



Research Article

<https://doi.org/10.1631/ENG.ITEE.2026.0120>

Social radar Doppler: multi-scenario validation of social media-based crisis early warning

Hetian CHEN¹, Lili FAN^{2✉}, Guangyu SHI³

¹School of Software, Xinjiang University, Urumqi 830046, China

²School of AI, Beijing Institute of Technology, Beijing 100081, China

³School of AI, Beijing Institute of Technology, Zhuhai 519088, China

Abstract: Social media is an essential data source for risk perception of complex social events. However, existing early warning methods mainly rely on popularity peaks, task-specific semantic labels, or observable propagation structures, often ignoring propagation rhythm changes prior to a crisis. This paper proposes an unsupervised online framework for social media crisis early warning. Aggregating heterogeneous behaviors into a unified intensity sequence, it detects anomalously accelerated propagation phases relative to historical baselines through causal smoothing, rate-of-change extraction, rolling quantile thresholds, persistence constraints, and cooling mechanisms. Compared to intensity-driven methods, it focuses on “whether the system is accelerating toward a risk state” rather than “whether it has reached a high popularity state;” compared to semantic- and structure-driven methods, it has weaker dependencies on annotations, ontologies, and complete propagation graphs. Strict causal replay experiments on the Russia–Ukraine crisis Weibo dataset and the Douban Movie Short Comments dataset demonstrate that the framework captures aggregative rising processes before key events early with low false alarm interference, exhibiting structural reusability across scenarios. Results indicate that detecting propagation dynamics changes offers a lightweight, interpretable, and transferable path for online early warning in complex systems.

Key words: Social media; Crisis early warning; Online anomaly detection; Causal dynamics; Social radar Doppler

1 Introduction

Major social events are typically accompanied by the rapid evolution of information diffusion, collective emotional shifts, and social mobilization. With the proliferation of social media platforms, large-scale user-generated data have provided new possibilities for real-time observation of social system dynamics. Compared with traditional statistics or retrospective reports, social media exhibits higher timeliness and broader participation, thus becoming a vital data source for risk perception and early warning (Imran et al., 2015; Gallotti et al., 2020; Alam et al., 2021).

Existing social media-based event detection and early

warning methods can generally be categorized into three types. The first type comprises methods based on popularity and temporal anomalies, which identify emergent events primarily through sudden surges in keyword frequency, post volume, or topic intensity. These methods are simple, efficient, and suitable for online deployment, but their core mechanism still relies on detecting increases in signal intensity. Consequently, they are more adept at identifying events that have already formed significant peaks (Weng and Lee, 2011; Xie et al., 2016). The second type includes methods based on content and semantic understanding, which characterize “what users are talking about” through topic modeling, sentiment analysis, and event extraction. While highly interpretable, these methods typically rely on stable semantic patterns, labeled data, or domain knowledge, and often face generalization challenges during the early, sparse, and noisy stages of unexpected events (Lau et al., 2012; Diao and Jiang, 2013; Parekh et al., 2024). The third type focuses on propagation structures and network evolution, modeling the information diffusion process through retweet chains, reply trees, or graph structures. These methods closely mirror social propagation mechanisms but typically require relatively complete interaction graphs or propagation chains,

✉ Lili FAN, lilifan@bit.edu.cn

✉ Hetian CHEN, <https://orcid.org/0009-0003-1067-7971>

Lili FAN, <https://orcid.org/0000-0002-8767-7123>

Guangyu SHI, <https://orcid.org/0009-0000-0892-867X>

CLC number: TP277; TP391

Received: Apr. 27, 2026; Revision accepted: July 7, 2026;

Crosschecked: July 22, 2026; Published online: Aug. 14, 2026

© The Authors 2026. Published by Zhejiang University Press Co., Ltd. This is an open access article distributed under the terms of the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

thereby limiting their application in real-time early warning scenarios (Ma J et al., 2017; Bian et al., 2020).

Although the aforementioned methods have made significant progress across the dimensions of heat, semantics, and structure, most of them still treat “whether the current state is abnormal” as their core detection objective. That is, they focus on whether the system has already exhibited significant popularity, clear semantic patterns, or abnormal propagation structures. For early warning, however, the more critical question is not “whether it is already abnormal now,” but “whether the system is rapidly accelerating toward an abnormal state.” In the initial stages of event formation, absolute popularity has often not yet significantly increased, semantic patterns are not yet stable, and propagation structures have not fully unfolded. Thus, relying solely on state detection often fails to capture the risk formation process in a timely manner. Therefore, the key methodological gap addressed in this study is not merely event detection after anomalies emerge, but the causal online measurement of acceleration in propagation dynamics before clear intensity, semantic, or structural peaks are fully formed.

To address this limitation, this paper proposes an unsupervised online framework for social media crisis early warning, namely, social radar Doppler (SRD). This method shifts the early warning task from traditional intensity monitoring to the detection of changes in propagation dynamics, focusing on identifying anomalously accelerated stages of social information propagation relative to historical rhythms. Specifically, under strict causal constraints, SRD aggregates heterogeneous social behaviors into a unified social intensity sequence, suppresses high-frequency noise through causal smoothing, and characterizes the dynamic features of the propagation process with the rate of change. On this basis, combined with rolling quantile thresholds, persistence constraints, and cooling mechanisms, it transforms continuous dynamic signals into stable multi-level early warning outputs. Compared with existing methods, SRD exhibits three main characteristics: first, it does not rely on supervised labels, domain-specific rules, or complete propagation graphs, thus achieving favorable cross-scenario reusability; second, the entire computational process strictly complies with causal online settings, eliminating future information leakage; third, by focusing on relative changes in propagation rhythms rather than absolute intensity levels, it can identify early signals of potential risks before the formation of popularity peaks.

The main contributions of this paper are summarized as follows:

1. We propose an SRD framework for social media crisis early warning, which formulates early warning as a propagation-dynamics monitoring problem and provides a causal online mechanism for identifying sustained acceleration patterns in collective behavior.
2. We design a strictly causal, unsupervised online early warning mechanism. Through causal smoothing, rolling quantile thresholds, persistence constraints, and cooling mechanisms, this mechanism achieves stable identification of anomalously accelerated propagation phases and generates multi-level early warning outputs.
3. The experimental results provide evidence of the frame-

work’s effectiveness and its potential reusability across different propagation scenarios.

2 Related works

2.1 Heat- and temporal-anomaly-driven social media analysis

Event detection based on popularity and temporal anomalies constitutes one of the earliest systematic approaches in social media analysis. These methods generally represent social media streams as time series of keyword frequency, post volume, or topic intensity, to identify emergent events through burst detection, anomaly detection, or sliding window statistical mechanisms. Representative works include a Twitter event detection method based on wavelet transformations and graph partitioning (Weng and Lee, 2011), as well as the TopicSketch framework designed for real-time scenarios (Xie et al., 2016). Subsequent research has also attempted to incorporate distributed semantic representations to alleviate the overreliance on word frequencies inherent in traditional burst methods (Li et al., 2025).

The primary advantage of this approach resides in its simplicity, low deployment overhead, and high sensitivity to large-scale unexpected events, making it particularly suitable for platform-level monitoring and real-time early warning. However, its core criterion is still primarily derived from sudden quantitative spikes, meaning that it is more adept at identifying events that have already formed significant peaks. For evolutionary processes with low initial popularity, low discussion volume, and high risk, these methods are usually insufficiently sensitive. Even with semantic enhancements or deep anomaly detection methods (Gui et al., 2025), they mostly aim to enhance the stability of intensity anomaly detection, rather than altering the fundamental paradigm of using peak formation as the primary basis.

2.2 Content- and semantics-driven social media analysis

For public opinion analysis driven by content and semantic understanding, the core question is no longer “whether the volume is increasing,” but “what users are actually discussing.” Relevant research includes online topic models for tracking emerging events on Twitter (Lau et al., 2012), probabilistic methods that jointly model topics, events, and users (Diao and Jiang, 2013), and works that apply structured event extraction to epidemic prediction (Parekh et al., 2024). Concurrently, research in aspect-based sentiment analysis, Chinese textual affective structure analysis (Xiong et al., 2025), emotion recognition, multimodal semantic understanding, and large language model (LLM)-aided event detection continues to enhance the ability to characterize social media content (Tang et al., 2024; Ma CB et al., 2025; Peng SC et al., 2025; Qiu et al., 2025; Zhang et al., 2025).

The strength of the semantic approach lies in its high interpretability, as it can reveal the topics, emotions, or event semantics corresponding to risk signals. This makes it easier to construct interpretable event profiles and incorporate domain knowledge. However, such methods are highly dependent on

labeled datasets, event ontologies, domain knowledge, or model transferability. Short-text noise, cross-platform language differences, and domain shifts can all affect their stability. More importantly, these methods often focus on post-level or localized event-level semantics, making it difficult to directly construct continuous and dynamic early warning representations across time and scenarios.

2.3 Propagation-structure- and network-evolution-driven social media analysis

Propagation-structure- and network-evolution-driven methods focus on how information propagates and how the underlying interaction topology develops over time. Representative studies include propagation tree-based rumor detection methods (Ma J et al., 2017), as well as bi-directional graph convolutional network (Bi-GCN), which jointly models top-down and bottom-up propagation features (Bian et al., 2020). Further research indicates that information cascade paths, topological structures, and temporal unfolding processes jointly shape the macroscopic patterns of information propagation (Zhou et al., 2021; Peng XY et al., 2025).

The advantage of these methods is that they better align with real social diffusion mechanisms. For tasks such as rumor identification, organized propagation, and abnormal diffusion, structural patterns in propagation chains, reply trees, or interaction graphs are often more discriminative than pure text content. However, their limitations are notable: they typically rely on relatively complete interaction data, propagation chains, or reply structures; moreover, the high computational overhead of graph construction and model training limits their direct deployment in large-scale online early warning. Especially in the early stages of an event, propagation structures remain underdeveloped.

In summary, prior studies have developed three mature paths for social media event analysis: popularity anomalies, semantic understanding, and propagation structures. These lines provide important foundations, while still leaving room for an online-computable representation of propagation acceleration that can operate before intensity, semantic, or structural signals become fully formed.

2.4 Social radar and dynamic situation awareness

Unlike event detection tailored for single tasks, social radar research emphasizes organizing heterogeneous and continuously arriving social data into a dynamically updated observation stream for situational understanding, risk identification, and decision support. This perspective can be traced back to continuous observation and parallel execution in complex social systems (Wang, 2007), and has been further developed in social radar workflows and open-source intelligence studies (Lasswell, 1948; Schramm and Roberts, 1954; Mathieu et al., 2012; Ramakrishnan et al., 2014). However, existing social radar research mainly focuses on scenario construction, workflow organization, and multi-source sensing. In contrast, the present study focuses on a narrower methodological issue: how to transform heterogeneous social behaviors into a causal, online-computable dynamic indicator that characterizes deviations in propagation rhythm.

3 Social radar Doppler framework

This section provides the formal definition of the SRD framework. Distinct from traditional anomaly detection based on intensity levels, SRD models early warning as a problem of detecting changes in propagation dynamics. Its core lies in identifying anomalously accelerated phases of social information propagation relative to its historical rhythm. Under strict causal settings, SRD maps heterogeneous social behaviors into a set of online-computable dynamic variables and generates stable early warning outputs through a multi-level threshold mechanism.

3.1 Problem formulation and core definition

Let $t = 1, 2, \dots, T$ denote the daily-scale time index. For any given scenario, let $S(t)$ be the social intensity sequence, $\tilde{S}(t)$ be its smoothed result, and $v(t)$ be the corresponding social velocity. The goal of SRD is to identify anomalously accelerated phases during the propagation process online, relying solely on current and historical observations.

The SRD indicator is defined as follows:

$$f_d(t) = \frac{v(t)}{\sigma_{\text{rob}} + \varepsilon}, \quad t = 2, 3, \dots, T, \quad (1)$$

where σ_{rob} is the scene-level fixed robust scale estimated from a short calibration window, and $\varepsilon > 0$ is a numerical stability term. This indicator captures the normalized deviation of the current rate of propagation change relative to the scenario's background fluctuation scale. Unlike intensity-based detection, the core criterion of SRD is not whether the system is already anomalous, but whether a dynamic deviation occurs in the propagation process.

3.2 Social signal construction

To unify discrete, heterogeneous, and dimensionally varied social media behaviors into a one-dimensional (1D) input that can be processed online, this paper first aggregates scenario-related observations by natural days. Let $x_j(t)$ be the value of the j^{th} raw behavioral feature on day t , such as content production, interactive diffusion, and participation scale. After feature transformation and linear weighting, the raw social intensity is defined as

$$A_{\text{raw}}(t) = \sum_{j=1}^d w_j \phi_j(x_j(t)), \quad (2)$$

where $\phi_j(\cdot)$ is the feature transformation function, w_j is the corresponding weight, and d is the feature dimension. For count-based variables, $\phi_j(\cdot)$ typically takes $\log(1+x)$ to mitigate the effects of long-tail distributions and extreme values.

Because this paper focuses on strict online early warning settings, social intensity does not use global normalization over the entire time range, but rather causal rolling normalization. Given a normalization window length W_{norm} , the normalized historical window at time t is defined as

$$\mathcal{J}(t) = \{\max(1, t - W_{\text{norm}} + 1), \dots, t\}, \quad t = 1, 2, \dots, T. \quad (3)$$

Based on this, the rolling minimum and rolling maximum are

respectively defined as

$$A_{\min}(t) = \min_{\tau \in \mathcal{J}(t)} A_{\text{raw}}(\tau), \quad (4)$$

$$A_{\max}(t) = \max_{\tau \in \mathcal{J}(t)} A_{\text{raw}}(\tau). \quad (5)$$

Thus, the social intensity is defined as

$$S(t) = 100 \cdot \frac{A_{\text{raw}}(t) - A_{\min}(t)}{A_{\max}(t) - A_{\min}(t) + \varepsilon_{\text{norm}}}, \quad t = 1, 2, \dots, T, \quad (6)$$

where $\varepsilon_{\text{norm}} > 0$ is a numerical stability term. Therefore, $S(t) \in [0, 100]$. This normalization approach serves two purposes. First, it strictly restricts the intensity calculation at each moment to depend only on current and historical observations, thereby avoiding future information leakage. Second, it allows the intensity representation to adaptively update according to the scenario's recent historical range, preventing its scale from being fixed by distant extreme values within the entire observation period. The resulting $S(t)$ is no longer a simple superposition of raw behaviors, but a dynamically normalized social intensity sequence relative to the recent background, serving as input for subsequent causal smoothing and social velocity calculations.

It should be noted that the primary differences between scenarios do not lie in the dynamic structure of SRD itself, but mainly in the feature composition and weight configuration of $A_{\text{raw}}(t)$. In other words, what SRD reuses is the unified dynamic calculation framework—from intensity to velocity and finally to the Doppler-inspired indicator—rather than assuming that all scenarios share identical raw input semantics.

3.3 Causal smoothing and social velocity

Although $S(t)$ obtained from Eq. (6) is uniformly mapped to a fixed interval, it still contains high-frequency disturbances such as platform diurnal rhythms, overall traffic fluctuations, and collection noise. Performing differencing directly on $S(t)$ makes the sequence overly sensitive to single-day spikes, thereby affecting the characterization of short- to medium-term propagation trends. Therefore, before constructing dynamic variables, causal smoothing is applied to the intensity sequence.

We first define the smoothed intensity sequence and derived dynamic variables used in SRD. Let $\tilde{S}(t)$ denote the causally smoothed version of $S(t)$ obtained via exponential moving average (EMA). Let $v(t)$ denote the first-order temporal difference of $\tilde{S}(t)$, referred to as the social velocity. Finally, let $\mathcal{C}$ denote the calibration window used for robust scale estimation, which contains only the initial warm-up observations and is excluded from evaluation.

This paper employs EMA to smooth $S(t)$. Let the smoothing window span be τ_{EMA} . Then, the smoothing coefficient is

$$\alpha_{\text{EMA}} = \frac{2}{\tau_{\text{EMA}} + 1}, \quad (7)$$

and the corresponding causal recursive form is

$$\tilde{S}(t) = \alpha_{\text{EMA}} S(t) + (1 - \alpha_{\text{EMA}}) \tilde{S}(t-1), \quad t \geq 2, \quad (8)$$

with the initial value $\tilde{S}(1) = S(1)$. Since this recursion relies solely on current and historical observations, it satisfies the strict online setting.

Based on the smoothed intensity $\tilde{S}(t)$, social velocity is defined as the first-order difference:

$$v(t) = \tilde{S}(t) - \tilde{S}(t-1), \quad t = 2, 3, \dots, T, \quad (9)$$

with $v(1) = 0$.

Consequently, SRD transforms the raw intensity sequence into a dynamic metric describing its temporal rate of change. Compared to $S(t)$, which describes the current level, $v(t)$ directly reflects the changing trend of the propagation process: a positive value indicates rising intensity, and a negative value indicates falling intensity. Therefore, subsequent warning judgments rely not on the intensity level itself, but on the abnormal deviation of this change rate relative to background fluctuations.

3.4 Scene-level robust scaling and social Doppler indicator

After obtaining the social velocity $v(t)$, a key issue is how to scale it so that propagation rates of change across different scenarios are comparable. Because the magnitude of background fluctuations in social media sequences can vary significantly between scenarios, the same raw rate of change does not correspond to the same degree of anomaly. Therefore, to enhance robustness (Liu YS et al., 2024), instead of using a fixed scale shared across all scenarios, SRD introduces a scene-level robust scale to normalize propagation velocity.

Given the predefined calibration window $\mathcal{C}$ at the beginning of each scenario, which is solely used to estimate the background fluctuation scale, let the median of the velocity sequence within this window be

$$m_v = \text{median}\{v(t) : t \in \mathcal{C}\}. \quad (10)$$

The robust scale is defined using the median absolute deviation (MAD):

$$\sigma_{\text{rob}} = 1.4826 \cdot \text{median}\{|v(t) - m_v| : t \in \mathcal{C}\}. \quad (11)$$

Here, the coefficient 1.4826 makes MAD comparable to the standard deviation under Gaussian assumptions, while preserving its robustness to outliers and short-term spikes.

Unlike online scale estimation based on rolling windows, SRD determines σ_{rob} once during the calibration phase and keeps it fixed during subsequent detection processes. This design ensures that propagation rates of change in subsequent periods are always measured against the same background scale, avoiding threshold drift issues caused by rolling statistics being elevated as the event develops.

With the scene-level robust scale σ_{rob} defined in Eq. (11), the SRD indicator is obtained using Eq. (1).

Thus, $f_a(t)$ no longer represents absolute change rates, but rather the normalized deviation of the current propagation velocity relative to the background fluctuation scale. A larger value indicates a more significant acceleration of the propagation process relative to its normal rhythm. Compared to traditional detection methods based on intensity or unnormalized rates of change, this scaling treatment allows early warning judgments to focus on relative acceleration rather than absolute magnitude. Here, the term ‘‘Doppler’’ is used as an interpretive analogy rather than a literal physical frequency-shift

model. Just as the Doppler effect perceives velocity changes via frequency shift, the proposed indicator perceives propagation acceleration via the normalized rate of change in social intensity.

3.5 Quantile-based multi-level early warning

After calculating the social Doppler indicator $f_d(t)$ and social velocity $v(t)$, SRD converts continuous dynamic signals into discrete online early warning decisions. This process involves three steps: dynamic threshold construction, candidate warning generation, and temporal consistency constraints.

To adapt to the non-stationary nature of social media sequences, SRD does not use fixed thresholds but constructs dynamic thresholds based on recent historical statistics. Given a rolling window length W , the causal window at time t is defined as

$$\mathcal{I}(t) = \{\max(t_0, t - W + 1), \dots, t\}, \quad t = t_0, t_1, \dots, T, \quad (12)$$

where t_0 is the starting time of formal detection. Let the empirical quantile operator be

$$\mathcal{Q}_{p,W}[x](t) = \text{Quantile}_p(\{x(\tau)\}_{\tau \in \mathcal{I}(t)}), \quad (13)$$

where $\text{Quantile}_p(\cdot)$ represents the empirical p^{th} quantile. For each warning level $k \in \{1, 2, 3\}$, let its corresponding quantile level for $f_d(t)$ be $p^{(k)}$, and that for $v(t)$ be $q^{(k)}$; the dynamic thresholds are then defined as

$$\theta_f^{(k)}(t) = \mathcal{Q}_{p^{(k)},W}[f_d](t), \quad \theta_v^{(k)}(t) = \mathcal{Q}_{q^{(k)},W}[v](t). \quad (14)$$

This quantile threshold updates dynamically based on local history, ensuring that early warning judgments are always made relative to the current background distribution, thereby maintaining stability in non-stationary environments.

Based on this, SRD employs a dual-threshold mechanism to generate candidate warning levels $\ell^{(0)}(t) \in \{0, 1, 2, 3\}$:

$$\begin{cases} C_k(t) := [f_d(t) > \theta_f^{(k)}(t)] \wedge [v(t) > \theta_v^{(k)}(t)], & k \in \{1, 2, 3\}, \\ \ell^{(0)}(t) = \max(\{k \in \{1, 2, 3\} : C_k(t)\} \cup \{0\}). \end{cases} \quad (15)$$

This design requires both the scaled deviation $f_d(t)$ and the raw rate of change $v(t)$ to exceed their respective thresholds, focusing the warning on phases that possess both a significant upward trend and anomalous acceleration. Compared to single-sequence thresholds, this dual constraint effectively suppresses false triggers caused by isolated spikes or scale amplification.

The multi-level thresholds correspond to warning signals of varying intensities, where $L_1/L_2/L_3$ indicate weak, moderate, and strong acceleration phases, respectively. To enhance warning stability, SRD applies temporal consistency constraints to the candidate sequence, including persistence filtering and cooling mechanisms.

First, a minimum duration length m is introduced. For any continuous non-zero interval $[a, b]$, if $b - a + 1 < m$, such intervals are treated as short-term noise and eliminated, resulting in the persistence-filtered sequence $\ell^{(1)}(t)$.

Next, a cooling window c is introduced to prevent the same propagation process from being repeatedly triggered. Let

$$t^*(t) = \max\{\tau < t : \ell(\tau) > 0\}. \quad (16)$$

The final warning is then defined as

$$\ell(t) = \ell^{(1)}(t) \cdot \mathbb{I}(t - t^*(t) > c), \quad t = t_0, t_1, \dots, T. \quad (17)$$

This mechanism consolidates continuous trigger signals into a finite number of stable warning intervals, reducing repetitive alarms on the timeline. Finally, the final level $\ell(t)$ is mapped into three warning indicator sequences:

$$\alpha_{L_k}(t) = \mathbb{I}(\ell(t) = k), \quad k \in \{1, 2, 3\}. \quad (18)$$

Thus, under strict causal settings, SRD realizes a complete online modeling workflow from continuous dynamic metrics to a stable multi-level warning output.

To further clarify the methodological distinction between representative early warning baselines and the proposed framework, Table 1 summarizes their respective detection philosophies. While the cumulative sum (CUSUM) and early aberration reporting system (EARS) can both capture dynamic changes to some extent, their core mathematical mechanisms differ fundamentally from SRD. Specifically, although CUSUM implicitly captures trend changes by accumulating state deviations, this cumulative-sum approach may exhibit delayed responses to emerging changes when substantial accumulation is required before threshold crossing. Similarly, while EARS is a rate-based method sensitive to recent incidence surges, it primarily responds to first-order rate spikes (i.e., abrupt increases in current activity levels). In contrast, SRD explicitly targets sustained acceleration dynamics. Rather than emphasizing isolated rate spikes or delayed cumulative shifts, SRD's Doppler-inspired indicator $f_d(t)$ is specifically designed to identify and track temporally consistent acceleration patterns in collective behavior. Consequently, the three methods differ not only in their triggering mechanisms but also in the behavioral signals they are intended to capture.

Table 1 Comparison of detection philosophies

Method	Detection philosophy	Primary monitoring variable
CUSUM	Cumulative deviation detection	Activity deviation from the baseline
EARS	Abrupt anomaly detection	Observed activity level
SRD (ours)	Sustained acceleration detection	$f_d(t)$

4 Datasets

To validate the effectiveness and transferability of the proposed method across different propagation scenarios, this paper selects two social media datasets with distinct differences: one is a sudden-event-driven crisis scenario dataset and the other is a non-crisis cultural consumption dataset. All analyses are conducted based on time series aggregated at the natural-day granularity and only use anonymized population-level statistical information.

4.1 RUW: Russia–Ukraine crisis Weibo dataset

The main experiment of this paper is based on the Russia–Ukraine crisis Weibo dataset (RUW) released by Fung and Ji (2022). This dataset collects public discussion data

on the Weibo platform surrounding the 2022 Russia–Ukraine conflict, including main posts, comments, timestamps, and interaction information using focused crawling strategies (Liu JF et al., 2025). Data collection began in late February 2022 and retroactively covered historical records for approximately three months before the outbreak of the conflict. According to the original data description, the first version of the dataset contains 27 341 main posts, over 3.5 million comments, and 107 797 unique users.

In this paper, we aggregate the original records by natural days based on Beijing Time (UTC+8) to construct daily time series. Specifically, the following three types of basic behavioral signals are extracted: content publication volume, which is the number of main posts published daily; interaction diffusion volume, which is the total count of user interactions (such as comments and reposts) that occurred on that day, counted by the actual time of occurrence; participating user scale, which is the number of unique users who authored main posts that day. These multi-dimensional behavioral signals are uniformly represented as daily sequences and serve as inputs for the subsequent construction of social intensity (see Section 3). This paper does not model individual users; all analyses are conducted purely at the aggregate statistical level.

As the primary experimental scenario, the RUW evaluates the method’s performance on event-level early warning tasks.

4.2 Douban Movie Short Comments dataset (cross-scenario validation)

To further evaluate the method’s applicability in non-crisis scenarios, this paper selects the Douban Movie Short Comments (DMSC) dataset (Utmhikari, 2019; Chung-I, 2020) for cross-scenario validation. This dataset contains over 2.1 million short comments for 28 movies, with 1–5 star ratings. Since different movies vary significantly in genre, audience, and release cycles, directly mixing them might introduce confounding factors. Therefore, this paper adopts a single-movie observation strategy, specifically selecting *Avengers: Age of Ultron* as a representative case to build a time-series subset.

This subset contains 54 153 comments from 50 650 unique users between Apr. 4, 2015 and Jan. 22, 2017. Aggregated by natural day, the time span covers 660 days (659 with records).

In this scenario, three daily observational signals are constructed: comment publication volume (new comments daily), participating user scale (unique users daily), and rating statistics (aggregate indicators like average rating).

Although this dataset differs from RUW in feature composition, both are ultimately mapped into a unified 1D social intensity sequence. The proposed method relies only on the temporal variation of this aggregated intensity, maintaining independence from specific input feature semantics.

Because the Douban scenario lacks multiple discrete sudden events, this paper avoids the event set evaluation approach. Instead, the movie’s release date (Apr. 22, 2015) serves as a retrospective time anchor to analyze the temporal relationship between warning signals and the propagation launch phase. This dataset primarily validates the method’s generalization under different propagation structures, rather than replacing the event-level evaluation in the RUW scenario.

5 Experimental design

This section details the evaluation protocol and statistical criteria of SRD under unified causal online settings. To ensure comparability among different methods, this paper standardizes the data replay approach, event definitions, warning matching rules, quiet day construction, baseline methods and parameter settings, and evaluation metrics. Unless otherwise specified, all experimental results are obtained under strict online replay conditions without using future information.

5.1 Causal online replay protocol

This paper employs a strict causal online replay method to evaluate all methods. Assuming that the input is a daily time series, the replay process advances step by step by natural days. At any time t , methods are only allowed to use observation data up to t and their causally derived variables to calculate statistics, thresholds, and warning states. Using any future information is strictly prohibited.

Each method outputs a series of warning intervals along the timeline, represented as $\mathcal{A} = \{[a_i, b_i]\}_{i=1}^A$, where A is the total number of warning periods, and $[a_i, b_i]$ denotes the start and end of the i^{th} warning period. To reduce triggers from short-term fluctuations and avoid double-counting the same risk process, this paper applies consistent post-processing rules to all methods: a minimum duration constraint, where only warning intervals lasting no less than 2 days are retained, and a cooling mechanism, where after a warning period ends, a 2-day cooling period is introduced, during which no new warnings are permitted to trigger.

This unified processing ensures a consistent statistical criterion across different methods at the warning output level, thereby ensuring that performance differences primarily reflect the detection mechanisms themselves rather than discrepancies in post-processing strategies.

5.2 Event anchor construction

In the RUW scenario, evaluation is based on a set of event anchors. First, the original date is obtained according to the event timeline and sorted chronologically. To prevent high-density events from being double-counted, a chain merging strategy is adopted: if the interval between two adjacent event dates does not exceed 10 days, they are grouped into the same event cluster.

After merging, the event set is obtained: $\mathcal{E} = \{e_1, e_2, \dots, e_E\}$, where each event cluster uses the earliest date D_e as the representative time point. Under this setting, the RUW scenario yields $E = 7$ events.

For the Douban Movie scenario, since it does not contain multiple discrete sudden events, this paper does not use the event set evaluation approach. Instead, the movie’s release date is used as a retrospective anchor to analyze the temporal relationship between warning signals and the propagation launch.

5.3 Event-warning matching rules

To evaluate whether a warning provides effective advance notification before an event occurs, this paper defines

a maximum lead window $L_{\max}$ for each event. A warning is considered a valid candidate for event e only if the warning start point a_i satisfies

$$a_i \in [D_e - L_{\max}, D_e]. \quad (19)$$

We uniformly set $L_{\max} = 30$ days, meaning that only warnings triggered within 30 days prior to the event are considered valid candidates. Warnings occurring on the day of the event or later are not counted as early warnings.

A greedy one-to-one matching strategy is applied to the candidate set: each event can be matched with at most one warning period, and each warning period can be matched with at most one event. For each event, the warning start point closest to the event time is selected:

$$a^* = \max\{a_i : a_i \in [D_e - L_{\max}, D_e]\}. \quad (20)$$

If no warning meets the condition, the event is marked as a miss. For successfully matched events, the lead time (LT) is defined as

$$LT(e) = D_e - a^*, \quad LT(e) \in \{1, 2, \dots, L_{\max}\}. \quad (21)$$

This matching mechanism ensures that a single warning interval is not repeatedly used for multiple events, nor are multiple warnings attributed to the same event repeatedly.

5.4 Quiet days and false alarm metrics

Besides event hit capability, the stability of warning methods during non-event periods is equally important. To this end, this paper constructs a set of quiet days $\mathcal{Q}$ in the RUW scenario. Specifically, extend 3 weeks at both ends of the original time series to increase non-event samples; for each event anchor D_e , exclude its lead window:

$$[D_e - L_{\max}, D_e]. \quad (22)$$

Under our setting, the number of quiet days is $N_{\text{quiet}} = 33$.

The false alarm coverage days are defined as

$$N_{\text{fa_day}} = \sum_{t \in \mathcal{Q}} \mathbb{I}(\exists i : t \in [a_i, b_i]). \quad (23)$$

Thus, the false alarm rate (FAR) is estimated as

$$\widehat{\text{FAR}} = \frac{N_{\text{fa_day}}}{N_{\text{quiet}}}. \quad (24)$$

Simultaneously, to characterize the frequency of false alarms, this paper introduces the average run length (ARL):

$$\text{ARL}_0 \approx \frac{1}{\widehat{\text{FAR}}}, \quad (25)$$

where ARL_0 denotes the ARL under the in-control (quiet-day) state.

When zero false alarms are observed, the rule-of-three is applied to provide a 95% upper confidence limit:

$$\text{FAR}_U^{95\%} \approx \frac{3}{N_{\text{quiet}}}, \quad \text{ARL}_{0,L}^{95\%} \approx \frac{N_{\text{quiet}}}{3}, \quad (26)$$

where $\text{FAR}_U^{95\%}$ is the 95% upper confidence bound for FAR, whereas $\text{ARL}_{0,L}^{95\%}$ is the 95% lower confidence bound for in-control ARL.

This statistical treatment prevents overly optimistic interpretation of zero false alarm results.

5.5 Baselines and parameter settings

This paper compares the proposed SRD method with several classic statistical warning methods, including standard score (Z-score), EARS (C2+C3), and CUSUM. All methods are evaluated using the same daily input data, a unified causal online replay protocol, and consistent warning post-processing rules, ensuring fairness and comparability of the results. For the SRD method, a pre-calibration window is first set at the beginning of the sequence to estimate the scene-level robust scale (the first 15 days for RUW and the first 14 days for Douban Movie); this phase is excluded from formal evaluation. During formal detection, two sets of rolling window lengths $W \in \{15, 30\}$ are used to construct quantile thresholds to examine warning performance under different historical receptive windows. The causal smoothing parameter is uniformly set to $\tau_{\text{EMA}} = 7$. In addition, all methods share identical post-processing parameters for warnings: a minimum duration of 2 days and a cooling period of 2 days. Through these standardized configurations, performance differences among methods reflect the warning mechanisms themselves, rather than differences in experimental settings.

5.6 Evaluation metrics

Let the total number of events be E , the total number of warning periods be A , and the number of successfully matched events and warning periods be denoted as $\text{TP}_{\text{events}}$ and $\text{TP}_{\text{alerts}}$ respectively, with both being equal. Thus, event recall, warning precision, and the F1 score are defined as

$$\text{Recall} = \frac{\text{TP}_{\text{events}}}{E}, \quad \text{Precision} = \frac{\text{TP}_{\text{alerts}}}{A}, \quad (27)$$

$$\text{F1} = \frac{2 \cdot \text{Recall} \cdot \text{Precision}}{\text{Recall} + \text{Precision}}. \quad (28)$$

For all successfully matched events, this paper further reports the median and interquartile range of the lead time $LT(e)$ to characterize the earliness of warnings in typical cases and their dispersion. Let $\mathcal{E}_{\text{matched}}$ denote the set of successfully matched events. Meanwhile, we report

$$\text{R7d} = \frac{1}{\text{TP}_{\text{events}}} \sum_{e \in \mathcal{E}_{\text{matched}}} \mathbb{I}(LT(e) \leq 7), \quad (29)$$

which is the proportion of matched events with a lead time of no more than 7.

Regarding robustness on quiet days, this paper reports the false alarm rate $\widehat{\text{FAR}}$ and the derived ARL_0 . Here, $\widehat{\text{FAR}}$ measures the proportion of non-event periods erroneously covered by warnings, and ARL_0 characterizes the average interval between two false alarms under control conditions.

For proportional indicators, this paper uniformly reports Wilson 95% confidence intervals. Specifically, the sample size for Recall is E , for Precision is A , for R7d is $\text{TP}_{\text{events}}$, and for FAR is N_{quiet} . These settings enable comparison of the statistical uncertainty of various proportional metrics within a unified framework.

6 Results

This section reports the experimental results of SRD under the unified causal online replay protocol defined in Section 5. It first presents the core comparative results in the

RUW main experiment, then analyzes the warning behavior of the main configuration via timeline results, and finally further examines the method's cross-scenario applicability in the Douban Movie scenario. Unless otherwise specified, all results employ scene-level robust scale estimation, and the calibration window is excluded from the formal evaluation.

6.1 Main experimental results on RUW

The event anchors in the RUW scenario result in $E = 7$ events after 10-day chain merging. Event-level evaluation uses a one-to-one "one event, one warning" matching strategy, with the pairing window defined as $[D_e - L_{\max}, D_e]$, where $L_{\max} = 30$. All methods are run under identical daily inputs, the same causal online replay protocol, and consistent warning post-processing rules. Quiet days are constructed according to the criteria described in Section 5, using FAR to measure false alarm coverage.

Table 2 presents the core results in the RUW main experiment. Overall, SRD ($W=15$) achieved the most balanced performance across event coverage, warning accuracy, and false alarm control, with the highest $F_1 = 0.615$, alongside a Recall of 0.571 and Precision of 0.667. This result indicates that under the current evaluation parameters, SRD ($W=15$) can successfully detect key events while maintaining good warning period accuracy.

Table 2 Core evaluation results in the RUW scenario

Method	Recall	Precision	F1	LT _{med}	FAR
SRD ($W=15$)	0.571	0.667	0.615	12.5	0.0000
SRD ($W=30$)	0.286	0.667	0.400	13.0	0.0000
Z-Score ($W=15$)	0.429	0.600	0.500	12.0	1.0000
EARS ($W=15$, C2+C3)	0.714	0.417	0.526	7.0	0.0667
CUSUM ($W=15$)	0.143	1.000	0.250	14.0	0.0556

Recall is the event hit rate, Precision is the positive predictive value (PPV) of warning periods, and LT is the lead time (days) with the subscript "med" denoting median. The best results are in bold

The differences in results under various window lengths reflect the temporal structural characteristics of the RUW scenario. Compared to SRD ($W=15$), SRD ($W=30$) maintains the same Precision but experiences a drop in Recall from 0.571 to 0.286, and a drop in F1 from 0.615 to 0.400. For sequences like RUW, characterized by phased rapid increases, a shorter rolling window better captures structural acceleration changes prior to events. With a longer window, local anomalies are more easily smoothed into the broader historical background, thereby reducing the probability of continuous threshold crossings and successful detections.

Baseline methods show different trade-offs. EARS achieves the highest Recall (0.714), but its Precision is only 0.417, indicating that its higher event coverage costs more unmatched warning periods. CUSUM maintains a Precision of 1.000, but its Recall is only 0.143, reflecting highly conservative triggering. Although Z-Score detects some events, its FAR (1.0000) implies that non-event intervals are almost entirely covered by false triggers. In contrast, SRD ($W=15$) achieves a superior overall balance in the main experiment.

It should be noted that the aforementioned baseline methods primarily rely on absolute intensity levels or unnormalized variations for detection. Their core criterion remains whether

the current state is already anomalous, rather than whether it is accelerating toward an anomalous state. Consequently, during the early stages before a significant peak forms, these methods are more prone to trigger delays or oversensitivity to short-term fluctuations. By incorporating dynamic variation and relative scale normalization, SRD ensures that warning judgments correspond more directly to structural acceleration phases within the propagation process.

In terms of quiet day robustness, both SRD ($W=15$) and SRD ($W=30$) exhibit an FAR of 0.0000. According to the statistical criteria in Section 5, when $N_{\text{quiet}} = 33$, the 95% upper confidence bound for FAR provided by the rule-of-three is $\text{FAR}_{\text{U}}^{95\%} \leq 0.091$, and the corresponding $\text{ARI}_{0,L}^{95\%}$ lower bound is 11. Combining this with its Recall performance, it is evident that SRD's low false alarm rate does not merely rely on high thresholds to suppress triggers. Instead, while maintaining a certain level of event coverage, it effectively reduces the warning burden during non-event periods through selective responses to propagation acceleration phases.

Fig. 1 illustrates the timeline results of SRD ($W=15$) in the RUW scenario. It can be observed that warning triggers are primarily concentrated in phases where f_d continuously rises and crosses the rolling quantile thresholds, with several warning intervals located before the event anchors. During these intervals, the intensity level has not yet reached its peak, but the rate of propagation change has already shown a significant uplift. This indicates that what SRD captures is not static intensity anomalies, but the propagation acceleration process prior to event formation. This phenomenon is consistent with the method's dynamic modeling objective.

Simultaneously, the Recall of 0.571 for SRD ($W=15$) indicates that some events remain undetected. A retrospective observation of the time series for these missed events reveals that their propagation processes resemble gradual accumulation rather than phased acceleration, lacking obvious local surges in the rate of change. This presents a structural mismatch with SRD's trigger mechanism. Therefore, these missed cases to some extent reflect the method's applicability boundaries rather than pure detection failures.

6.2 Cross-scenario validation on Douban Movie

To examine whether SRD can transcend crisis contexts and apply to more general social propagation processes, this paper further conducts cross-scenario validation in the Douban Movie scenario. This scenario does not contain a multi-event crisis sequence such as RUW, so this paper uses the release date of *Avengers: Age of Ultron* (Apr. 22, 2015) as a retrospective anchor to examine the correspondence between warning timelines and propagation launch stages.

Fig. 2 presents the timeline results for the movie scenario, demonstrating that SRD generates a warning response before the release date, with the most recent valid warning occurring 5 days in advance. This result indicates that the method is not dependent on crisis contexts themselves, but can respond to anomalous accelerations in social propagation processes relative to background rhythms.

Meanwhile, compared to the RUW scenario, the number of warning periods in the movie discussion sequence

significantly increases. This indicates that in cultural consumption scenarios characterized by multi-wave diffusion and stronger background fluctuations, propagation acceleration signals may be more frequent, resulting in higher alarm density. Thus, the output of SRD reflects not only potential risks but also the inherent vitality of propagation dynamics.

It must be emphasized that these cross-scenario results provide supplementary evidence regarding the method's transferability, rather than an event-level evaluation equivalent to the RUW scenario. Because only one movie is used in the Douban validation, this experiment should be regarded as illustrative rather than statistically conclusive. The Douban Movie experiment demonstrates that the dynamic anomalies characterized by SRD are not confined to sudden crisis contexts; however, its statistical performance in different scenarios

remains subject to background fluctuation structures and task definitions, which warrants further validation in more systematic, multi-object evaluations.

7 Discussion

The RUW main experimental results indicate that SRD's effectiveness does not primarily stem from direct responses to intensity levels, but from characterizing the relative acceleration phases of the propagation process. Unlike statistical detection methods based on absolute level deviations, the constructed Doppler indicator f_d describes the normalized deviation of the propagation change rate relative to background fluctuations, and judgements are made using causal rolling quantile thresholds referenced against local history. This shifts




Fig. 1 RUW scenario analysis: the normalized propagation-rate indicator $f_d(t)$, causal rolling quantile thresholds, and generated warning points under scene-level robust scaling with a 15-day rolling window ($W = 15$). P90, P95, and P97 denote the rolling 90th-, 95th-, and 97th-percentile thresholds of $f_d(t)$, respectively; E1–E3 denote the Maxar satellite images, recognition of Donetsk and Luhansk (D/L), and Istanbul talks, respectively. Vertical dashed lines denote event anchors




Fig. 2 Cross-scenario validation on the Douban Movie dataset. The release date serves as the retrospective anchor, and shaded regions indicate warning intervals

the focus of warning triggers toward whether the propagation rhythm exhibits structural deviation, rather than whether the current level has already reached an anomalous state. Results show that this mechanism can identify key acceleration phases prior to event formation, thereby achieving a relatively balanced performance across event coverage, warning lead time, and false alarm control.

From a broader modeling perspective, this result highlights a critical distinction: existing methods predominantly use “whether the current state is anomalous” as their core criterion, whereas in early warning scenarios, the more crucial question is “whether the system is deviating from its conventional evolutionary trajectory.” SRD’s performance demonstrates that modeling early warning as a propagation dynamics change detection problem helps secure practically meaningful early response capabilities while maintaining low false alarm interference.

Unlike CUSUM, which relies on accumulated deviations from a reference state, and EARS, which primarily reacts to abrupt increases in observed activity, SRD is designed to capture temporally consistent acceleration patterns in collective behavioral dynamics. This distinction in detection philosophy may contribute to the earlier warning tendencies observed across multiple evaluation scenarios. By emphasizing sustained acceleration rather than isolated spikes or cumulative level shifts, SRD is able to respond to emerging propagation dynamics before large-scale activity surges become fully observable.

The impact of window length on the results further uncovers the temporal structural characteristics of the RUW scenario. Compared to SRD ($W = 30$), SRD ($W = 15$) performs better in Recall and F1 while maintaining a stable Precision, suggesting that early signals in this scenario are closer to localized phased accelerations rather than long-term trend drifts. A shorter window preserves the contrast between localized changes and historical background, whereas a longer window more easily absorbs short-term acceleration segments into the overall distribution, thereby weakening the persistence of triggering conditions.

Furthermore, the results indicate that SRD is not equally sensitive to all types of event evolution processes. Its effectiveness relies heavily on specific propagation characteristics: first, the presence of identifiable local acceleration phases; second, a significant deviation of this acceleration relative to background fluctuations; third, a certain temporal persistence of the acceleration process. Consequently, for scenarios characterized by slow accumulation, a lack of abrupt change rate mutations, or high noise interference, the method’s responsiveness may be constrained. This phenomenon reflects SRD’s mechanistic boundaries regarding warning targets—specifically, it is better suited to identify phases where propagation dynamics undergo structural shifts, rather than to cover all types of macroscopic event evolution.

The cross-scenario results from Douban Movie further indicate that the dynamic anomalies characterized by SRD are not limited to sudden crisis contexts; they can similarly generate responses during propagation launch phases. However, the significant increase in warning periods within this scenario suggests that in environments with multi-wave diffusion and stronger background fluctuations, alarm density will rise with

variations in propagation rhythms. It is important to note that these results primarily provide preliminary evidence of the method’s transferability, and its statistical stability requires further validation in larger-scale, multi-scenario settings.

Overall, the advantage of SRD lies not in achieving comprehensive coverage of all events, but in providing stable responses to propagation phases with structural acceleration features at the cost of low quiet-day interference. This outcome indicates that in social media early warning tasks, the core issue is not determining whether a system is already anomalous, but recognizing whether it is deviating from its conventional rhythm. Therefore, modeling early warning as a propagation dynamics change detection problem offers a more explanatory analytical perspective for understanding risk formation processes in complex social systems.

8 Conclusions

This paper proposes an unsupervised online framework for social media early warning—SRD—to identify anomalous acceleration phases of propagation relative to historical rhythms. Under strict causal constraints, SRD uniformly maps heterogeneous social behaviors into a social intensity sequence and achieves online detection of propagation dynamics deviations through change rate modeling, robust scale normalization, and a rolling quantile threshold mechanism.

Experimental results on the RUW scenario and the Douban Movie cross-scenario validation demonstrate that SRD can respond to key acceleration phases before event formation with low false alarm interference, achieving a relatively balanced performance among event coverage, warning lead time, and alarm stability. Furthermore, the results show that this method is more applicable to propagation processes with structural acceleration features, while its responsiveness has certain boundaries for scenarios with slow evolution or weak rhythmic changes.

More generally, the findings of this paper suggest that effective social media early warning may benefit from focusing on deviations in propagation dynamics rather than solely on observed activity levels or state anomalies. Compared to traditional detection methods based on intensity levels or state assessments, modeling early warnings as a propagation dynamics change detection problem corresponds more directly to the process characteristics of risk formation, thereby enhancing the practical utility of warnings while reducing false alarm interference.

Future work could further incorporate richer signal representations, real-time analysis agents (Han et al., 2025), and multi-source information fusion mechanisms, while preserving the core dynamic modeling framework, to enhance the method’s applicability and robustness in complex real-world environments. For the non-crisis validation setting, future work will include more movies and diverse genres to provide more systematic evidence of cross-scenario generalizability.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No. 2023YFB3209801) and the National Natural Science Foundation of China (No. 52402490).

Author contributions

Hetian CHEN and Lili FAN designed the study. Hetian CHEN and Guangyu SHI processed the data. Hetian CHEN drafted the paper. Lili FAN and Guangyu SHI revised the paper. All the authors finalized the paper.

Conflict of interest

All the authors declare that they have no conflict of interest.

Data availability

The RUW and Douban Movie Short Comments dataset analyzed in this study are publicly available from the sources cited in Section 4. The processed daily time series and analysis results are available from the corresponding author upon reasonable request.

Declaration on the use of generative AI tools

During the preparation of this work, the authors did not use any generative AI tools. All content was independently created, written, and revised by the authors, who take full responsibility for the integrity and accuracy of the paper.

References

- Alam F, Sajjad H, Imran M, et al., 2021. CrisisBench: benchmarking crisis-related social media datasets for humanitarian information processing. *Proc 15th Int AAAI Conf on Web and Social Media*, p.923-932. <https://doi.org/10.1609/icwsm.v15i1.18115>
- Bian T, Xiao X, Xu TY, et al., 2020. Rumor detection on social media with bi-directional graph convolutional networks. *Proc 34th AAAI Conf on Artificial Intelligence*, 34(1):549-556. <https://doi.org/10.1609/aaai.v34i01.5393>
- Chung-I, 2020. Douban-Sentiment-Analysis. <https://github.com/Chung-I/Douban-Sentiment-Analysis> [Accessed on June 26, 2026].
- Diao QM, Jiang J, 2013. A unified model for topics, events and users on Twitter. *Proc Conf on Empirical Methods in Natural Language Processing*, p.1869-1879. <https://doi.org/10.18653/v1/D13-1192>
- Fung YR, Ji H, 2022. A Weibo dataset for the 2022 Russo-Ukrainian crisis. <https://doi.org/10.48550/arXiv.2203.05967>
- Gallotti R, Valle F, Castaldo N, et al., 2020. Assessing the risks of 'infodemics' in response to COVID-19 epidemics. *Nat Hum Behav*, 4(12):1285-1293. <https://doi.org/10.1038/s41562-020-00994-6>
- Gui JY, Ma ZN, Zhou H, et al., 2025. Deep anomaly detection of temporal heterogeneous data in AIOps: a survey. *Front Inform Technol Electron Eng*, 26(9):1551-1576. <https://doi.org/10.1631/FITEE.2400467>
- Han SJ, Zhang JS, Shen YQ, et al., 2025. FinSphere: a real-time stock analysis agent with instruction-tuned large language models and domain-specific tool integration. *Front Inform Technol Electron Eng*, 26(10):1822-1831. <https://doi.org/10.1631/FITEE.2500414>
- Imran M, Castillo C, Diaz F, et al., 2015. Processing social media messages in mass emergency: a survey. *ACM Comput Surv*, 47(4):67. <https://doi.org/10.1145/2771588>
- Lasswell HD, 1948. The structure and function of communication in society. In: Bryson L (Ed.), *The Communication of Ideas*. Harper and Row, New York, p.37-51.
- Lau JH, Collier N, Baldwin T, 2012. On-line trend analysis with topic models: #Twitter trends detection topic model online. *Proc COLING*, p.1519-1534.
- Li X, Gao RJ, Wang R, 2025. Bursty event detection in social networks by integrating word features and Word2Vec model. *Proc 4th Int Conf on Intelligent Systems, Communications and Computer Networks*, p.300-306. <https://doi.org/10.1145/3732945.3732989>
- Liu JF, Wu YC, Liu ZX, 2025. A focused crawling strategy based on comprehensive priority evaluation of hyperlinks and improved Bayesian classifier. *Front Inform Technol Electron Eng*, 26(12):2569-2582. <https://doi.org/10.1631/FITEE.2400939>
- Liu YS, Chen Y, Zhang JY, et al., 2024. Toward an accurate mobility trajectory recovery using contrastive learning. *Front Inform Technol Electron Eng*, 25(11):1479-1496. <https://doi.org/10.1631/FITEE.2300647>
- Ma CB, Wang YX, Wu J, et al., 2025. Explicit and implicit data augmentation for social event detection. *Proc 63rd Annual Meeting of the Association for Computational Linguistics*, p.8402-8415. <https://doi.org/10.18653/v1/2025.acl-long.412>
- Ma J, Gao W, Wong KF, 2017. Detect rumors in microblog posts using propagation structure via kernel learning. *Proc 55th Annual Meeting of the Association for Computational Linguistics*, p.708-717. <https://doi.org/10.18653/v1/P17-1066>
- Mathieu J, Fulk M, Lorber M, et al., 2012. Social Radar Workflows, Dashboards, and Environments. RTO-MP-HFM-201, The MITRE Corporation.
- Parekh T, Mac A, Yu JR, et al., 2024. Event detection from social media for epidemic prediction. *Proc Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, p.5758-5783. <https://doi.org/10.18653/v1/2024.naacl-long.322>
- Peng SC, Cao LH, Wang GJ, et al., 2025. A survey on textual emotion cause extraction in social networks. *Dig Commun Netw*, 11(2):524-536. <https://doi.org/10.1016/j.dcan.2024.07.004>
- Peng XY, Wu JR, Liu RM, et al., 2025. Rumor detection on social media with temporal propagation structure optimization. *Proc 31st Int Conf on Computational Linguistics*, p.3865-3878. <https://doi.org/10.48550/arXiv.2412.06208>
- Qiu ZT, Ma CB, Wu J, et al., 2025. Text is all you need: LLM-enhanced incremental social event detection. *Proc 63rd Annual Meeting of the Association for Computational Linguistics*, p.4666-4680. <https://doi.org/10.18653/v1/2025.acl-long.233>
- Ramakrishnan N, Butler P, Muthiah S, et al., 2014. 'Beating the news' with EMBERS: forecasting civil unrest using open source indicators. *Proc 20th ACM SIGKDD Int Conf on Knowledge Discovery and Data Mining*, p.1799-1808. <https://doi.org/10.1145/2623330.2623373>
- Schramm W, Roberts DF, 1954. How communication works. In: Schramm WL (Ed.), *The Process and Effects of Mass Communication*. University of Illinois Press, Urbana, p.3-26.
- Tang BH, Lin BD, Yan HL, et al., 2024. Leveraging generative large language models with visual instruction and demonstration retrieval for multimodal sarcasm detection. *Proc Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, p.1732-1742. <https://doi.org/10.18653/v1/2024.naacl-long.97>
- Utmhikari, 2019. Douban Movie Short Comments Dataset. <https://www.kaggle.com/datasets/utmhikari/doubanmovieshortcomments> [Accessed on June 26, 2026].
- Wang FY, 2007. Toward a paradigm shift in social computing: the ACP approach. *IEEE Intell Syst*, 22(5):65-67. <https://doi.org/10.1109/MIS.2007.4338496>
- Weng JS, Lee BS, 2011. Event detection in Twitter. *Proc 5th Int AAAI Conf on Web and Social Media*, p.401-408. <https://doi.org/10.1609/icwsm.v5i1.14102>
- Xie W, Zhu FD, Jiang J, 2016. TopicSketch: real-time bursty topic detection from Twitter. *IEEE Trans Knowl Data Eng*, 28(8):2216-2229. <https://doi.org/10.1109/TKDE.2016.2556661>
- Xiong SF, Zhang GP, Fan XB, 2025. MAL: multilevel active learning with BERT for Chinese textual affective structure analysis. *Front Inform Technol Electron Eng*, 26(6):833-846. <https://doi.org/10.1631/FITEE.2400242>
- Zhang YZ, Wang MY, Wu YX, et al., 2025. DialogueLLM: context and emotion knowledge-tuned large language models for emotion recognition in conversations. *Neur Netw*, 192(1):107901. <https://doi.org/10.1016/j.neunet.2025.107901>
- Zhou F, Xu X, Trajcevski G, et al., 2021. A survey of information cascade analysis: models, predictions, and recent advances. *ACM Comput Surv*, 54(2):27. <https://doi.org/10.1145/3433000>