twitter: social amplification and attenuation of risk in the early stages of the COVID-19 pandemic. Comput. Hum. Behav. 126, 106983. https://doi.org/10.1016/j.chb.2021.106983.

Zhang, C., Luo, L., Xu, W., Ledwith, V., 2008. Use of local Moran's I and GIS to identify pollution hotspots of Pb in urban soils of Galway Ireland. Sci. Total Environ. 398 (1–3), 212–221. https://doi.org/10.1016/j.scitotenv.2008.03.011.

Zhang, J., Ma, X., Zhang, J., Sun, D., Zhou, X., Mi, C., Wen, H., 2023. Insights into geospatial heterogeneity of landslide susceptibility based on the SHAP-XGBoost model. J. Environ. Manage. 332, 117357. https://doi.org/10.1016/j.jenvman.2023.117357.

Zhang, P., Zhang, H., Kong, F., 2024a. Research on online public opinion in the investigation of the "7–20" extraordinary rainstorm and flooding disaster in Zhengzhou, China. Int. J. Disaster Risk Reduct. 105, 104422. https://doi.org/10.1016/j.ijdrr.2024.104422.

Zhang, P., Zhang, H., Kong, F., 2024b. Study on the evolution of online public opinion and government response strategies for the "7–20" extraordinary rainstorm and flooding disaster in Zhengzhou, China. Nat. Hazards. https://doi.org/10.1007/s11069-024-06904-7.

Zhang, W., Clark, R., Zhou, T., Li, L., Li, C., Rivera, J., Zhang, L., Gui, K., Zhang, T., Li, L., Pan, R., Chen, Y., Tang, S., Huang, X., Hu, S., 2024c. 2023: weather and climate extremes hitting the globe with emerging features. Adv. Atmos. Sci. 41 (6), 1001–1016. https://doi.org/10.1007/s00376-024-4080-3.

Zhang, Y., Xu, Y., Dong, F., 2025. An enhanced few-shot text classification approach by integrating topic modeling and prompt-tuning. Neurocomputing 617, 129082. https://doi.org/10.1016/j.neucom.2024.129082.

Zhang, P., Jiang, N., Li, Y., Wei, Z., Hui, P., 2026a. Leveraging large language models to explain provincial heterogeneity in online discussion of water pollution in China: a socioeconomic and environmental perspective. J. Clean. Prod. 543, 147517. https://doi.org/10.1016/j.jclepro.2026.147517.

Zhang, P., Li, Y., Wei, Z., Hui, P., 2026b. The geography of climate concern: a large-scale analysis of public discourse on extreme heat in China using social media and explainable AI. Environ. Impact Assess. Rev. 117, 108227. https://doi.org/10.1016/j.eiar.2025.108227.

Zhang, P., Li, Y., Wei, Z., Hui, P., Jiang, N., 2026c. Digital grievances: a geospatial analysis of public discourse on doctor-patient conflict in China. Appl. Geogr. 189, 103954. https://doi.org/10.1016/j.apgeog.2026.103954.

Zhang, P., Wei, Z., Kong, F., 2026d. Reconfiguring responsibility: an empirical analysis of crisis discourse and situational crisis communication on Douyin. Int. J. Disaster Risk Sci. https://doi.org/10.1007/s13753-026-00693-2.

Zhang, P., Wei, Z., Zhou, M., Evans, J., Hui, P., 2026e. Beyond aggregate feedback: a spatially-aware computational framework for understanding citizen-state interaction on social media. ACM Trans. Social Comput. https://doi.org/10.1145/3790251.

Zou, L., Liao, D., Lam, N.S.N., Meyer, M.A., Gharaibeh, N.G., Cai, H., Zhou, B., Li, D., 2023. Social media for emergency rescue: an analysis of rescue requests on Twitter during Hurricane Harvey. Int. J. Disaster Risk Reduct. 85, 103513. https://doi.org/10.1016/j.ijdrr.2022.103513.