

Relational Prompt-based Pre-trained Language Models for Social Event Detection

PU LI*, Kunming University of Science and Technology, China

XIAOYAN YU*, Beijing Institute of Technology, China

HAO PENG, Beihang University, China

YANTUAN XIAN[†], Kunming University of Science and Technology, China

LINQIN WANG, Kunming University of Science and Technology, China

LI SUN, North China Electric Power University, China

JINGYUN ZHANG, Beihang University, China

PHILIP S. YU, University of Illinois at Chicago, USA

Social Event Detection (SED) aims to identify significant events from social streams, and has a wide application ranging from public opinion analysis to risk management. In recent years, Graph Neural Network (GNN) based solutions have achieved state-of-the-art performance. However, GNN-based methods often struggle with missing and noisy edges between messages, affecting the quality of learned message embedding. Moreover, these methods statically initialize node embedding before training, which, in turn, limits the ability to learn from message texts and relations simultaneously. In this paper, we approach social event detection from a new perspective based on Pre-trained Language Models (PLMs), and present RPLM_{SED} (Relational prompt-based Pre-trained Language Models for Social Event Detection). We first propose a new pairwise message modeling strategy to construct social messages into message pairs with multi-relational sequences. Secondly, a new multi-relational prompt-based pairwise message learning mechanism is proposed to learn more comprehensive message representation from message pairs with multi-relational prompts using PLMs. Thirdly, we design a new clustering constraint to optimize the encoding process by enhancing intra-cluster compactness and inter-cluster dispersion, making the message representation more distinguishable. We evaluate the RPLM_{SED} on three real-world datasets, demonstrating that the RPLM_{SED} model achieves state-of-the-art performance in offline, online, low-resource, and long-tail distribution scenarios for social event detection tasks.

*Equal Contributions.

[†]Corresponding author.

Authors' addresses: P. Li, Faculty of Information Engineering and Automation, and Yunnan Key Laboratory of Artificial Intelligence, Kunming University of Science and Technology, Kunming 650500, China; email: lip@stu.kust.edu.cn; X. Yu, School of Computer Science and Technology, Beijing Institute of Technology, Beijing 100081, China; email: xiaoyan.yu@bit.edu.cn; H. Peng, School of Cyber Science and Technology, Beihang University, No. 37 Xue Yuan Road, Haidian District, Beijing, 100191, China; email: penghao@buaa.edu.cn; Y. Xian, Faculty of Information Engineering and Automation, and Yunnan Key Laboratory of Artificial Intelligence, Kunming University of Science and Technology, Kunming 650500, China; email: xianyt@kust.edu.cn; L. Wang, Faculty of Information Engineering and Automation, and Yunnan Key Laboratory of Artificial Intelligence, Kunming University of Science and Technology, Kunming 650500, China; email: linqinwang7767@163.com; L. Sun, School of Control and Computer Engineering, North China Electric Power University, Beijing 102206, China; email: ccesunli@ncepu.edu.cn; J. Zhang, School of Cyber Science and Technology, Beihang University, No. 37 Xue Yuan Road, Haidian District, Beijing, 100191, China; email: zhangjingyun@buaa.edu.cn; P. S. Yu, Department of Computer Science, University of Illinois at Chicago, Chicago 60607, IL; email: psyu@uic.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2024/9-ART

<https://doi.org/0000001.0000001>

, Vol. 1, No. 1, Article . Publication date: September 2024.

2 • P. Li et al.

CCS Concepts: • **Information systems** → **Web searching and information discovery**; **Data mining**; *Social networks*; • **Computing methodologies** → *Artificial intelligence*; *Neural networks*.

Additional Key Words and Phrases: Social event detection, pre-trained language models, prompt learning, clustering constraint.

ACM Reference Format:

Pu Li, Xiaoyan Yu, Hao Peng, Yantuan Xian, Linqin Wang, Li Sun, Jingyun Zhang, and Philip S. Yu. 2024. Relational Prompt-based Pre-trained Language Models for Social Event Detection. 1, 1 (September 2024), 43 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

Social events refer to the various incidents or activities occurring in reality that elicit widespread public interest and discourse across platforms such as Twitter, Weibo, Facebook, Tumblr, Telegram, etc. These events encompass a broad spectrum of topics, including news reports, sudden public incidents, social trends, market dynamics, and more [15]. Social Event Detection (SED) aims to group messages reporting on the same incident to identify significant events within them from social streams [6, 38]. Owing to its capability to provide deep insights into public issues and trends, it is extensively applied in fields such as public opinion analysis [59], financial market analysis [55], risk management [46], and political analysis [49].

Social data's incompleteness, ambiguity, and streaming nature present more significant challenges for social event detection than traditional text mining tasks. Early social event detection methods typically focused on utilizing text contents [3, 54, 69, 77, 81, 89] or attributes extracted from texts [22, 79, 80] to identify and categorize events. These methods focus on categorizing messages into different events but fail to detect events from social streams. In recent years, Graph Neural Network (GNN) based social event detection methods have demonstrated tremendous potential [11, 16, 57, 59, 60, 65, 68]. These methods typically model social messages into graphs to integrate semantic and structural information. By leveraging GNNs to learn message representations from graphs and employing clustering techniques for event detection, these approaches significantly enhance the performance of social event detection. For instance, Peng et al. [57] construct social messages as a Heterogeneous Information Network (HIN) and, upon obtaining representations through Graph Convolutional Networks (GCN), employ a pairwise sampling method to address the challenges presented by the numerous event categories and the sparse number of instances per category. Cui et al. [16] respectively construct semantic and temporal views by leveraging hashtags and time distribution of messages, then integrate representations from semantic and temporal views using a hashtag-based multi-view graph attention mechanism. Cao et al. [11] construct messages as a HIN graph and then transform it into a homogeneous graph, utilizing the learning and inductive capabilities of GNNs to achieve incremental event detection. Moreover, Peng et al. [60], recognizing that different relations between messages may have varying degrees of value, models social messages into a weighted multi-relational graph, employing reinforcement learning to determine the optimal thresholds for various relations.

However, GNN-based social event detection methods still exhibit the following limitations.

- (1) Graphs constructed from common attributes (e.g., entities, hashtags, users, etc.) between messages face issues of missing and noisy edges [45]. These issues are primarily manifested in the lack of connections between message nodes within the same event and the existence of connections between nodes from different events. Whether modeling social messages as HIN graph (Figure 1(a)), homogeneous graph (Figure 1(b)), or multi-relational graph (Figure 1(c)), none of these approaches can handle well above issues. This challenge is multifaceted. On the one hand, intra-event messages may have no connections. This impedes GNNs from capturing associations between these nodes, resulting in an incomplete representation of related information. On the other hand, attribute co-occurrence is probable for inter-event messages, potentially introducing noise and impacting the quality of node representation. Besides, GNN-based methods fail to transmit information adequately in sparse graph scenarios, leading to suboptimal message representations.

, Vol. 1, No. 1, Article . Publication date: September 2024.

- (2) GNN-based methods initialize node embeddings statically via pre-trained word embedding models [9, 52, 62] before training. This practice prevents the acquisition of better semantic embeddings from the message content in subsequent learning processes [82]. Meanwhile, GNNs inherently separate the utilization of relations/edges and message content when learning message embeddings from explicit message graphs. The role of edges primarily select nodes for aggregation rather than directly participating in the processing of message representations. This approach may not fully consider the complexity of relations between messages and the potential interactions between these relations and message content [17].
- (3) GNN-based methods exhibit insufficient constraints on the training process, treating messages merely as targets to be pulled closer or pushed apart [27].




Fig. 1. Social message modeling methods. **(a) HIN graph** models social messages into the HIN graph by extracting attributes from the messages (e.g., entities, hashtags, users, etc.) as relational nodes connecting to the related message nodes. **(b) Homogeneous graph** is obtained by mapping a HIN with message nodes retained and adding edges between message nodes that share common adjacent relation nodes. **(c) Multi-relational graph**: Similar to (b), but transforms relation nodes into different types of edges between message nodes. **(d) Message pairs**: Converting multi-relational edges in the multi-relational graph into multi-relational prompts. Regardless of the actual co-occurrence information between messages, corresponding multi-relational prompts will exist.

Herein, to circumvent the aforementioned challenges of GNN-based methods, we put forward a perspective of learning expressive message representations with Pre-trained Language Models (PLMs), instead of GNNs. Thus, we present a novel **Relational prompt-based Pre-trained Language Model for Social Event Detection**, namely $RPLM_{SED}$. Concretely, we **first** propose a new pairwise message modeling strategy to model social messages. As illustrated in Figure 1(d), we sample other messages for each message to form pairs and convert the multiple relations between the two into a multi-relational sequence to preserve structural information. For any two messages, each type of relation between them and the existence or non-existence of these relations can be comprehensively represented as a multi-relational sequence. This strategy addresses the issues of missing and noisy edges in the message graph and achieves the joint modeling of the messages' structure and semantics. **Secondly**, a new multi-relational prompt-based pairwise message learning mechanism is proposed to learn more detailed and comprehensive message representations from semantic and structural information between messages. It breaks away from traditional methods that rely on explicit relational graph learning and aggregation of message representations, and dividing message representation learning and aggregation into two independent parts. We transform the multi-relational sequences into multi-relational prompt embeddings and integrate them with the content embeddings of messages. By consistently encoding the embeddings of message pairs along with their corresponding multi-relational prompt embeddings through PLMs [30, 43, 47, 87], we achieve the simultaneous utilization of message content and structural information, as well as the dynamic embedding and encoding of message content. To obtain more robust and stable message representations from their corresponding message pairs, a similarity-based message representation aggregation approach is introduced to aggregate the representations extracted from message pairs with high confidence levels, serving as the final representation of the

message. This approach also avoids the over-smoothing of representations and enhances the similarity between representations within the same event. **Finally**, we design a new clustering constraint encompassing inter-cluster and intra-cluster losses. Unlike previous optimization methods that solely consider messages as targets, the objective of the clustering constraint optimization is focused on event clusters. The inter-cluster loss aims to distance the cluster centers of different events, whereas the intra-cluster loss ensures that messages of the same event are gathered close to their respective cluster center. This clustering constraint strengthens the cohesiveness within clusters and the separation between them, thereby effectively improving the distinguishability of the message representations. By leveraging the inherent abilities of PLMs to learn, retain, and expand knowledge, we achieve incremental social event detection.

We conduct comprehensive experiments on three datasets, Events2012 [51], Events2018 [50], and Arabic-Twitter [2]. Analyses and comparisons are conducted with one community detection method [39], eight GNN-based methods [11, 16, 57, 60, 65–68], and seven pre-training-based methods [30, 43, 52] across online, offline, low-resource, and long-tail distribution scenarios. In all scenarios, $RPLM_{SED}$ demonstrates SOTA performance compared to all benchmark methods. Additionally, we compare and analyze the performance of $RPLM_{SED}$ based on different PLMs and report the time consumption for each part of $RPLM_{SED}$. Finally, we analyze the performance of three density-based clustering algorithms [4, 10, 19] on the representations obtained by $RPLM_{SED}$. The codes for all baseline models and $RPLM_{SED}$, along with all datasets, are publicly accessible on GitHub¹.

In summary, the contributions of this paper are as follows:

- We present a novel relational prompt-based pre-trained language model for social event detection ($RPLM_{SED}$). To the best of our knowledge, this is the first SED method based on PLM that integrates both semantic information from social messages and structural relations between them. It breaks away from traditional GNN-based methods, no longer relying on explicit structural relations between messages to learn representations.
- A new pairwise message modeling strategy is proposed to solve the issues of missing and noisy edges in the message graph, while simultaneously achieving joint modeling of message content and structural information.
- A new multi-relational prompt-based pairwise message learning mechanism is proposed to learn more detailed and comprehensive message representations from message pairs with multi-relational prompts. It leverages the powerful semantic understanding and context encoding capabilities of PLMs, consistently encoding both message text pairs and their corresponding multiple relations, thus enabling the simultaneous utilization of textual content and structural information.
- A new clustering constraint is designed to enhance the discriminability of message representations. This clustering constraint loss pushes the cluster centers of different events apart and constrains messages from the same event to be closer to their corresponding cluster center.
- Comprehensive experiments and analyses have been conducted on three datasets, illustrating that the proposed approach surpasses the current SOTA methods for social event detection regarding effectiveness.

The structure of this paper is as follows: Section 2 outlines the background and preliminaries of our work. In Section 3, we describe the technical details of the proposed model, named $RPLM_{SED}$. Section 4 presents the experimental setup, and Section 5 discusses the experiment's results. Section 6 provides an overview of related works. Finally, we conclude the paper in Section 7.

¹https://github.com/RingBDStack/RPLM_SED

2 BACKGROUND AND OVERVIEW

This section first provides an overview of the problems and challenges encountered in social event detection. Subsequently, we detail the core definitions related to social event detection. Additionally, Table 1 presents a comprehensive list of the primary symbols used throughout this paper.

Table 1. Glossary of notations.

Notation	Description
S	A social stream
$M; m$	A message block; A message in a message block or event
E_i	Set of events from message block M_i
e_k	k -th event in E_i
\mathcal{G}	A multi-relational message graph
\mathcal{M}	The set of all nodes in \mathcal{G}
w	The window size for maintaining the model
$\theta; \theta^{plm}$	The parameters of model; The parameters with the PLM section of model
ζ	The parameter update weight during model maintenance
\mathcal{R}_t	Edge set with relational type t in \mathcal{G}
$r_{i,s}^t$	An edge between m_i and sampled node m_s with relational type t
$r_{i,s}$	The multi-relational sequence between m_i and sampled node m_s
$D; d$	Preprocessed message blocks; A message pair
o	The number of heads in structured attention mechanism
$\hat{x}_i; \hat{x}_s$	The representations of m_i and sampled node m_s obtained through structured attention mechanism and averaging
q_i	The similarity between m_i and sampled node m_s
$X_i; Q_i$	A set of candidate representations of m_i ; A set of similarities corresponding to X_i
I_i	A set of indices satisfying the similarity score
α	The similarity threshold
\bar{x}_i	The final message representation of m_i
B	The training batch size
$H; h$	The center feature representation matrix; The central feature representation
L_{pair}	The pairwise cross-entropy loss
L_{intra}	The intra-cluster loss
L_{inter}	The inter-cluster loss
L	Total loss
$\kappa; \lambda; \mu$	Similarity pairwise cross-Entropy loss weight; Intra-cluster loss weight; Inter-cluster loss weight

2.1 Problem and Challenges

When addressing social event detection tasks, the most straightforward approach is to learn representations of social messages and subsequently cluster these representations. To acquire higher-quality message representations for clustering, social event detection faces the following three main challenges:

Challenge 1: How to effectively address the issues of missing and noisy edges in the social message graph for social event detection?

The issues of missing and noisy edges in the social message graph are multifaceted. They include the lack of direct connections between nodes of the same event, connections between nodes of different events, and the presence of isolated nodes. The emergence of these issues can be attributed to various factors, such as describing

the same event using different vocabulary or sentences, an individual user expressing opinions on other events, the event having just occurred, or being rarely discussed. However, GNN-based methods [3, 22, 54, 69, 77, 79–81, 89] learn message representations from explicit message graph. Noisy edges in the graph can adversely affect the learned message representations, while the absence of effective edges may result in incomplete message representations. Additionally, isolated nodes within the graph cannot be effectively processed. Hence, effectively addressing the issues of missing and noisy edges in the social message graph remains one of the challenges in social event detection.

Challenge 2: How to concurrently leverage the content and structural information of messages to learn their representations?

GNN-based methods typically initialize message embeddings statically before training, which precludes the acquisition of better semantic embeddings from message content during the training process. In addition, the representation aggregation process of GNNs can be comprehended as being guided by the edges/relations to determine which message nodes are selected for subsequent aggregation. However, the edges themselves do not directly participate in the processing of message representations. These approaches inherently separate the utilization of message content from structural information, failing to adequately account for the complexity of inter-message relations and the potential interactions between these relations and message content. Therefore, exploring a new learning mechanism to address the complexities of social streams is imperative, enabling a more effective utilization of both message content and structural information.

Challenge 3: How to obtain more distinguishable message representations for social event detection?

In social event detection, events are detected through clustering algorithms after obtaining message representations. Both distance-based [48] and density-based [10, 19] clustering algorithms fundamentally rely on the Euclidean distance between vectors. Ensuring sufficient closeness among messages belonging to the same event while maintaining adequate distance from messages of different events can further improve the accuracy of event detection. Therefore, enhancing the discriminability of message representations is essential. However, some common optimization methods, such as the triplet loss adopted in KPGNN [11] and FinEvent [60], are designed to select one positive and one negative sample for each message. The goal is then to ensure that the distance between each message and its positive sample is significantly smaller than its distance to the negative sample. Alternatively, the pairwise loss proposed in QSGNN [65], which constrains the distance between all negative sample pairs to be greater than that of positive samples, offers more stringent constraints compared to triplet loss. Both triplet loss and pairwise loss approaches treat messages as objects to be pushed away or drawn closer, which may not effectively differentiate all event clusters nor guarantee compactness within event clusters. Thus, it is necessary to design a more effective optimization method to enhance the distinguishability of message representations and thereby improve event detection accuracy.

2.2 Problem Definition

In this section, we formalize the definitions of social stream, social event, multi-relational message graph, social event detection, and incremental social event detection as follows:

Definition 2.1. $S = M_0, M_1, \dots, M_{i-1}, M_i, \dots$ represents a **social stream** consisting of continuous time-ordered social message blocks. **Message block** $M_i = \{m_j \mid 1 \leq j \leq |M_i|\}$ consists all of messages in a time block $[t_i, t_{i+1})$, where $|M_i|$ is the total number of messages in M_i , and m_j is a message within M_i .

Definition 2.2. A **social event** denoted as $e = \{m_j \mid 1 \leq j \leq |e|\}$ constitutes a collection of interrelated social messages addressing a common real-world occurrence. We assume that each social message is affiliated with at most one event.

Definition 2.3. We define a **multi-relational message graph** as $\mathcal{G} = \{\mathcal{M}, \{\mathcal{R}_t\}_{t=1}^T\}$, where \mathcal{M} is a set of nodes representing social messages, $\{\mathcal{R}_t\}_{t=1}^T$ are sets of multi-relational edges with type t . $r_{i,j}^t \in \mathcal{R}_t$ denotes an edge/relation between message m_i and m_j with relational type t . Relational types may include occurrences of entities, hashtags, user mentions, or message pairs posted within a predefined period.

Definition 2.4. Given a message block M_i , an algorithm for **social event detection** learns a model $f(M_i; \theta) = E_i$, where $E_i = \{e_k \mid 1 \leq k \leq |E_i|\}$ represents a set of events within M_i . θ denotes the parameters associated with the model f .

Definition 2.5. Given a social stream S , an **incremental social event detection** algorithm learns a series of event detection models $f_0, \dots, f_{t-w}, f_t, f_{t+w}, \dots$, such that $f_t(M_i; \theta_t, \theta_{t-w}) = E_i$ is the model for all message blocks in $\{M_i \mid t+1 \leq i \leq t+w\}$. Here, $E_i = \{e_k \mid 1 \leq k \leq |E_i|\}$ is a set of events contained in the message block M_i , w is the window size for continuously updating the model, θ_t and θ_{t-w} are the parameters of f_t and f_{t-w} , respectively. Note that f_t inherits part of the parameters θ_{t-w} from its predecessor f_{t-w} , and it extends and updates the knowledge learned. Specifically, f_0 , which does not extend any previous model, is referred to as the initial model.

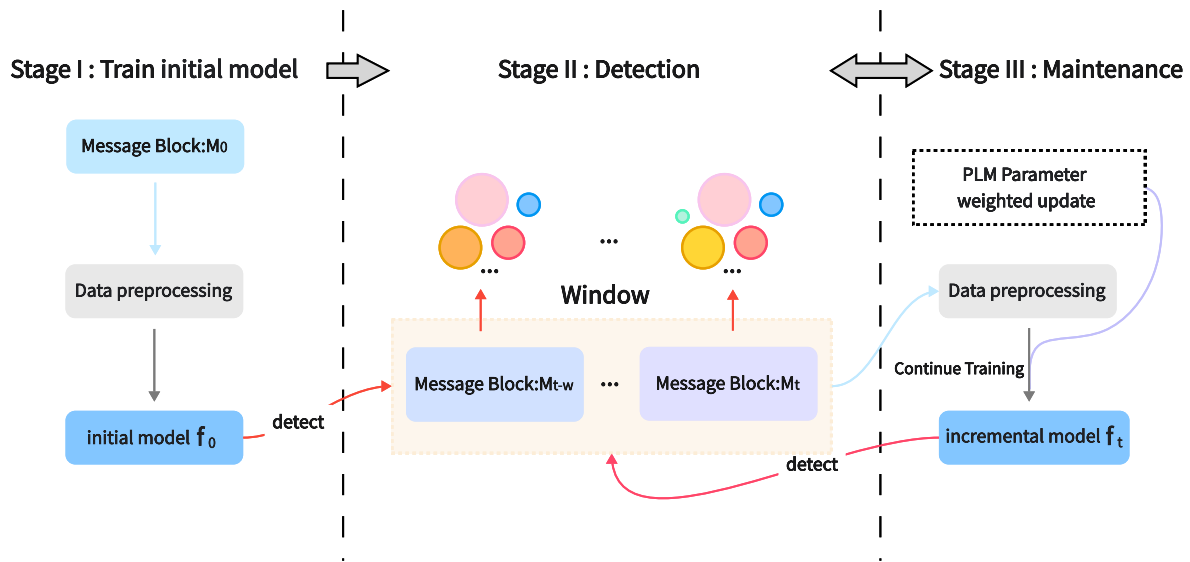


Fig. 2. Incremental life-cycle in HP-Event architecture. **Stage I.** The initial data block M_0 is utilized to train the initial model. **Stage II.** The model detects events within each message block of the window separately. **Stage III.** The model is maintained through the current window's message blocks. Then, return to the detection phase to detect events of each message block in the next window.

3 METHODOLOGY

This section provides a detailed exposition of the proposed method $RPLM_{SED}$. We elucidate the lifecycle of the $RPLM_{SED}$ model in Section 3.1. To leverage messages' structure and semantics, we propose a new Pairwise Message Modeling Strategy in Section 3.2. Section 3.3 and Section 3.4 introduce the Multi-relational Prompt-based

Pairwise Message Learning Mechanism, and loss function, respectively. Additionally, we provide and explain the overall algorithmic process in Section 3.5 and analyze the time complexity of $RPLM_{SED}$ in Section 3.6.

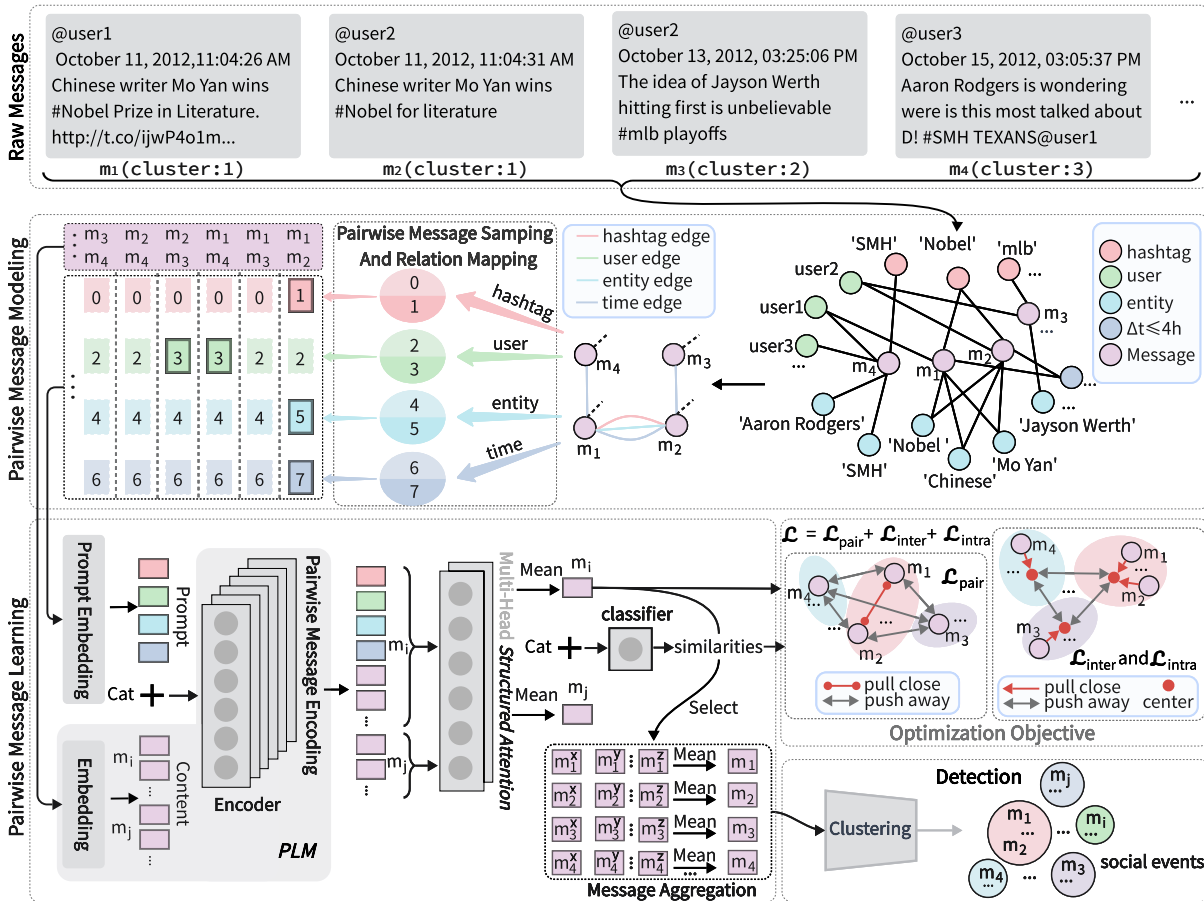


Fig. 3. The architecture of the $RPLM_{SED}$. **Raw Message** shows real messages from Twitter. **Pairwise Message Modeling** (Section 3.2) initially constructs a multi-relational message graph, then sample message to form pairs and map relations to the multi-relational sequence. **Pairwise Message Learning** (Section 3.3) embeds and encodes message content and relational sequence, then extracts message representations from the encoding vector. the messages' similarity is assessed by the classifier. In the detection phase, candidate representations of each message are **selected** to **aggregate** the final representations. **Optimization Objective** (Section 3.4) comprises pairwise cross-entropy, inter-cluster and intra-cluster losses, jointly guiding model training.

3.1 Incremental event detection life-cycle

As shown in Figure 2, to detect events from social streams incrementally, we propose a three-stage event detection framework, including initial model training, event detection, and model maintenance:

- (1) **Initial model training:** we train the initial social event detection model with the first message block M_0 using the proposed $RPLM_{SED}$.

- (2) **Event detection:** event detection is performed on incoming message blocks M_{t+1}, \dots, M_{t+w} with the trained model f_t . Specifically, we use model f_t to detect events in each message block separately and evaluate the results.
- (3) **Model maintenance:** the model enters the maintenance stage after each event detection stage. We first perform a weighted fusion of the parameters of PLM in RPLM with the initial parameters of PLM, thereby striking a balance between retaining some existing knowledge and assimilating new knowledge. Subsequently, we use message blocks M_{t+1}, \dots, M_{t+w} to maintain the model f_{t+w} .

3.2 Pairwise Message Modeling Strategy

To jointly model social messages' textual content and structural information, we model social messages into message pairs with multi-relational sequences. **Firstly**, we construct a Heterogeneous Information Network (HIN) [57, 60], where each social message is treated as a message node. Concurrently, we extract a set of named entities, hashtags, and users (including users and mentioned users) as relational nodes from each message. Edges between the message nodes and the corresponding relational nodes denote their associations. For instance, the message m_1 depicted in the 'Raw Messages' part of Figure 3 includes three entities ('Nobel', 'Mo Yan', 'Chinese'), one user ('user1'), and one hashtag ('Nobel'). Edges are established between m_1 and its related nodes. By merging duplicate nodes, a HIN graph is formed. In this graph, message nodes, entity nodes, user nodes, and hashtag nodes are represented by m , e , u , and h , respectively. **Subsequently**, the relational nodes in the heterogeneous information graph are converted into different types of edges between message nodes, thereby constructing a multi-relational message graph that includes single-type nodes and multiple-type relational edges. Additionally, edges are added between message nodes published within a 4-hour timeframe to capture temporal correlations. **Finally**, after establishing the multi-relational message graph, we further implement pairwise modeling for message nodes within the graph. For any given message node m_i , we sample y positive and negative samples to form message pairs in the training phase and sample n nodes in the detection phase, where n is three times y . In the training stage, we prefer to select message nodes of the same event with richer connections to m_i as positive sampling message nodes, whereas messages from different events with sparser connections to m_i are chosen as negative sampling message nodes. This balanced and targeted sampling strategy helps to obtain more valuable sample pairs, enabling the model to learn more effective information during training. In the detection phase, nodes that have richer connections to m_i are sampled to form message pairs. Considering that the multiple relations between messages may have varying impacts. These multiple relations are preserved comprehensively through a multi-relational sequence. We define an external vocabulary to represent multi-relational edges $\{\mathcal{R}_t\}_{t=1}^T$ in a multi-relational message graph. Each type of relation/edge R_t occupies two distinct discrete values, indicating whether this relation/edge exists between messages. Consequently, The different states (exist or not) of different relations between messages correspond to unique discrete values. As illustrated in the 'Pairwise Message Learning' part of Figure 3, given a social message pair (m_i, m_s) , we map the edge relations between them into the corresponding multi-relational sequence as follows:

$$h_{i,s} = \begin{cases} 1 & \text{if } \exists r_{i,s}^h \in \mathcal{R}_h \\ 0 & \text{otherwise} \end{cases}, u_{i,s} = \begin{cases} 3 & \text{if } \exists r_{i,s}^u \in \mathcal{R}_u \\ 2 & \text{otherwise} \end{cases}, e_{i,s} = \begin{cases} 5 & \text{if } \exists r_{i,s}^e \in \mathcal{R}_e \\ 4 & \text{otherwise} \end{cases}, d_{i,s} = \begin{cases} 7 & \text{if } \exists r_{i,s}^d \in \mathcal{R}_d \\ 6 & \text{otherwise} \end{cases}, \quad (1)$$

$$r_{i,s} = \{h_{i,s}, u_{i,s}, e_{i,s}, d_{i,s} \mid i, s \in |\mathcal{G}_k[\mathcal{M}]|\}$$

Here, $r_{i,s}^h$, $r_{i,s}^u$, $r_{i,s}^e$, and $r_{i,s}^d$ denote the hashtag, user, entity, and temporal edges between message node m_i and m_s , respectively. Correspondingly, \mathcal{R}_h , \mathcal{R}_u , \mathcal{R}_e , and \mathcal{R}_d represent the sets of hashtag, user, entity, and temporal edges in the multi-relational message graph \mathcal{G} . $r_{i,s}$ represents the multi-relational sequence between m_i and m_s . We represent one such sample pair and the batch of all nodes in the pairwise modeled multi-relational message

10 • P. Li et al.

graph as follows:

$$d_j = \{r_{i,s}, c_{i,s}, m_i, m_s \mid m_i, m_s \in \mathcal{G}_k[\mathcal{M}]\}, \quad (2)$$

$$D_k = \{d_j \mid 1 \leq j \leq |\mathcal{G}_k[\mathcal{M}]|\}, \quad (3)$$

where m_s is the sampled message node. The $c_{i,s}$ denotes the true cluster relation label for a message pair, which is assigned a value of 1 if the pair of messages belongs to the same event, and 0 otherwise. $r_{i,s}$ represents the multi-relational sequence between m_i and m_s . $\mathcal{G}_k[\mathcal{M}]$ and $|\mathcal{G}_k[\mathcal{M}]|$ denote the set of message nodes in the multi-relational message graph \mathcal{G} and its size, respectively. The d_j is a data pair of samples, and D_k is the batch of data after pairwise modeling for the k -th time block.

3.3 Multi-relational Prompt-based Pairwise Message Learning Mechanism

Owing to current GNN-based SED methods to learn message representations from explicit message relations, a challenge arises where GNNs cannot directly learn effective knowledge when messages of the same event share no common attributes yet have the same semantics. Conversely, when messages from two different events share common attributes, GNN-based methods may introduce noise from irrelevant messages, thereby affecting message representations. Furthermore, GNNs struggle to handle situations with missing structural relations. Consequently, we propose a new multi-relational prompt-based pairwise message learning mechanism. The core of the mechanism lies in harnessing the powerful semantic and contextual relationship understanding abilities of PLMs to learn the deeper connections and differences between messages from both textual content and structural information, thereby obtaining more detailed and comprehensive message representations. As illustrated in the ‘Pairwise Message Learning’ part of Figure 3, this mechanism comprises three key components: *Pairwise Message Encoding*, *Multi-Head Structured Attention*, and *Similarity-based Representation Aggregation*. The *Pairwise Message Encoding* module employs a PLM as the encoder to learn message representations from both message content embeddings and multi-relational prompt embeddings. Subsequently, we utilize the *Multi-Head Structured Attention* to extract more critical and clearer message representations. Finally, the final message representations are obtained through *Similarity-based Representation Aggregation*.

3.3.1 Pairwise Message Encoding. To incorporate the multi-relational sequence between messages as prompt information into the encoding process of PLMs, we transform it into multi-relational prompt embeddings. Considering that the multi-relational sequence serves as soft prompts, it reveals the relations between messages without directly conveying semantic information. Sharing the embedding space of the PLMs with message content may blur the semantic boundary between multi-relational prompts and textual content, thereby diminishing their guiding value. We design an independent and learnable prompt embedding layer to distinguish between message content embeddings and multi-relational prompt embeddings. This design prevents potential conflicts in the embedding space, enabling the model to adaptively learn optimized representations of different relational prompts during training, thereby enhancing the model’s capability to understand and process complex relationships in social messages.

In the encoding phase, we combine the embeddings of multi-relational prompts with pairwise message content embeddings, delving into the connections and distinctions in semantic and structural information between messages. Multiple messages are sampled for each message to form pairs, allowing each message to be repeatedly learned. This process enhances the learned message representations, making them more closely related to messages from the same event. Conversely, it makes message representations more distinctly differentiated from messages of different events. The specific encoding operation for the message pair m_i and m_s is as follows:

$$x_i, x_s = \text{Extract}(\Phi(\Psi(r_i^s); \Psi^{Plm}(m_i, m_s))), \quad (4)$$

, Vol. 1, No. 1, Article . Publication date: September 2024.

where $\Phi(\cdot)$ denotes the encoding operation of PLMs, $\Psi(\cdot)$ and $\Psi^{plm}(\cdot, \cdot)$ denote the multi-relational embedding operation and the message pair embedding operation. The r_i^s represents the multi-relational sequence of the message pair. The $\text{Extract}(\cdot)$ operation utilizes token type IDs to separately extract the encoded representation x_i of message m_i and the encoded representation x_s for sampled message m_s from the encoded representation of message pair. It is essential to note that, to ensure that x_i encompasses both message content and structural information, the multi-relational prompt encoded representation is treated as a part of x_i .

3.3.2 Multi-Head Structured Attention. To further extract key features from message representations, we introduce the multi-head structured attention mechanism [31, 75]. This mechanism can attend to different subspaces of the message representations. It highlights those message features closely related to the event while ignoring the irrelevant and noisy parts. This process is defined as:

$$\hat{x}_i = \text{Mean}(\Lambda(x_i, o)), \quad \hat{x}_s = \text{Mean}(\Lambda(x_s, o)). \quad (5)$$

Here, $\text{Mean}(\cdot)$ represents the operation of averaging representations, and $\Lambda(\cdot, \cdot)$ denotes the structured attention computation operation, with o being the number of attention heads. Finally, the resulting \hat{x}_i and \hat{x}_s are concatenated and fed into a classifier to discern the similarity between the two messages:

$$q_{i,s} = \sigma(\mathbf{W}^T \cdot (\hat{x}_i; \hat{x}_s) + b), \quad (6)$$

where \mathbf{W} is a weight matrix of dimension $(\cdot, 1)$, and b is a bias term. The linear layer maps the concatenated representation to a scalar, which is then mapped to the interval $(0, 1)$ through the sigmoid function $\sigma(\cdot)$, representing the similarity score q_i between two messages. Simultaneously, \hat{x}_i will also participate as a candidate representation of the message m_i in the message representation aggregation. We will introduce this in detail in Section 3.3.3.

3.3.3 Similarity-based Representation Aggregation. We propose a similarity-based representation aggregation approach to ensure that the final representation of m_i is as similar as possible to the message representations of the same event. The essence of this approach lies in identifying and filtering out ambiguous or noisy candidate message representations generated by negative sample pairs. Subsequently, aggregating the candidate message representations extracted from pairs of messages with high similarity makes the resulting final message representation more robust and stable.

In the detection phase, we sample n other message nodes for each message node m_i . Consequently, for each specific message m_i , we obtain a set of n candidate representations $X_i = [\hat{x}_{i,1}, \hat{x}_{i,2}, \dots, \hat{x}_{i,n}]$, along with the corresponding set of similarity $Q_i = [q_{i,1}, q_{i,2}, \dots, q_{i,n}]$. Subsequently, select the indices from Q_i that satisfy the conditions and average the candidate representations corresponding to these indices to obtain the final representation of m_i . We formalize this process as follows:

$$I_i = \{j \mid q_{i,j} > \alpha, j \in \{1, 2, \dots, n\}\}, \quad \bar{x}_i = \frac{1}{|I_i|} \sum_{j \in I_i} \hat{x}_{i,j}. \quad (7)$$

Here, I_i represents the set of indices obtained by filtering the set Q_i , α is the similarity retention threshold, and \bar{x}_i denotes the final representation of message m_i . Particularly, we choose the candidate representation corresponding to the highest similarity index in Q_i as the final representation for those messages in the similarity set Q_i with only one or no indices meeting the condition.

3.4 Optimization Objective

During the model training process, we enforce intra-event message proximity and inter-event message separation through *Pairwise Cross-Entropy Loss*. Additionally, to enhance intra-cluster cohesion and inter-cluster dispersion, we design an additional *Clustering Constraint*.

3.4.1 Pairwise Cross-Entropy Loss. Considering the pairwise training characteristic of the model, we construct the pairwise cross-entropy loss. This loss minimizes the divergence between the similarity of messages and the true cluster relationship labels, thereby drawing messages within the same cluster closer while pushing messages from different clusters further apart. Moreover, the pairwise cross-entropy loss also optimizes the model's classifier, enabling it to predict whether two messages belong to the same event more accurately. The pairwise cross-entropy loss is defined as:

$$L_{pair} = \frac{1}{B} \sum_{i=1}^B c_i \cdot \log(q_i) + (1 - c_i) \cdot \log(1 - q_i), \quad (8)$$

where B denotes the batch size, c_i represents the true cluster relation label of the message pair, and q_i represents the similarity of messages as predicted by the classifier.

3.4.2 Clustering Constraint. Given that the pairwise cross-entropy loss treats messages as objects to be pushed away or pulled closer, it may not effectively separate the various event clusters or ensure compactness within clusters. Hence, we further incorporate a new clustering constraint to enhance the distinguishability of message representations. This clustering constraint is composed of two parts: intra-cluster loss and inter-cluster loss. The intra-cluster loss aims to ensure that message representations from the same event are closer. Specifically, we initialize an updatable central feature matrix, mapping real event labels to the row indices of the matrix, where each row represents the central feature representation of a particular event. With the input of message representations, we treat the first incoming message representation as the initial central feature representation of that event category. In subsequent processes, the central feature representation is updated in a weighted manner. This process is formulated as follows:

$$h_j = \begin{cases} \hat{x}_i & \text{if } \text{Sum}(\hat{h}_j) = 0 \\ \beta \cdot \hat{h}_j + (1 - \beta) \cdot \hat{x}_i & \text{otherwise} \end{cases}, \quad (9)$$

$$\mathbf{H}_k = [h_1, h_2, \dots, h_j, \dots, h_{|E_k|}]^T, \quad (10)$$

where \hat{x}_i represents a message representation, and β is the retention weight for the central feature. \hat{h}_j and h_j are the central feature representations before and after updating, respectively. $\text{Sum}(\cdot)$ is the feature summation operation, \mathbf{H}_k is the center feature matrix, and $|E_k|$ is the number of event categories in the current message block. After updating the center feature matrix, we average the Euclidean distance between all message representations in the batch with their corresponding central feature representations to derive the final loss value. The intra-cluster loss is defined as:

$$L_{intra} = \frac{1}{B} \sum_{i=1}^B \mathbf{D}(x_i, h_j), \quad (11)$$

where $\mathbf{D}(\cdot, \cdot)$ calculates the euclidean distance. B denotes the batch size, and h_j represents message representation x_i corresponding central feature representation. Moreover, the inter-cluster loss aims to push apart the centroids of different event clusters. It is challenging to push away all messages from different clusters for each message, but we can achieve this by pushing apart the centroids of different clusters. We randomly shuffle the order of the central feature matrix and compute its Euclidean distance to the original matrix, utilizing a distance threshold to keep them apart. The inter-cluster loss is defined as:

$$H'_k = \text{Shuffled}(H_k), \quad L_{inter} = \frac{1}{|H_k|} \sum_{n=0}^{|H_k|} \text{Max}\{\gamma - \mathbf{D}(H_{k,n}, H'_{k,n}), 0\}. \quad (12)$$

Where $\text{Shuffled}(\cdot)$ denotes the shuffling operation performed on H_k once, i.e., random rearrangement of its order, $D(\cdot, \cdot)$ calculates the euclidean distance, and the $\text{Max}\{\cdot, \cdot\}$ operation selects the larger of two elements. γ represents the distance threshold. $H_{k,n}$ and $H'_{k,n}$ represent the n -th row of the original central feature matrix and the n -th row of the central feature matrix after shuffling the order. Through the constraints and guidance of inter-cluster and intra-cluster losses, representations of messages within the same cluster become more cohesive, while those of different clusters become more dispersed. The pairwise cross-entropy loss, inter-cluster loss, and intra-cluster loss complement each other. They jointly constrain and guide the model training, thereby achieving higher-quality message representations. The overall loss is defined as:

$$L = \kappa \cdot L_{pair} + \lambda \cdot L_{intra} + \mu \cdot L_{inter}, \quad (13)$$

where κ , λ , and μ represent the weights of the pairwise cross-entropy loss, intra-cluster loss, and inter-cluster loss, respectively.

3.5 Proposed RPLM_{SED}

As depicted in Algorithm 1, the life cycle of RPLM_{SED} is divided into three stages: the initial training phase, the detection phase, and the maintenance phase. During the initial training stage, we train the model using the initial message block (Algorithm 1, lines 17, 25-33). Subsequently, the model performs event detection on message blocks within the window (Algorithm 1, lines 3-13). To retain existing knowledge while adapting to new data, we perform a weighted update of the parameters in the PLMs within the model before the maintenance phase (Algorithm 1, line 21), then maintain the model by utilizing the message blocks within the window (Algorithm 1, lines 19, 20, 25-33). Once the model maintenance is completed, it will proceed to detect events in message blocks within the next window.

Lines 4-11 and 25-32 of Algorithm 1 separately introduce the pairwise message modeling strategy in the training and detection stages, including message node sampling and multiple relation mapping. Lines 2-14 of Algorithm 2 explain the multi-relational prompt-based pairwise message learning mechanism. After obtaining message representations (Algorithm 2, line 3) and their respective similarity (Algorithm 2, line 4), during the training phase, we construct the pairwise cross-entropy loss based on the similarity of messages and actual event relations between them. Additionally, the intra-cluster loss and inter-cluster loss (Algorithm 2, line 7) are constructed by the central feature matrix (Algorithm 2, line 6) and message representations. The model is optimized through backpropagation of the total loss (Algorithm 2, line 8, 9). During the detection phase, we update the candidate representation set and the similarity set for each message based on the obtained candidate representations of messages and their corresponding similarities. (Algorithm 2, lines 11, 12). Subsequently, we obtain the final representation of each message through the similarity-based representation aggregation approach (Algorithm 2, lines 17-20). Finally, events in a message block are detected, utilizing all messages' final representations through clustering algorithms (Algorithm 2, line 21). Herein, we employ two distinct clustering algorithms: the distance-based *K-Means* [48] and the density-based *HDBSCAN* [10] algorithms. *HDBSCAN* is an enhanced version of the *DBSCAN* [19] algorithm, incorporating concepts from hierarchical clustering. This method does not require pre-specifying the number of clusters. Instead, it identifies the optimal clustering outcome by constructing a density hierarchy. Compared to the original *DBSCAN*, *HDBSCAN* offers greater flexibility in parameter settings, effectively adapts to data from different density regions, and exhibits stronger robustness. In practice, density-based clustering algorithms are more suitable as they do not require pre-setting the number of clusters. This allows them to better adapt to incremental data changes and meet the needs for continuous or dynamic detection.

Algorithm 1: RPLM_{SED}: Relational Prompt-based Pre-trained Language Models for Social Event Detection.

Input: A social stream $S = M_0, M_1, M_2, \dots$, the ground truth label set: $L = L_0, L_1, L_2, \dots$, the window size: w , and the number of mini-batches B .

Output: Sets of events: E_0, E_1, E_2, \dots .

```

1 for  $t = 0, 1, 2, \dots$  do
2   /* Detect events from  $M_t$  */
3   if  $t > 0$  then
4      $\mathcal{G} \leftarrow$  Construct multi-relational message graph;
5     for  $a = 1, 2, 3, \dots, n$  do
6        $(m_i, m_s) \leftarrow \{m_i, \text{Sample}(\mathcal{G}[M_t]) \mid \forall m_i \in M_t\}$ ;
7       // Sampling messages that have richer connections with  $m_i$  to form pairs.
8        $r_{i,s} \leftarrow$  Eq. 1;
9        $d \leftarrow$  Eq. 2;
10      Append  $d$  to  $D$  (Eq. 3);
11    end
12     $D, \text{Train} = \text{False} \Rightarrow$  Execute Algorithm 2;
13  end
14  /* Initial model training or the model maintenance. */
15  if  $t \% w == 0$  then
16    if  $t == 0$  then
17       $M, L^{gt} = M_0, L_0$ ; //  $M_0$  is used to train initial model  $f_0$ .
18    else
19       $M = M_{t-w+1} + \dots + M_t$ ; // Using messages in the window to maintain the model.
20       $L^{gt} = L_{t-w+1} + \dots + L_t$ ;
21       $\theta_t^{PLM} = \zeta \cdot \theta_{t-w}^{PLM} + (1 - \zeta) \cdot \theta_0^{PLM}$ ; // Parameter weighted update of PLM in RPLMSED.
22      ( $\theta_{t-w}^{PLM}$  and  $\theta_0^{PLM}$  represent the PLM parameters in RPLMSED and the initial parameters of the
23      PLM, respectively.  $\zeta$  denotes the retention weight threshold.)
24    end
25    /* Pairwise message modeling. */
26     $\mathcal{G} \leftarrow$  Construct multi-relational message graph;
27    for  $a = 1, 2, 3, \dots, y$  do
28       $(m_i, m_s^+), (m_i, m_s^-) \leftarrow \{m_i, \text{Sample}(\mathcal{G}[M], L^{gt}) \mid \forall m_i \in M\}$ ;
29      // Utilizing ground truth labels to choose positive and negative samples.
30       $r_{i,s} \leftarrow$  Eq. 1;
31       $d^+, d^- \leftarrow$  Eq. 2;
32      Append  $d^+, d^-$  to  $D$  (Eq. 3);
33    end
34     $D, \text{Train} = \text{True} \Rightarrow$  Execute Algorithm 2;
35  end

```

Algorithm 2: Model Training and Detection.

```

1 /* Multi-relational prompt-based pairwise message learning. */
2 for  $b = 1, 2, 3, \dots, B$  do
3   get representations of the message and sampled message from each message pair via Eq. 4, Eq. 5;
4   Get the similarity of the two in each message pair via Eq. 6;
5   if Train then
6     Update central feature matrix  $H \leftarrow$  Eq. 9;
7      $L_{pair}, L_{intra}, L_{inter} \leftarrow$  Eq. 8, Eq. 11, Eq. 12;
8      $L \leftarrow$  Eq. 13;
9     Back propagate to update model parameters;
10  else
11    Update sets of each message candidate representations:  $X_1, X_2, \dots, X_i$ ;
12    Update corresponding similarity sets:  $Q_1, Q_2, \dots, Q_i$ ;
13  end
14 end
15 /* Similarity-based representation aggregation and Clustering. */
16 if not Train then
17   for  $X, Q$  in  $[X_1, X_2, X_3, \dots, X_i], [Q_1, Q_2, Q_3, \dots, Q_i]$  do
18      $\bar{x}_i \leftarrow$  Eq. 7;
19     Append final message representation  $\bar{x}_i$  into the clustering feature matrix Matrix;
20   end
21    $E_t = \{e_0, e_1, e_2, \dots\} \leftarrow$  Cluster(Matrix);
22 end

```

3.6 Time Complexity

The time complexity of RPLM_{SED} consists of pairwise message modeling, multi-relational prompt-based pairwise message learning mechanism, and message representation clustering. During pairwise message modeling, constructing a HIN graph and transforming it into a multi-relational message graph necessitates traversing all edges, resulting in a time complexity of $O(N_e)$, where N_e is the number of edges in the graph. The time complexity for sampling messages to form pairs is $O(N \cdot n) = O(N)$, where N is the number of messages in the graph, and n is the number of message pairs sampled for each message. Thus, the total time complexity of this phase is $O(N_e + N) = O(N_e)$. In the multi-relational prompt-based pairwise message learning mechanism, firstly, embedding operations for message pair content and multi-relation sequences require time $O(N \cdot n \cdot s) = O(N)$, where s is the average input sequence length. Secondly, the PLM comprises multiple layers of *Transformers*, where each layer's primary source of time complexity stems from the self-attention mechanism. The time complexity in the encoding phase is $O(N \cdot n \cdot s^2) = O(N \cdot s^2)$. Notably, as the scale of PLM increases, higher layer counts and hidden dimensions lead to increased time complexity. Thirdly, the time complexities of the structured attention mechanism and classifier are $O(N \cdot n \cdot s) = O(N)$ and $O(N)$, respectively. Finally, the time complexity for selecting and aggregating final message representations is $O(N \cdot n) = O(N)$. Therefore, the overall time complexity of the pairwise message learning mechanism is $O(N + N \cdot s^2 + N + N + N) = O(N \cdot s^2)$. In the clustering stage, the time complexity for clustering message representations using the K-means algorithm is $O(I \cdot K \cdot N \cdot d) = O(N)$, where I and K denote the number of iterations and the given number of clusters, respectively. In contrast, the time complexity of clustering with the HDBSCAN algorithm is $O(N \cdot \log N)$.

4 EXPERIMENTAL SETUP

This section provides a detailed overview of the experimental setup, including dataset selection, comparison baselines, evaluation metrics, and experiment parameter configurations. Specifically, in Section 4.1, we describe the details of the datasets employed in this study. In Section 4.2, we list and discuss all the baseline methods included in the comparison, aiming to conduct a comprehensive comparative analysis. Furthermore, in Section 4.3, we present the hardware and software configurations for our experiments. We also offer a comprehensive description of the parameter settings for $RPLM_{SED}$ and the implementation approaches for baseline methods, ensuring the reproducibility of our experimental results. Finally, Section 4.4 elaborates on the evaluation metrics adopted for the experiments.

4.1 Datasets

As shown in Table 2, we evaluate the performance of the $RPLM_{SED}$ model on three publicly available social event datasets: Events2012, Events2018, and Arabic-Twitter. We test the model's performance in offline and incremental scenarios using the Events2012 and Events2018 datasets and explore the model's performance under low-resource language conditions using the Arabic Twitter dataset. Moreover, to observe the model's performance on long-tail distribution issues, we extracted 100 event categories from both the Events2012 and Events2018 datasets. Considering that for the long-tail recognition task, an imbalanced training set alongside balanced validation and test sets should be provided, we select an additional 20 and 30 tweets for each event for the validation and test sets, respectively. After filtering out unavailable and duplicate tweets, the detailed information of these datasets is as described:

- **Events2012** [51]. The Events2012 dataset contains 68,841 annotated English tweets covering 503 different event categories, encompassing tweets over a consecutive 29-day period. The **Events2012_100** contains 100 events with a total of 15,019 tweets, where the maximal event comprises 2,377 tweets, and the minimally has 55 tweets, with an imbalance ratio of approximately 43.
- **Events2018** [50]. The Events2018 includes 64,516 annotated French tweets covering 257 different event categories, with data spanning over a consecutive 23-day period. The **Events2018_100** contains 100 events with a total of 19,944 tweets, where the maximal event comprises 4,189 tweets and the minimally has 27 tweets, an imbalance ratio of approximately 155.
- **Arabic-Twitter** [2]. The Arabic-Twitter dataset comprises 9,070 annotated Arabic tweets, covering seven catastrophic-class events from various periods.

Table 2. Statistics of the datasets.

Dataset	Messages	Events	Timespan	Max. tweet count/Event	Min. tweet count/Event	Avg. tweet count/Event
Events2012	68841	503	29 days	7133	1	137
Events2018	64516	257	23 days	12567	1	251
Events2012_100	15019	100	-	2377	55	150
Events2018_100	19944	100	-	4189	27	199
Arabic-Twitter	9070	7	-	2883	519	1296

4.2 Baselines

To verify the effectiveness of our proposed $RPLM_{SED}$, we compare it with popular social event detection methods in recent years. These baseline methods include:

- **PP-GCN** [57]: PP-GCN is an offline fine-grained social event detection method based on Graph Convolutional Networks (GCN) [32].
- **EventX** [39]: EventX is a fine-grained event detection method based on community detection suitable for online scenarios.
- **KPGNN** [11]: KPGNN is a knowledge-preserving incremental social event detection framework based on heterogeneous GNN, leveraging the inductive learning capability of GNN to represent and detect events effectively.
- **QSGNN** [65]: QSGNN focuses on open-set social event detection tasks. It extends knowledge from known to unknown by utilizing the best-known samples (i.e., setting stricter constraints in inter-class distances and directional relationships) and reliable knowledge transfer (generating and selecting high-quality pseudo labels).
- **FinEvent** [60]: FinEvent is an event detection framework that combines reinforcement learning with GNNs, achieving incremental and cross-lingual social event detection. Concurrently, this framework proposes the Deep Reinforcement Learning Guided DBSCAN Model (DRL-DBSCAN) [86]. In this paper, the FinEvent baseline, based on the K-means clustering method, is denoted as **FinEvent^k**, while the FinEvent baseline leveraging DGL-DBSCAN is referred to as **FinEvent^d**.
- **CLKD** [67]: CLKD is a cross-lingual social event detection framework based on knowledge distillation, jointing knowledge distillation, and GNNs to improve social event detection performance in low-resource language.
- **ETGNN** [66]: ETGNN is an evidential temporal-aware graph neural network for social event detection, which integrates view-specific graph quality and temporal information uncertainty into social event detection.
- **UCL_{SED}** [68]: UCL_{SED} is an uncertainty-guided class imbalance learning framework for imbalanced social event detection tasks. It focuses on boundary learning in latent space and classifier learning with high-quality uncertainty estimation.
- **MVGAN** [16]: MVGAN is a multi-view graph attention network for social event detection. It enriches event semantics through node aggregation and multi-view fusion in heterogeneous message graphs.

Additionally, we enumerate the methods employed in previous works that utilize word embedding models and pre-trained models for event detection while also proposing new approaches for implementing incremental social event detection through PLMs:

- **Word2Vec** [52]: Word2Vec is a pre-trained word embedding technique. It employs the average of the embeddings of all words within a message to serve as the message embedding.
- **BERT** [30]: BERT is a pre-trained model for natural language processing (NLP) proposed by researchers at Google in 2018. The cornerstone of this model's innovation lies in its utilization of the encoder architecture from the Transformer model to train context-based word embeddings, incorporating a Masked Language Model (MLM) design. Consistent with previous methodologies, we employ BERT to obtain message representations for clustering directly. **BERT-FT** achieves streaming event detection by applying fine-tuning to BERT. Given the inapplicability of cross-entropy in open event sets, the encoding process is optimized through online triplet loss and the implementation of proposed clustering constraints. Analogous to standard incremental event detection methods, BERT-FT, after its training or maintenance, detects events within message blocks inside the window.
- **RoBERTa** [43]: RoBERTa is an enhanced version of the BERT model, introduced by the Facebook AI research team in 2019. It is trained using larger datasets, increased batch sizes, and more granular parameter tuning. Compared to BERT, this model adopts a dynamic masking strategy and implements more refined adjustments to hyperparameters. Similarly, we utilize RoBERTa to encode messages for clustering directly.

RoBERTa-FT shares the same training approach as BERT-FT, with the distinction between them lying solely in using different PLMs.

- **BERT-LR**: BERT-LR is a variant of $RPLM_{SED}$ that fine-tunes BERT within $RPLM_{SED}$ using Low-Rank Adaptation (LoRA) [28], an efficient fine-tuning technique for PLMs. The core concept of this approach is to introduce low-rank weights into the parameters of the pre-trained model to balance parameter efficiency and performance during model fine-tuning. **RoBERTa-LR** is another variant of $RPLM_{SED}$, which, in contrast to BERT-LR, is based on RoBERTa.

4.3 Experimental Setting and Implementation

In the English dataset Events2012, we select the RoBERTa (RoBERTa-large) [43] released by Facebook AI in 2019 as the backbone architecture for the $RPLM_{SED}$ model. In the French dataset Events2018 and the Arabic dataset Arabic-Twitter, we use the multilingual version of the TwHIN-BERT (TwHIN-BERT-large) [87]. TwHIN-BERT is a multilingual tweet language model that is trained on 7 billion Tweets from over 100 distinct languages. It is trained with not only text-based self-supervision (e.g., MLM) but also with a social objective based on the rich social engagements within a Twitter Heterogeneous Information Network (TwHIN). During the data preprocessing phase, for each message in the training set, we sample 60 positive and negative samples. For each message in the validation and test sets, 180 samples are sampled. The experimental setup includes setting the number of training epochs to 2, the window size to 3, the learning rate to $2e - 5$, and the batch size to 35. During the model maintenance phase, the parameter update weight ζ is set to 0.4, the inter-cluster loss distance threshold γ is defined as 10, and the loss weights κ , λ , and μ are set to 1, 0.01, and 0.005, respectively. The similarity threshold α for the message feature aggregation stage is set to 0.9, and the number of heads in structured attention mechanism o is set to 2. The embedding and encoding dimensions of the model are consistent with the chosen PLM. In terms of optimizer selection, the entire model is trained using Adam. All experiments are repeated 5 times to ensure the stability of the results, and the mean and standard deviation of the experiments are reported.

We implement all models using Python 3.8 and PyTorch 2.0. All experiments are conducted on a 32-core Intel Core i9-13900K@3.80GHz, with 128GB RAM and an NVIDIA GeForce RTX 4090 GPU. For PP-GCN², KPGNN³, QSGNN⁴ and FinEvent⁵, we follow their open-source implementations. We implement EventX with Python 3.7. To ensure fairness in comparison, unless specifically noted, we uniformly apply the K-Means [48] clustering algorithm to all baselines and $RPLM_{SED}$, excluding EventX, setting the total number of categories equal to the number of event categories in the real world. Additionally, we train the FinEvent framework from scratch on the French dataset, rather than utilizing its proposed cross-lingual social message representation learning method, ‘Crlme’. Furthermore, for the DGL-DBSCAN method proposed by FinEvent, we conduct a more comprehensive comparison by further clustering the message representations generated by $RPLM_{SED}$ through the HDBSCAN [10] clustering method, and conduct a comprehensive comparison and analysis of the results. Where the minimum cluster size parameter of HDBSCAN is set to 8.

4.4 Evaluation Metrics

To evaluate the effectiveness of $RPLM_{SED}$, we adopt three widely-used metrics to assess the quality of clustering: Normalized Mutual Information (NMI) [20], Adjusted Mutual Information (AMI) [76], and the Adjusted Rand Index (ARI) [76]. The NMI metric is employed to gauge the clustering quality, primarily focusing on the congruence

²<https://github.com/RingBDStack/PPGCN>

³<https://github.com/RingBDStack/KPGNN>

⁴<https://github.com/RingBDStack/open-set-social-event-detection>

⁵<https://github.com/RingBDStack/FinEvent>

between our model's clustering output and the actual ground truth categories:

$$\text{NMI} = \frac{2 \cdot \text{MI}(P, T)}{H(P) + H(T)}, \quad (14)$$

where P and T represent the predicted label set and the true label set, respectively. $\text{MI}(P, T)$ denotes the mutual information between them, and $H(P)$ and $H(T)$ respectively denote their entropies. AMI represents an enhancement of the NMI metric. This improvement is achieved by adjusting for the expected mutual information derived from random label assignments. Therefore, AMI provides values that are closer to 0 in scenarios where labeling is performed randomly, thereby offering a more accurate assessment of the information shared between clusters:

$$\text{AMI} = \frac{\text{MI}(P, T) - \mathbb{E}[\text{MI}(P, T)]}{\max\{H(P), H(T)\} - \mathbb{E}[\text{MI}(P, T)]}, \quad (15)$$

where $\mathbb{E}[\text{MI}(P, T)]$ represents the expected value of the mutual information between P and T . The Rand Index measures the similarity between two data clustering outcomes. It is particularly used to compare the similarity between algorithmic clustering results and benchmark clusterings. The Rand Index values range from 0 to 1, where 1 indicates identical clustering. ARI improves upon the original Rand Index by ensuring the expected value of random label assignments is 0 and adjusts the range to be between -1 and 1:

$$\text{ARI} = \frac{\text{RI} - \text{ERI}}{\text{MaxRI} - \text{ERI}}, \quad (16)$$

where RI is the original Rand Index, ERI represents the Expected Rand Index, and MaxRI denotes the maximum value in the Rand Index.

5 RESULTS AND DISCUSSION

This section primarily revolves around several key questions to empirically evaluate the performance of the RPLM_{SED} model:

- Q1: How does the performance of the proposed PLM-based method RPLM_{SED} compare to current mainstream SED methods in both offline and online detection scenarios (Sect 5.1)?
- Q2: How do different models fare in the low-resource scenario (Section 5.2)?
- Q3: How do different models address the long-tail distribution issue of social data (Section 5.3)?
- Q4: How do various PLM-based approaches perform in online scenarios (Section 5.4)?
- Q5: How does the performance of the RPLM_{SED} model, based on different PLMs, manifest (Section 5.5)?
- Q6: How does the time consumption of RPLM_{SED} based on different scale PLMs (Section 5.6)?
- Q7: How do different components and hyperparameters in RPLM_{SED} influence overall effectiveness (Section 5.7, Section 5.8)?
- Q8: How does the performance of RPLM_{SED} compare to the strongest baseline in a specific scenario (Section 5.9)?
- Q9: How does the performance of RPLM_{SED} based on different cluster algorithms (Section 5.10)?

5.1 Model Effectiveness

In this section, we compare and analyze RPLM_{SED} and all baseline methods under offline and online detection scenarios. To ensure consistency in the evaluation, RPLM_{SED}^k is compared with baseline methods that employ the K-Means clustering approach. In contrast, RPLM_{SED}^d is separately compared with FinEvent^d.

5.1.1 Offline Evaluation. In the offline detection scenario, the training, validation, and test sets share the same events. For the Events2012 and Events2018 datasets, experiments are conducted using data from the first week, with 70% for training, 20% for testing, and 10% for validation. As illustrated in Table 3, the RPLM_{SED} outperforms

Table 3. Evaluation on the Closed Set. The best results are bolded, and the second-best results are underlined.

Blocks	Metrics	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	RPLM ^k _{SED}	RPLM ^d _{SED}	Improve
M_0 in Events 2012	NMI	.70 ± .02	.68 ± .00	.76 ± .02	.79 ± .01	.84 ± .01	<u>.86 ± .01</u>	<u>.86 ± .01</u>	.87 ± .01	↑ .01
	AMI	.56 ± .01	.29 ± .00	.64 ± .02	.68 ± .01	.72 ± .01	.77 ± .01	<u>.81 ± .00</u>	.84 ± .00	↑ .07
	ARI	.20 ± .01	.05 ± .00	.34 ± .01	.38 ± .01	.40 ± .02	<u>.56 ± .02</u>	.41 ± .01	.80 ± .02	↑ .24
M_0 in Events 2018	NMI	.60 ± .01	.57 ± .00	.66 ± .03	.71 ± .02	.73 ± .02	.78 ± .01	.93 ± .00	<u>.92 ± .00</u>	↑ .15
	AMI	.49 ± .02	.56 ± .00	.60 ± .02	.64 ± .02	.69 ± .01	.74 ± .01	.92 ± .01	<u>.91 ± .01</u>	↑ .18
	ARI	.17 ± .01	.01 ± .00	.27 ± .03	.33 ± .02	.35 ± .02	.61 ± .02	<u>.94 ± .01</u>	.95 ± .01	↑ .34

Table 4. Incremental evaluation NMI on Events2012. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	RPLM ^k _{SED}	RPLM ^d _{SED}	Improve
M_1	.23 ± .00	.36 ± .00	.39 ± .00	.43 ± .01	.38 ± .01	<u>.84 ± .01</u>	.46 ± .01	.91 ± .02	↑ .07
M_2	.57 ± .00	.68 ± .00	.79 ± .01	.81 ± .02	.81 ± .00	.84 ± .01	<u>.89 ± .02</u>	.91 ± .01	↑ .07
M_3	.55 ± .01	.63 ± .00	.67 ± .00	.78 ± .01	.71 ± .01	<u>.89 ± .00</u>	.88 ± .00	.93 ± .00	↑ .04
M_4	.46 ± .01	.63 ± .00	.67 ± .02	.71 ± .02	.71 ± .01	.71 ± .01	<u>.78 ± .01</u>	.83 ± .01	↑ .12
M_5	.48 ± .00	.59 ± .00	.73 ± .00	.75 ± .00	.76 ± .00	.83 ± .00	.85 ± .01	<u>.85 ± .02</u>	↑ .02
M_6	.57 ± .01	.70 ± .00	.82 ± .01	.83 ± .01	.84 ± .00	.83 ± .00	<u>.87 ± .01</u>	.92 ± .00	↑ .08
M_7	.37 ± .00	.51 ± .00	.55 ± .01	.57 ± .01	.56 ± .00	<u>.73 ± .01</u>	.68 ± .01	.88 ± .01	↑ .15
M_8	.55 ± .02	.71 ± .00	.80 ± .00	.79 ± .01	.87 ± .01	.87 ± .02	<u>.88 ± .01</u>	.88 ± .00	↑ .01
M_9	.51 ± .00	.67 ± .00	.74 ± .02	.77 ± .02	.78 ± .02	.79 ± .01	<u>.87 ± .01</u>	.92 ± .01	↑ .13
M_{10}	.55 ± .02	.68 ± .00	.80 ± .01	.82 ± .02	.81 ± .01	.82 ± .01	<u>.90 ± .00</u>	.91 ± .01	↑ .09
M_{11}	.50 ± .01	.65 ± .00	.74 ± .01	.75 ± .01	.76 ± .00	.75 ± .00	<u>.87 ± .01</u>	.88 ± .01	↑ .12
M_{12}	.45 ± .01	.61 ± .00	.68 ± .01	.70 ± .01	.76 ± .01	.67 ± .01	<u>.85 ± .01</u>	.92 ± .00	↑ .16
M_{13}	.47 ± .01	.58 ± .00	.69 ± .01	.68 ± .02	.67 ± .00	<u>.79 ± .00</u>	.77 ± .01	.91 ± .01	↑ .12
M_{14}	.44 ± .01	.57 ± .00	.69 ± .00	.68 ± .01	.74 ± .00	<u>.82 ± .00</u>	.82 ± .01	.88 ± .00	↑ .06
M_{15}	.39 ± .01	.49 ± .00	.58 ± .00	.59 ± .01	.64 ± .00	.69 ± .01	<u>.75 ± .01</u>	.83 ± .01	↑ .14
M_{16}	.55 ± .01	.62 ± .00	.79 ± .01	.78 ± .01	.80 ± .00	<u>.90 ± .01</u>	.84 ± .00	.93 ± .01	↑ .03
M_{17}	.48 ± .00	.58 ± .00	.70 ± .01	.71 ± .01	.73 ± .00	<u>.83 ± .00</u>	.80 ± .01	.86 ± .01	↑ .03
M_{18}	.47 ± .01	.59 ± .00	.68 ± .02	.70 ± .01	.72 ± .01	.74 ± .01	<u>.82 ± .01</u>	.83 ± .00	↑ .09
M_{19}	.51 ± .02	.60 ± .00	.73 ± .01	.73 ± .00	.76 ± .02	.66 ± .01	<u>.88 ± .01</u>	.91 ± .00	↑ .15
M_{20}	.51 ± .01	.67 ± .00	.72 ± .02	.73 ± .02	.73 ± .00	.80 ± .00	.85 ± .00	<u>.82 ± .01</u>	↑ .05
M_{21}	.41 ± .02	.53 ± .00	.60 ± .00	.61 ± .01	.65 ± .01	<u>.74 ± .01</u>	.76 ± .00	.74 ± .02	↑ .02

all baseline methods across all metrics (NMI, AMI, ARI). Compared to EventX, on the Events2012 dataset, RPLM^k_{SED} shows improvements of 26%, 180%, and 710% in NMI, AMI, and ARI, respectively. On the Events2018 dataset, RPLM^k_{SED} exhibits enhancements of 64%, 64%, and 9350% in NMI, AMI, and ARI, respectively. RPLM_{SED} significantly outperforms EventX on both datasets, especially in terms of ARI, indicating that EventX tends to produce a higher number of clusters than the actual events. In contrast, RPLM_{SED} demonstrates greater stability. Compared with PP-GCN, KPGNN, QSGNN, and FinEvent^k, RPLM^k_{SED} shows significant improvements on the Events2012 dataset, with increases of 2%-23% in NMI, 13%-45% in AMI, and 1%-103% in ARI. Furthermore, on the Events2018 dataset, improvements include 27% to 55% in NMI, 33% to 87% in AMI, and 170% to 456% in ARI. Compared to FinEvent^d on the Events2012 dataset, RPLM^d_{SED} exhibits improvements of 1%, 9%, and 43% in NMI, AMI, and ARI, respectively. On the Events2018 dataset, the increases are 18%, 23%, and 56% in NMI, AMI, and

ARI, respectively. The superior performance of $RPLM_{SED}$ compared to GNN-based methods can be attributed to its approach of not relying on explicit structural relations for representation learning, which also enhances its robustness. Furthermore, by concurrently leveraging the structure and semantics of messages, $RPLM_{SED}$ is capable of capturing the interplay between structural and semantic information of messages, thereby facilitating a more accurate determination of event relations between messages. Conversely, GNN-based models are impacted during the learning process by the presence of noisy edges between different event messages or the absence of effective connections among messages of the same event, thereby affecting the quality of message representation. Additionally, the static initialization of message nodes impedes their ability to dynamically refine the embeddings of message content throughout the training process in the manner exhibited by $RPLM_{SED}$.

Table 5. Incremental evaluation AMIs on Events2012. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	$RPLM_{SED}^k$	$RPLM_{SED}^d$	Improve
M_1	.21 ± .00	.06 ± .00	.37 ± .00	.41 ± .02	.36 ± .01	<u>.84 ± .01</u>	.44 ± .01	.91 ± .01	↑ .07
M_2	.55 ± .02	.29 ± .02	.78 ± .01	.80 ± .01	.77 ± .00	.84 ± .01	<u>.88 ± .01</u>	.91 ± .00	↑ .07
M_3	.52 ± .01	.18 ± .01	.74 ± .00	.76 ± .01	.82 ± .01	<u>.89 ± .01</u>	.87 ± .00	.93 ± .00	↑ .04
M_4	.42 ± .01	.19 ± .01	.64 ± .01	.68 ± .01	.67 ± .02	.69 ± .00	<u>.76 ± .01</u>	.81 ± .01	↑ .12
M_5	.46 ± .01	.14 ± .00	.71 ± .01	.73 ± .00	.74 ± .00	.82 ± .00	<u>.83 ± .01</u>	.84 ± .02	↑ .02
M_6	.52 ± .02	.27 ± .00	.79 ± .01	.80 ± .01	.81 ± .00	.82 ± .02	<u>.85 ± .01</u>	.91 ± .00	↑ .09
M_7	.34 ± .00	.13 ± .00	.51 ± .01	.54 ± .00	.53 ± .00	<u>.72 ± .00</u>	.66 ± .01	.88 ± .01	↑ .16
M_8	.49 ± .02	.21 ± .00	.76 ± .00	.75 ± .01	.84 ± .01	.87 ± .01	.86 ± .01	<u>.86 ± .00</u>	↓ .01
M_9	.46 ± .02	.19 ± .00	.71 ± .02	.75 ± .02	.75 ± .00	.78 ± .01	<u>.86 ± .01</u>	.91 ± .00	↑ .13
M_{10}	.51 ± .02	.24 ± .00	.78 ± .01	.80 ± .03	.78 ± .00	.81 ± .00	<u>.88 ± .01</u>	.90 ± .01	↑ .09
M_{11}	.46 ± .01	.24 ± .00	.71 ± .01	.72 ± .01	.73 ± .00	.74 ± .00	<u>.85 ± .01</u>	.87 ± .01	↑ .13
M_{12}	.42 ± .01	.16 ± .00	.66 ± .01	.68 ± .00	.75 ± .01	.67 ± .02	.84 ± .00	<u>.77 ± .00</u>	↑ .09
M_{13}	.43 ± .01	.16 ± .00	.67 ± .01	.66 ± .01	.64 ± .00	<u>.79 ± .00</u>	.75 ± .01	.91 ± .00	↑ .12
M_{14}	.41 ± .01	.14 ± .00	.65 ± .00	.66 ± .01	.72 ± .00	<u>.82 ± .01</u>	.81 ± .01	.88 ± .01	↑ .06
M_{15}	.35 ± .01	.07 ± .00	.54 ± .00	.55 ± .01	.61 ± .00	.67 ± .01	<u>.73 ± .01</u>	.82 ± .01	↑ .15
M_{16}	.52 ± .01	.19 ± .00	.77 ± .01	.76 ± .02	.75 ± .01	<u>.90 ± .00</u>	.82 ± .00	.93 ± .01	↑ .03
M_{17}	.45 ± .00	.18 ± .00	.68 ± .01	.69 ± .01	.71 ± .02	<u>.82 ± .00</u>	.79 ± .00	.86 ± .01	↑ .04
M_{18}	.45 ± .01	.16 ± .00	.66 ± .02	.68 ± .01	.70 ± .00	.74 ± .00	<u>.81 ± .01</u>	.82 ± .01	↑ .08
M_{19}	.48 ± .02	.16 ± .00	.71 ± .01	.70 ± .01	.75 ± .01	.66 ± .00	<u>.86 ± .01</u>	.90 ± .01	↑ .15
M_{20}	.45 ± .02	.18 ± .00	.68 ± .02	.69 ± .02	.68 ± .00	.78 ± .00	.82 ± .00	<u>.80 ± .01</u>	↑ .04
M_{21}	.38 ± .02	.10 ± .00	.57 ± .00	.58 ± .00	.63 ± .01	.64 ± .01	.74 ± .01	<u>.71 ± .01</u>	↑ .10

5.1.2 Online Evaluation. In the online detection scenario, we utilize the data from the first week as the M_0 to train the initial model, constructing subsequent message blocks daily. It is crucial to note that for offline baselines like PP-GCN, direct adaptation to the online scenario is not feasible. We circumvent this by retraining the model using previous message blocks as the training set and predicting the current message block. Unlike offline detection, the $RPLM_{SED}$ method enters the maintenance phase after predicting the message block within the window. During the maintenance phase, $RPLM_{SED}$ initially updates the parameters of the PLMs section through a parameter weighting method, thereby enabling rapid adaptation to new knowledge while preserving historical knowledge. Subsequently, the model is maintained using message blocks within the current window.

As shown in Tables 4 to 9, $RPLM_{SED}$ significantly outperforms all baseline methods in incremental event detection, consistently leading in evaluation metrics across all message blocks. Overall, in comparison with EventX, on the Events2012 dataset, $RPLM_{SED}^k$ achieves improvements of 35%, 360%, and 480% in terms of NMI,

AMI, and ARI, respectively. On the Events2018 dataset, $RPLM_{SED}^k$ shows enhancements of 64%, 346%, and 4105% in NMI, AMI, and ARI, respectively. This can be attributed to the fact that EventX relies on community detection whilst neglecting the rich semantic information present in social messages. Compared to PP-GCN, KPGNN, QSGNN, and FinEvent^k, $RPLM_{SED}^k$ exhibits improvements ranging from 12% to 71%, 12% to 81%, and 21% to 56% in NMI, AMI, and ARI, respectively, on the Events2012 dataset. On the Events2018 dataset, $RPLM_{SED}^k$ demonstrates enhancements of 22% to 34%, 22% to 36%, and 55% to 78% in NMI, AMI, and ARI, respectively. Compared to FinEvent^d, $RPLM_{SED}^d$ achieves improvements of 11%, 11%, and 28% in terms of NMI, AMI, and ARI, respectively, on the Events2012 dataset. On the Events2018 dataset, $RPLM_{SED}^d$ shows enhancements of 11%, 12%, and 32% in NMI, AMI, and ARI, respectively. Regardless of the clustering method employed, our model maintained optimal performance. These improvements further underscore the effectiveness of $RPLM_{SED}$. Within the model, the multi-relational prompt-based pairwise message learning mechanism acquires more effective message representations by learning the connections and differences between messages in pairs. Subsequently, through the similarity-based representation aggregation method, message representations are filtered and aggregated, further enhancing the robustness of message representations. Additionally, converting the structural relations between messages into multi-relational prompts allows the model to better cope with scenarios characterized by a scarcity of structural information. Furthermore, at each maintenance stage, the parameters of the PLM within $RPLM_{SED}$ are updated through a weighted method. This endows $RPLM_{SED}$ with enhanced capabilities for knowledge adaptation, preservation, and expansion, thereby exhibiting stronger generalization abilities and stability when dealing with social data streams.

As depicted in Figure 4, to intuitively compare and further demonstrate the effectiveness of our proposed $RPLM_{SED}$ model, we select message representations of three message blocks (M_2 , M_9 , and M_{19}) from the Events2012

Table 6. Incremental evaluation ARIs on Events2012. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	$RPLM_{SED}^k$	$RPLM_{SED}^d$	Improve
M_1	.05 ± .00	.01 ± .00	.07 ± .01	.09 ± .01	.05 ± .00	<u>.90 ± .00</u>	.07 ± .01	.96 ± .01	↑ .06
M_2	.67 ± .03	.45 ± .02	.76 ± .02	.78 ± .02	.67 ± .01	<u>.90 ± .01</u>	.83 ± .02	.91 ± .01	↑ .00
M_3	.47 ± .01	.09 ± .01	.58 ± .01	.62 ± .01	.58 ± .00	<u>.89 ± .01</u>	.66 ± .01	.94 ± .00	↑ .05
M_4	.24 ± .01	.07 ± .01	.29 ± .01	.31 ± .02	.27 ± .02	.27 ± .01	<u>.39 ± .02</u>	.61 ± .00	↑ .32
M_5	.34 ± .00	.04 ± .00	.47 ± .03	.48 ± .02	.43 ± .01	.63 ± .02	<u>.67 ± .01</u>	.74 ± .02	↑ .11
M_6	.55 ± .03	.14 ± .00	.72 ± .03	.65 ± .03	.65 ± .00	<u>.74 ± .00</u>	.65 ± .01	.93 ± .00	↑ .19
M_7	.11 ± .02	.02 ± .00	.12 ± .00	.13 ± .02	.09 ± .01	<u>.45 ± .01</u>	.19 ± .01	.91 ± .01	↑ .46
M_8	.43 ± .04	.09 ± .00	.60 ± .01	.57 ± .02	.65 ± .02	<u>.72 ± .01</u>	.68 ± .02	.79 ± .01	↑ .07
M_9	.31 ± .02	.07 ± .00	.46 ± .02	.48 ± .03	.43 ± .00	<u>.68 ± .00</u>	.56 ± .01	.95 ± .01	↑ .27
M_{10}	.50 ± .07	.13 ± .00	.70 ± .06	.65 ± .02	.62 ± .02	.74 ± .00	<u>.84 ± .01</u>	.93 ± .01	↑ .19
M_{11}	.38 ± .02	.16 ± .00	.49 ± .03	.50 ± .03	.42 ± .01	.60 ± .00	<u>.61 ± .06</u>	.91 ± .01	↑ .31
M_{12}	.34 ± .03	.07 ± .00	.48 ± .01	.47 ± .01	.44 ± .00	.26 ± .00	.69 ± .03	<u>.48 ± .03</u>	↑ .21
M_{13}	.19 ± .01	.04 ± .00	.29 ± .03	.28 ± .02	.21 ± .02	<u>.75 ± .02</u>	.32 ± .03	.97 ± .00	↑ .22
M_{14}	.29 ± .01	.10 ± .00	.42 ± .02	.39 ± .01	.43 ± .01	<u>.81 ± .01</u>	.53 ± .02	.85 ± .01	↑ .04
M_{15}	.15 ± .00	.01 ± .00	.17 ± .00	.15 ± .01	.16 ± .00	<u>.46 ± .00</u>	.31 ± .02	.78 ± .00	↑ .32
M_{16}	.51 ± .03	.08 ± .00	.66 ± .05	.58 ± .04	.56 ± .01	<u>.88 ± .01</u>	.61 ± .01	.95 ± .01	↑ .07
M_{17}	.35 ± .03	.12 ± .00	.43 ± .05	.40 ± .02	.36 ± .01	<u>.81 ± .01</u>	.44 ± .01	.89 ± .00	↑ .08
M_{18}	.39 ± .03	.08 ± .00	.47 ± .04	.48 ± .01	.44 ± .01	.52 ± .01	<u>.56 ± .01</u>	.67 ± .02	↑ .15
M_{19}	.41 ± .02	.07 ± .00	.51 ± .03	.45 ± .03	.44 ± .00	.35 ± .01	<u>.62 ± .03</u>	.90 ± .01	↑ .39
M_{20}	.41 ± .01	.11 ± .00	.51 ± .04	.47 ± .02	.43 ± .02	<u>.71 ± .01</u>	.64 ± .02	.75 ± .02	↑ .04
M_{21}	.20 ± .03	.01 ± .00	.20 ± .01	.24 ± .01	.23 ± .00	.48 ± .00	.37 ± .02	<u>.44 ± .01</u>	↓ .04

Table 7. Incremental evaluation NMI on Events2018. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	RPLM ^k _{SED}	RPLM ^d _{SED}	Improve
M_1	.49 ± .01	.34 ± .00	.54 ± .01	.57 ± .01	.57 ± .01	.70 ± .01	<u>.81 ± .01</u>	.89 ± .01	↑.19
M_2	.45 ± .00	.37 ± .00	.56 ± .02	.58 ± .01	.60 ± .01	.74 ± .01	<u>.77 ± .01</u>	.84 ± .01	↑.10
M_3	.56 ± .03	.37 ± .00	.52 ± .03	.57 ± .01	.62 ± .02	.64 ± .00	<u>.73 ± .01</u>	.76 ± .01	↑.12
M_4	.54 ± .03	.39 ± .00	.55 ± .01	.58 ± .03	.58 ± .01	.72 ± .01	.77 ± .01	<u>.75 ± .00</u>	↑.05
M_5	.54 ± .02	.53 ± .00	.58 ± .02	.61 ± .02	.59 ± .02	.64 ± .00	.72 ± .01	<u>.67 ± .01</u>	↑.08
M_6	.52 ± .02	.44 ± .00	.59 ± .03	.60 ± .01	.57 ± .01	.67 ± .00	.75 ± .01	<u>.73 ± .02</u>	↑.08
M_7	.56 ± .04	.42 ± .00	.63 ± .02	.64 ± .01	.67 ± .01	.78 ± .01	<u>.79 ± .00</u>	.88 ± .00	↑.10
M_8	.56 ± .03	.54 ± .00	.58 ± .02	.57 ± .02	.62 ± .02	.66 ± .02	.80 ± .01	<u>.77 ± .01</u>	↑.14
M_9	.54 ± .02	.45 ± .00	.48 ± .02	.52 ± .02	.48 ± .01	.57 ± .01	.61 ± .00	<u>.58 ± .02</u>	↑.04
M_{10}	.56 ± .06	.52 ± .00	.57 ± .01	.60 ± .01	.60 ± .01	.65 ± .01	.77 ± .01	<u>.77 ± .01</u>	↑.12
M_{11}	.59 ± .03	.48 ± .00	.54 ± .01	.60 ± .01	.54 ± .01	.63 ± .00	<u>.66 ± .00</u>	.68 ± .01	↑.05
M_{12}	.60 ± .02	.51 ± .00	.55 ± .04	.61 ± .02	.59 ± .02	.70 ± .01	.80 ± .01	<u>.77 ± .01</u>	↑.10
M_{13}	.61 ± .01	.44 ± .00	.60 ± .02	.59 ± .04	.64 ± .01	.67 ± .00	.70 ± .00	<u>.68 ± .00</u>	↑.03
M_{14}	.60 ± .02	.52 ± .00	.66 ± .01	.68 ± .02	.65 ± .00	.65 ± .00	<u>.69 ± .00</u>	.70 ± .00	↑.05
M_{15}	.57 ± .03	.49 ± .00	.60 ± .01	.63 ± .02	.60 ± .01	.70 ± .01	<u>.71 ± .01</u>	.73 ± .01	↑.03
M_{16}	.53 ± .02	.39 ± .00	.52 ± .02	.51 ± .03	.61 ± .01	<u>.75 ± .01</u>	.74 ± .02	.95 ± .02	↑.20

Table 8. Incremental evaluation AMI on Events2018. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	RPLM ^k _{SED}	RPLM ^d _{SED}	Improve
M_1	.48 ± .00	.11 ± .00	.54 ± .01	.56 ± .01	.55 ± .01	.70 ± .00	<u>.81 ± .01</u>	.89 ± .01	↑.19
M_2	.44 ± .02	.12 ± .00	.55 ± .01	.57 ± .01	.60 ± .01	.74 ± .00	<u>.77 ± .01</u>	.84 ± .00	↑.10
M_3	.55 ± .03	.11 ± .00	.55 ± .02	.56 ± .02	.62 ± .01	.64 ± .00	<u>.73 ± .01</u>	.77 ± .01	↑.13
M_4	.54 ± .04	.14 ± .00	.55 ± .01	.57 ± .03	.60 ± .01	.67 ± .00	.76 ± .01	<u>.75 ± .02</u>	↑.09
M_5	.53 ± .02	.24 ± .00	.57 ± .01	.59 ± .01	.57 ± .01	.64 ± .00	.70 ± .01	<u>.67 ± .01</u>	↑.06
M_6	.50 ± .03	.15 ± .00	.57 ± .02	.59 ± .01	.56 ± .01	.57 ± .00	.75 ± .01	<u>.72 ± .01</u>	↑.18
M_7	.55 ± .04	.12 ± .00	.61 ± .02	.63 ± .01	.62 ± .02	.67 ± .00	<u>.79 ± .00</u>	.87 ± .01	↑.20
M_8	.55 ± .02	.21 ± .00	.57 ± .02	.55 ± .02	.61 ± .01	.62 ± .00	.78 ± .02	<u>.76 ± .02</u>	↑.16
M_9	.48 ± .03	.16 ± .00	.46 ± .02	.46 ± .02	.47 ± .01	.52 ± .00	.60 ± .01	<u>.57 ± .02</u>	↑.08
M_{10}	.55 ± .04	.19 ± .00	.56 ± .02	.58 ± .01	.59 ± .01	.60 ± .00	<u>.75 ± .02</u>	.75 ± .01	↑.15
M_{11}	.57 ± .02	.18 ± .00	.53 ± .01	.59 ± .02	.52 ± .01	.54 ± .00	<u>.64 ± .00</u>	.68 ± .01	↑.14
M_{12}	.58 ± .02	.20 ± .00	.56 ± .02	.59 ± .02	.62 ± .02	.59 ± .00	.79 ± .01	<u>.77 ± .01</u>	↑.17
M_{13}	.59 ± .02	.15 ± .00	.60 ± .02	.58 ± .03	.65 ± .01	.64 ± .00	.69 ± .00	<u>.67 ± .00</u>	↑.04
M_{14}	.59 ± .01	.22 ± .00	.65 ± .00	.67 ± .02	.67 ± .01	.65 ± .00	<u>.68 ± .00</u>	.70 ± .01	↑.03
M_{15}	.55 ± .03	.22 ± .00	.58 ± .02	.61 ± .00	.57 ± .01	.65 ± .00	<u>.70 ± .01</u>	.71 ± .01	↑.06
M_{16}	.52 ± .02	.10 ± .00	.50 ± .01	.50 ± .03	.59 ± .01	.68 ± .00	<u>.73 ± .02</u>	.94 ± .02	↑.27

dataset. These are subsequently subjected to dimensionality reduction and visualization using t-SNE. Considering the long-tail distribution characteristic of social data, our attention is focused on the top seven events of each message block by quantity. From the visualization results obtained from the RPLM_{SED} (Figure 4 (d), (h), and (l)), it is clearly observable that event clusters exhibit stronger cohesion. Simultaneously, the separation between distinct event clusters is also more pronounced. This reflects the importance of the proposed clustering constraint,

which effectively enhances the distinguishability of message representations, thereby improving the accuracy of event detection.

Table 9. Incremental evaluation ARI on Events2018. The best results are bolded, and the second-best results are underlined.

Blocks	PP-GCN	EventX	KPGNN	QSGNN	FinEvent ^k	FinEvent ^d	RPLM _{SED} ^k	RPLM _{SED} ^d	Improve
M_1	.27 ± .03	.01 ± .00	.29 ± .02	.30 ± .01	.32 ± .02	.60 ± .01	<u>.85 ± .01</u>	.94 ± .01	↑ .34
M_2	.21 ± .01	.01 ± .00	.37 ± .01	.38 ± .02	.42 ± .02	.58 ± .01	<u>.79 ± .02</u>	.89 ± .01	↑ .31
M_3	.38 ± .03	.01 ± .00	.39 ± .04	.36 ± .02	.47 ± .01	.50 ± .02	<u>.76 ± .01</u>	.81 ± .00	↑ .31
M_4	.35 ± .05	.01 ± .00	.36 ± .04	.36 ± .01	.40 ± .03	.51 ± .02	.64 ± .01	<u>.64 ± .02</u>	↑ .13
M_5	.30 ± .01	.03 ± .00	.37 ± .02	.36 ± .02	.41 ± .01	.48 ± .01	<u>.49 ± .03</u>	.49 ± .01	↑ .01
M_6	.27 ± .02	.01 ± .00	.35 ± .04	.36 ± .01	.35 ± .02	.46 ± .02	.63 ± .02	<u>.60 ± .02</u>	↑ .17
M_7	.38 ± .05	.01 ± .00	.37 ± .02	.38 ± .02	.43 ± .01	.61 ± .01	<u>.79 ± .00</u>	.92 ± .01	↑ .31
M_8	.38 ± .04	.01 ± .00	.38 ± .02	.35 ± .02	.45 ± .02	.47 ± .01	.75 ± .02	<u>.67 ± .01</u>	↑ .28
M_9	.32 ± .04	.01 ± .00	.23 ± .02	.30 ± .04	.26 ± .01	.39 ± .01	.44 ± .01	<u>.41 ± .01</u>	↑ .04
M_{10}	.37 ± .04	.02 ± .00	.38 ± .02	.37 ± .02	.44 ± .03	<u>.52 ± .01</u>	.51 ± .02	.55 ± .00	↑ .03
M_{11}	.37 ± .04	.01 ± .00	.25 ± .02	.24 ± .03	.27 ± .01	<u>.40 ± .01</u>	.29 ± .01	.63 ± .02	↑ .23
M_{12}	.39 ± .03	.02 ± .00	.46 ± .02	.40 ± .04	.45 ± .01	<u>.57 ± .02</u>	.77 ± .01	.48 ± .01	↑ .20
M_{13}	.39 ± .01	.01 ± .00	.36 ± .05	.33 ± .03	.42 ± .03	.49 ± .04	<u>.61 ± .01</u>	.64 ± .01	↑ .15
M_{14}	.39 ± .06	.03 ± .00	.50 ± .03	.47 ± .03	.50 ± .03	.52 ± .02	.46 ± .02	<u>.50 ± .01</u>	↓ .02
M_{15}	.40 ± .06	.02 ± .00	.37 ± .02	.32 ± .02	.38 ± .02	.49 ± .01	<u>.51 ± .01</u>	.63 ± .01	↑ .14
M_{16}	.26 ± .03	.01 ± .00	.26 ± .02	.25 ± .02	.30 ± .03	<u>.46 ± .02</u>	.38 ± .03	.97 ± .01	↑ .51

5.2 Performance Evaluation on Low-Resource Language

Low-resource social event detection often faces the challenge of scarce training resources. Recent research introduced the CLKD [67] method, utilizing a knowledge distillation-based framework [26] to address this challenge. Specifically, CLKD initially employs pre-trained word embedding algorithms [9, 52, 62] to generate initial embeddings of messages, followed by mapping the data of low-resource languages into the word embedding space of high-resource languages through linear [5, 33] or non-linear [53] mapping techniques. Subsequently, a teacher model trained in a high-resource language (e.g., English) guides the student model in low-resource languages (e.g., Arabic) for learning and training. However, similar to other GNN-based methods, the backbone encoder of CLKD is still GNN. GNNs learn representations from explicit message graphs, within which some noisy edges are inevitably present. This results in the knowledge learned by the teacher model not being entirely effective and accurate, which may adversely affect the learning process of the student model. To validate our hypothesis, we conduct experiments on the low-resource Arabic-Twitter dataset. We choose TwHIN-BERT (TwHIN-BERT-large) as the backbone architecture for RPLM_{SED} and implement CLKD by using open-source code⁶. For other baseline methods, we use pre-trained word embedding algorithm⁷ to convert Arabic messages

Table 10. Evaluation on Arabic-Twitter. The best results are bolded, and the second-best results are underlined.

Metrics	PP-GCN	EventX	KPGNN	QSGNN	CLKD	FinEvent	RPLM _{SED}	Improve
NMI	.32 ± .02	.28 ± .00	.75 ± .02	.78 ± .02	.81 ± .02	<u>.82 ± .01</u>	.93 ± .01	↑ .11
AMI	.32 ± .02	.03 ± .00	.74 ± .01	.77 ± .01	.81 ± .02	<u>.82 ± .01</u>	.93 ± .01	↑ .11
ARI	.31 ± .02	.03 ± .00	.75 ± .02	.79 ± .02	.84 ± .02	<u>.86 ± .01</u>	.95 ± .01	↑ .09

⁶<https://github.com/RingBDStack/CLKD>

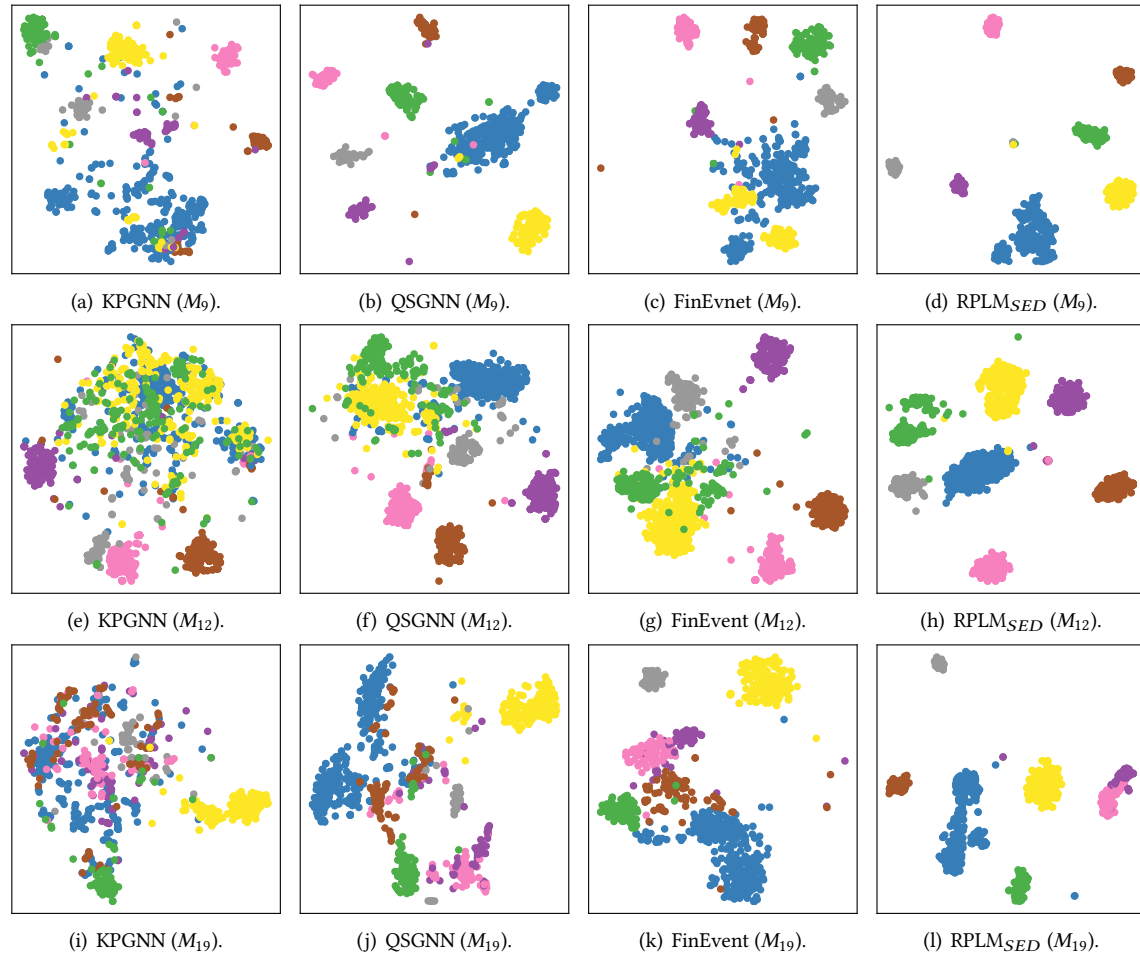


Fig. 4. Message representation visualization in the detection stage. Each row represents the t-SNE dimensionality reduction visualization of message representations for the same block obtained by different methods. Messages from the same event are represented with the same color.

into initial embeddings before training. To ensure a fair comparison, all methods utilize the K-Means clustering method.

As shown in Table 10, RPLM_{SED} outperforms the best results in the baseline by 13.41%, 13.41%, and 10.47% on NMI, AMI, and ARI, respectively. This is attributed to the inherent model advantage of RPLM_{SED} based on Multilingual Pretrained Language Models (MPLMs). Compared to GNNs, MPLMs can understand and encode different languages. This advantage enables RPLM_{SED} based on MPLMs to obtain better message embeddings when dealing with low-resource data, as well as stronger capabilities for encoding contextual messages. Secondly, the pairwise learning mechanism, by consistently encoding both the content and structural information of

⁷<https://github.com/bakrianoo/aravec>

messages, leverages more effective information. This enables the model to obtain robust and stable message representations, even in situations where training resources are limited.

5.3 Performance Evaluation on Long-tail Recognition Task

In real-world social data, the characteristics of a long-tail distribution are particularly evident. There are few head-dominated event classes, whereas low-frequency tail classes are more numerous. This imbalance in training data makes it difficult for tail events to be correctly classified. In recent research, ETGNN [66] employs the Dempster-Shafer theory [70] to integrate the uncertainty and temporal information of graphs constructed from different perspectives, thereby enhancing the accuracy of event detection. Furthermore, building upon ETGNN, the UCL_{SED} [68] framework is proposed. Compared to ETGNN, UCL_{SED} introduces uncertainty-guided contrastive learning loss, pushing boundary-ambiguous event classes further away to improve the distinguishability of representations. We also compare with MVGAN [16], which learns message representations from semantic and temporal views of messages and proposes a hashtag-based multi-view graph attention mechanism to fuse representations.

We conduct experiments on two long-tail datasets, Events2012_100 and Events2018_100. It is noteworthy that to accommodate the classification task, an event classifier is additionally incorporated into $RPLM_{SED}$. For the acquisition of message labels, we first filter out unreliable labels from their corresponding similarity sets based on similarity and then determine their final labels through a voting mechanism. In terms of evaluation metrics, we select Accuracy (ACC) and F1 value (F1), which are widely used in classification tasks. As illustrated in Table 11, compared with the optimal results among the baseline methods, $RPLM_{SED}$ exhibits an improvement of 4.34% in ACC and 5.43% in F1, respectively, on Events2012_100. On Events2018_100, the improvements are 19.92% and 16.46%, respectively. Recalling the pairwise message modeling process, we sample multiple messages for each message to construct message pairs. Utilizing the multi-relational prompt-based pairwise message learning mechanism, each message is learned multiple times from different message pairs. This allows messages of tail events to effectively learn a common representation of similar messages in terms of semantic and structural relations, as well as the subtle differences between different classes of messages. Additionally, the clustering constraint forces the representations of each event class to be as dispersed as possible in the latent space while ensuring that representations of messages within the same class are as close as possible. This further enhances the distinguishability of message representations, thereby enabling the event classifier to predict labels with greater accuracy. Furthermore, for each message, by filtering out unreliable labels from its corresponding set of candidate labels based on similarity, the label subsequently obtained through voting from the remaining labels is more reliable.

Table 11. Evaluation on the Long-tail Set. The best results are bolded, and the second-best results are underlined.

Dataset	Metrics	PPGCN	KPGNN	QSGNN	FinEvent	MVGAN	ETGNN	UCL_{SED}	$RPLM_{SED}$	Improve
Events 2012_100	ACC	.64 ± .00	.73 ± .01	.75 ± .01	.81 ± .01	.82 ± .00	.87 ± .01	<u>.92 ± .01</u>	.96 ± .01	↑ .04
	F1	.55 ± .00	.60 ± .00	.69 ± .01	.80 ± .01	.82 ± .01	.87 ± .00	<u>.92 ± .01</u>	.97 ± .01	↑ .05
Events 2018_100	ACC	.71 ± .01	.77 ± .01	.77 ± .01	.63 ± .00	.68 ± .01	.61 ± .01	<u>.79 ± .01</u>	.90 ± .01	↑ .11
	F1	.51 ± .00	.62 ± .01	.75 ± .01	.63 ± .01	.68 ± .01	.60 ± .01	<u>.79 ± .01</u>	.92 ± .01	↑ .12

5.4 PLM-based Model for Social Event Detection

In the preceding sections, we comprehensively compare $RPLM_{SED}$ against the currently dominant social event detection methods based on GNNs across various scenarios. To our knowledge, $RPLM_{SED}$ is the inaugural model

leveraging PLMs to achieve high-quality social event detection in diverse scenarios, including offline, online, low-resource settings, and long-tail distributions. To further explore the performance of Pre-trained Language Models (PLMs) in the social event detection tasks, we design experiments utilizing the pre-trained word embedding technology Word2Vec and two mainstream PLMs, BERT and RoBERTa. Specifically, for Word2Vec, BERT, and RoBERTa, which only use the message content as input, the final representation is obtained by averaging the embeddings of each word in the message. Additionally, we implement two distinct strategies for the original BERT and RoBERTa without fine-tuning and fine-tuning strategies. Furthermore, we propose variant models BERT-LR and RoBERTa-LR for RPLM_{SED} based on BERT and RoBERTa, respectively, which incorporate fine-tuning through LoRA. In practical applications, given the infeasibility of the K-Means clustering method due to its requirement for pre-specifying the number of events, we consistently employed the HDBSCAN clustering method. It is important to note that RPLM_{SED} is predicated on RoBERTa-base, and similarly, all other models also utilize their respective base versions.

As illustrated in Figure 5, among the methods for directly obtaining message representations from Word2Vec, BERT, and RoBERTa, BERT demonstrates the relatively best performance. This may be attributed to the higher compatibility between the knowledge acquired during the BERT pre-training process and the Events2012 dataset in the absence of model fine-tuning. Upon fine-tuning, BERT-FT and RoBERTa-FT exhibit significant enhancements, achieving nearly equivalent model performances. Nonetheless, a discernible performance gap remains when compared to RPLM_{SED}, attributable to their exclusive reliance on the semantic information of messages while neglecting the abundant structural information within the messages. Moreover, we observe that the performances of the variants BERT-LR and RoBERTa-LR fall short of those achieved through global fine-tuning of the models. In online event detection scenarios, models necessitate continual updates to incorporate new knowledge while concurrently sustaining an existing knowledge base. Utilizing LoRA to fine-tune models can effectively decrease the number of parameters requiring adjustments during the fine-tuning process, which aids in enhancing the efficiency of model updates and reducing computational costs. However, this strategy of reducing the volume of trainable parameters may adversely affect the model's long-term capabilities for knowledge learning, retention, and expansion, consequently leading to a decline in model performance. Overall, RPLM's model performance significantly surpasses that of other methods, further illustrating the efficacy of the multi-relational prompt-based pairwise message learning mechanism. On the other hand, the experiments also validate the feasibility of utilizing PLMs for social event detection tasks.




Fig. 5. Overall Performance of Various PLM-based Models on Events2012 in Online Scenario. The horizontal axis represents the average metric values of different methods, and the best results are bolded.

5.5 RPLM_{SED} with Different PLMs

In investigating the variations in the performance of RPLM_{SED} under different PLMs, we conduct experiments on the Events2012 dataset utilizing four distinct models: RoBERTa-large [43], RoBERTa-base, Twitter-RoBERTa-large [47], and BERT-base-uncased [30]. The experiments are designed to encompass three dimensions: model

28 • P. Li et al.

scale, pre-training corpora, and model architecture. To ensure consistency in comparison, the HDBSCAN method is utilized for clustering in all experiments.

As depicted in Figure 6, compared to $RPLM_{SED}$ based on RoBERTa-base, the model based on RoBERTa-large exhibits overall improvements of 3.23%, 3.45%, and 4.68% on the NMI, AMI, and ARI, respectively. Compared to



Fig. 6. $RPLM_{SED}$ performance comparison based on different PLMs. The horizontal axis represents the message blocks, while the vertical axis denotes each message block's corresponding metric value. Each message block's background color corresponds to the model's color with the highest score within that block.

, Vol. 1, No. 1, Article . Publication date: September 2024.

RoBERTa-base, RoBERTa-large possesses a larger model size, more trainable parameters, and a broader training dataset, resulting in superior model performance. Naturally, the RPLM_{SED} based on RoBERTa-large requires greater computational resources and training time. The primary distinction between Twitter-RoBERTa-large and RoBERTa-large lies in the difference in training corpora. Twitter-RoBERTa-large is specifically trained on Twitter data, whereas RoBERTa-large undergoes pre-training on a broader spectrum of text data. These two models have the same model scale and have been further fine-tuned on the same dataset, hence their overall performance is essentially consistent. When contrasting RoBERTa-base with BERT-base-uncased, two distinct types of pre-trained language models, we observe that RPLM_{SED} based on RoBERTa-base exhibits superior performance. In terms of overall performance, the RPLM_{SED} based on RoBERTa-base surpasses the RPLM_{SED} based on BERT-base-uncased by 4.36%, 4.81%, and 12.15% in NMI, AMI, and ARI, respectively. This advantage can be attributed to the pre-training of the RoBERTa-base on a larger dataset and optimized training strategies. Notably, even the RPLM_{SED} based on BERT-base-uncased outperforms the strongest baseline model, FinEvent, leading by 2.81% in NMI, 2.51% in AMI, and 8.75% in ARI.

5.6 Time Consumption of RPLM_{SED} based on Different Scale PLMs

This section reports the time consumption of different parts of RPLM_{SED} based on RoBERTa-large and RoBERTa-base on the Events2012. Figure 7 shows the message number of each block, the time consumption for pairwise messages modeling during the detection phase, and the execution time for detection operations. M_0 is solely used for initial model training and incurs no time consumption during detection. Additionally, Table 12 presents statistics on message numbers within each window, the time required for pairwise message modeling during training, and the time consumption for model training or maintenance in each window. Both Figures 7 and Table 12 show that the time expenditure for model training or detection correlates linearly with the number of messages in message blocks. Secondly, whether in the training or detection phase, the time spent on pairwise message modeling is significantly less than the time spent on model training or detection. Furthermore, the training and detection time of RPLM_{SED} based on RoBERTa-large is approximately 3.2 times that of RPLM_{SED} based on RoBERTa-base. This is attributed to RoBERTa-large's larger hidden dimensions and more Transformer layers, which enhance its performance compared to RoBERTa-base-based RPLM_{SED}. It should be noted that in addition to the employed PLM scale, factors such as message sampling numbers and hardware environment also markedly influence the overall time consumption of the model.

Table 12. The number of messages in each window and the time consumption of pairwise modeling and training (or maintenance) for each window in Events2012.

Message Number								
Window	M_0	$M_1 - M_3$	$M_4 - M_6$	$M_7 - M_9$	$M_{10} - M_{12}$	$M_{13} - M_{15}$	$M_{16} - M_{18}$	$M_{19} - M_{21}$
Number	20254	12048	5120	8201	5565	7477	5473	4702
Time Consumption (min)								
Window	M_0	$M_1 - M_3$	$M_4 - M_6$	$M_7 - M_9$	$M_{10} - M_{12}$	$M_{13} - M_{15}$	$M_{16} - M_{18}$	$M_{19} - M_{21}$
Pairwise Modeling	4.55	3.53	0.52	1.52	0.57	1.20	0.79	0.47
RoBERTa-base	55.06	32.88	14.73	22.28	15.08	19.95	14.7	12.58
RoBERTa-large	175.20	108.25	46.35	72.24	50.78	65.23	48.80	40.45




Fig. 7. The number of messages in each message block (histogram) and the time consumption of pairwise modeling and detection for each block (line chart) in Events2012.

5.7 Ablation Study

In this section, we conduct ablation experiments on four critical components of the $RPLM_{SED}$ model to demonstrate the indispensability of each component. To ensure consistency of results, all experiments are conducted on the Events2012 dataset and utilize the HDBSCAN clustering method. To investigate the impact of different components on model performance, we sequentially conduct ablation experiments on four parts: **Multi-relational Prompt Embedding**, **Similarity-based Representation Aggregation**, **Multi-Head Structured Attention**, and **Clustering Constraint**. The results of these experiments are reported in Table 13.

Multi-relational Prompt Embedding. As illustrated in ‘Prompt Embedding’ within the ‘Pairwise Message Learning’ part of Figure 3. The task of multi-relational prompt embedding is to transform the structural relations between messages into multi-relational prompt embeddings. When removing the multi-relational prompt embedding, we directly sample messages from the original message data to form message pairs. As shown in Table 13, after removing the multi-relational prompt embedding, the model experiences a decrease of 2.05%, 2.08%, and 2.07% in NMI, AMI, and ARI, respectively. Social messages contain rich structural information. After removing the multi-relational prompt embedding, due to the absence of multi-relational prompt guidance, the model can only learn representations from the semantics of messages, thereby affecting model performance.

Similarity-based Representation Aggregation. We investigate the variations in model performance under different representation aggregation methods, including similarity-based representation aggregation method, top-k sampling representation aggregation method, and all representation average aggregation method. For a set of candidate feature representations obtained for each message, under the top-k sampling representation aggregation method, we only sample the top 20 candidate representations with the highest similarities for averaging. This approach has certain limitations. A smaller value of k may lead to message representations that are not comprehensive enough. Conversely, a larger k may include candidate representations extracted from negative sample pairs, which often carry noise and irrelevant information, thereby affecting the quality of the final message representations. For the all-representation average aggregation method, averaging all candidate representations of messages may lead to overly smoothed representations and is also susceptible to noise from candidate representations extracted from negative sample pairs. Consequently, among the three methods, it performs the worst. In contrast, the similarity-based representation aggregation approach averages candidate representations of messages with high confidence, effectively filtering out noisy candidate representations

extracted from negative sample pairs while preventing message representations from becoming overly smoothed. Overall, the similarity-based representation aggregation approach outperforms the top-k sampling representation aggregation by 1.04%, 1.05%, and 2.89% in NMI, AMI, and ARI, respectively. Additionally, it surpassed the all representation average aggregation method by 3.68%, 3.84%, and 8.16% on NMI, AMI, and ARI, respectively.

Multi-Head Structured Attention. The role of the multi-head structured attention mechanism is to extract key features from different subspaces of message representations while eliminating irrelevant features and noise to enhance the clarity of the representations. When the multi-head structured attention mechanism is removed, the resulting candidate message representations may contain noise, affecting the classifier’s accuracy in discriminating message similarity. Additionally, it may also impact the quality of the final message representations. Consequently, when the multi-head structured attention mechanism is removed, the values of NMI, AMI, and ARI decrease by 2.06%, 2.31%, and 2.31%, respectively. This reduction substantiates the pivotal role of the attention mechanism in bolstering both representation clarity and classifier precision.

Clustering Constraint. Our proposed clustering constraint consists of inter-cluster and intra-cluster loss, which enhances the distinguishability of message representations by optimizing the training process. Upon removing the clustering constraint, we find that relying solely on pairwise cross-entropy loss to constrain model training is insufficient. This insufficiency stems from the fact that pairwise cross-entropy loss merely treats messages as objects to be pulled closer or pushed apart, making it difficult to ensure cohesiveness within event clusters and separation between them. After the removal of the Clustering Constraint, the NMI, AMI, and ARI decreased by 6.29%, 6.71%, and 9.25%, respectively.

Table 13. Ablation study for Multi-relational Prompt Embedding, Similarity-based Representation Aggregation, Multi-Head Structured Attention, and Clustering Constraint. The best results are bolded, and the second-best results are underlined.

	Multi-relational Prompt Embedding	Representation Aggregation Strategy	Multi-head Structured Attention	Clustering Constraint	Avg. Metrics		
					NMI	AMI	ARI
1	✓	<i>Similarity</i>	✓	✓	0.874	0.865	0.822
2	-	<i>Similarity</i>	✓	✓	0.856	0.847	<u>0.805</u>
3	✓	<i>Top – k</i>	✓	✓	<u>0.865</u>	<u>0.856</u>	0.799
4	✓	<i>Average</i>	✓	✓	0.843	0.833	0.760
5	✓	<i>Similarity</i>	-	✓	0.856	0.845	0.803
6	✓	<i>Similarity</i>	✓	-	0.819	0.807	0.746

5.8 Hyperparameter Study

We explore the impact of eight hyperparameters on model performance in the online detection scenario. All experiments are conducted on the Events2012 dataset and utilize the HDBSCAN clustering method. These hyperparameters include **Intra-cluster Loss Weight (λ)**, **Inter-cluster Loss Weight (μ)**, **Pairwise Cross-Entropy Loss Weight (κ)**, **Similarity Threshold (α)**, **Number of Structured Attention Heads (σ)**, **Parameter Update Weight (ζ)** during the maintenance phase, **Number of message sampled (y)**, and **Window Size (w)**. To ensure experimental consistency, all parameters except the one under study are kept fixed, with their initial values set to 0.005, 0.01, 1, 0.9, 2, 0.4, 60, and 3, respectively. Due to computational resource constraints, all hyperparameter experiments are conducted under the RPLM_{SED} based on RoBERTa-base. The experimental results are shown in Figure 8.

Intra-cluster Loss Weight (λ). When examining the effect of different intra-cluster loss weights on model performance, we select four weights: 0.001, 0.005, 0.01, and 0.015. As shown in Figure 8(a), the model performance remains stable when λ is set to 0.001 or 0.005. However, as λ increases to 0.01 or higher, a linear decline in model performance is observed. This suggests that an excessively high intra-cluster loss weight causes the model to

overly focus on making representations of the same event type more cohesive, neglecting the importance of separating different event types, thereby negatively impacting the final event detection accuracy.

Inter-cluster Loss Weight (μ). We also choose 0.001, 0.005, 0.01, and 0.015 as candidate values for the inter-cluster loss weight. As shown in Figure 8(b), when μ is set too high (such as 0.015) or too low (such as 0.001), model performance decreases. Especially when μ is set too low, the model lacks sufficient constraint to push apart the centroids of different clusters, resulting in event clusters not being well separated. Conversely, when μ is set too high, the model places more emphasis on separating different event clusters, leading to a slight decrease in the cohesiveness of messages within clusters, thus mildly affecting the model's performance.

Pairwise Cross-Entropy Loss Weight (κ). For pairwise cross-entropy loss weight, experiments are conducted on four values: 0.6, 0.8, 1.0, and 1.2. Figure 8(c) shows a balanced allocation among the weights of pairwise cross-entropy loss, inter-cluster loss, and intra-cluster loss is crucial. The model's overall performance decreases when the parameter κ is set to 0.6 or 0.8. This is attributed to the insufficient constraint imposed by the pairwise cross-entropy loss and a too-small weight, which inversely amplifies the constraints of inter-cluster loss and intra-cluster loss. This may cause similarity scores between messages obtained from the classifier to be inaccurate, thereby affecting the quality of the final message representation. Conversely, an excessively large κ value weakens the constraints imposed by both inter-cluster and intra-cluster losses. This, in turn, leads to message representations within the same event not being sufficiently tight and blurring the boundaries between different event clusters, thereby affecting the model's performance.

Similarity Threshold (α). In studying the similarity threshold, we test four thresholds: 0.1, 0.3, 0.6, and 0.9. Figure 8(d) shows that model performance improves with increasing α values. This phenomenon can be attributed to a higher α value implying stricter requirements for message pair similarity. Consequently, the selected candidate representations are more likely to be extracted from message pairs within the same cluster. This leads to more stable and robust aggregated message representations. In addition, the choice of a reasonable similarity threshold is also crucial. Setting α too high may result in message representations that lack comprehensiveness.

Number of Heads in Structured Attention Mechanism (o). The experiment examines settings of 1, 2, 3, and 4 attention heads. As illustrated in Figure 8(e), when there is only one attention head, the model is limited in capturing key message features, resulting in limited message representations and consequently affecting model performance. Conversely, too many attention heads may introduce noise, leading to a decline in message representation quality and impacting model performance.

Parameter Update Weight during Model Maintenance (ζ). We select 0.2, 0.4, 0.6, and 0.8 as candidate values for the parameter update weight. As shown in Figure 8(f), the model performs best when ζ is set to 0.4. When ζ is set to 0.2, there is a slight decrease in model performance, possibly due to the model retaining too much irrelevant historical information or information unrelated to current events. Conversely, when ζ values are too high (such as 0.6 and 0.8), the model retains too little historical knowledge, potentially making it difficult for the model to recognize and understand the connection between current and historical events, leading to a performance decline.

Number of message sampled (y). We investigate the impact of varying y on model performance, including four different values: 20, 40, 60, and 80. As shown in Figure 8(g), both excessively large and small y values decrease model performance. For events with a large number of messages, sampling too few may result in insufficiently comprehensive learned message representations. Conversely, social media messages have the long-tail distribution characteristic. When the sampling size is excessively large, the model may overfit the head events with a large number of messages, which neglects the representation learning of messages from long-tail events.

Window Size (w). The window size determines the extent of message blocks included during continuous detection and the message block numbers used for model maintenance. We analyze the impact of four different window sizes: 1, 2, 3, and 4. As observed in Figure 8(h), both excessively small (with w set to 1 or 2) and large window sizes (with w set to 4) result in varying degrees of decline in model performance. Smaller windows

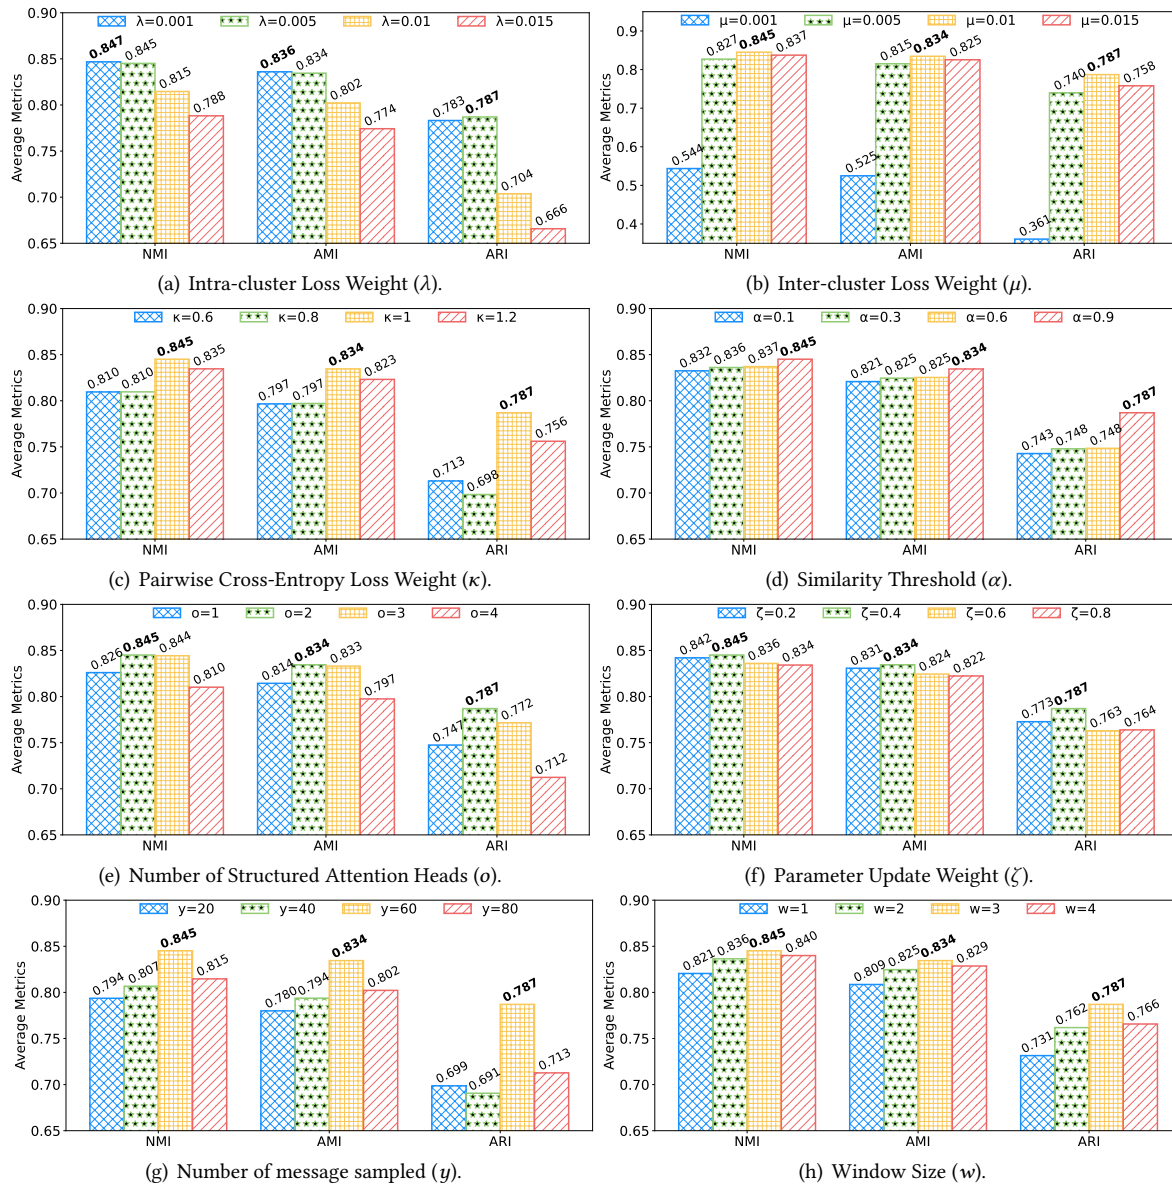


Fig. 8. RPLM_{SED} with different hyperparameters. The horizontal axis represents the metrics, while the vertical axis denotes each hyperparameter's corresponding average metric values. The best results are bolded.

are likely to contain more newly emerging or less frequently discussed events. These events' relevant message counts are relatively less, making the contextual information insufficiently clear. Consequently, during the model maintenance phase, the learned message representations of these events may not be comprehensive. In subsequent window message blocks, when the message numbers describing these events increase, the model struggles to recognize them effectively due to insufficient learning of these events, thereby affecting overall performance.

Moreover, the smaller window requires more frequent model maintenance. This may accelerate the forgetting of historical knowledge and lead to the overfitting of current knowledge. Conversely, when the window size is excessively large, some events may have already dissipated in subsequent message blocks. Continuing to fit messages related to these events during model maintenance is unnecessary and may introduce noise, thereby compromising the model's performance.

5.9 Case Study

To investigate the impact of message structural information on GNN-based methods and proposed $RPLM_{SED}$, we select M_{19} from Events2012 for comparison and analysis with the optimal baseline method FinEvent. The message graph of M_{19} contains more noisy edges and isolated nodes, and the structural information is also relatively scarce. This graph comprises 1,399 message nodes, including 92 isolated nodes, across 28 events, and possesses 21,374 edges based on entity relations (with 8,862 edges connecting messages from different events), 83 edges based on user relations (with 63 edges connecting messages from different events), and 623 edges based on hashtag relations (with 196 edges connecting messages from different events).

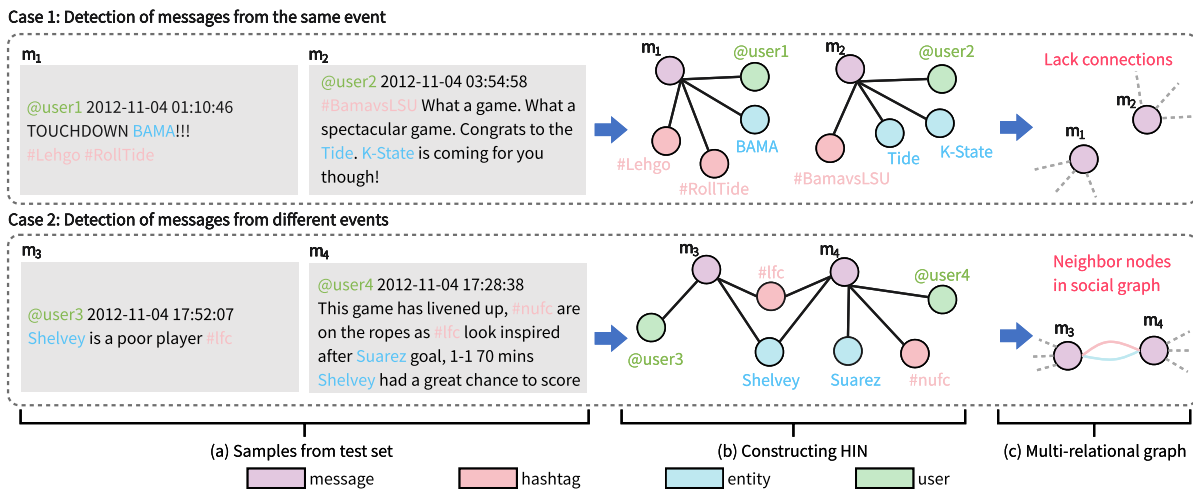


Fig. 9. Raw Messages of M_{19} within Events2012.

Firstly, we extract two pairs of messages from M_{19} where FinEvent predicts incorrectly while RPLM predicts accurately. For the 'case1' part in Figure 9, m_1 and m_2 , posted by different users, pertain to positive reactions towards the University of Alabama's performance in a specific game. The first message is focused on celebrating a particular scoring moment, while the second provides a broader perspective on the game and mentions a forthcoming challenge. Although m_1 and m_2 describe the same event, they share no common attributes other than belonging to the same time segment. Unfortunately, FinEvent utilizes time by transforming it into temporal features combined with message content, resulting in no connections between these two messages in the constructed multi-relational message graph. This results in FinEvent's inability to aggregate and learn effective information about both from the explicit message graph directly, thereby failing to group them correctly. Moreover, in the 'Case 2' part from Figure 9, m_3 articulates an opinion regarding a player, whereas m_4 delineates a football match. These two messages from different events share the same hashtag '#lfc' and entity 'Shelvey', which, in turn, leads to their erroneous classification together by FinEvent. In contrast, $RPLM_{SED}$ does not rely on the explicit structural relations between messages to learn representations. It transforms the structural relations

between messages into multi-relational prompt embeddings to complement message content embeddings, thereby enabling the joint utilization of message content and structural information. During the learning process, by consistently encoding the structural relations and semantic content between messages, it can better learn and understand the complex structural relations between messages and the interaction between structural relations and semantic content. In the absence of common attributes or the presence of misleading information, $RPLM_{SED}$ can still accurately group them by capturing the semantic and contextual relations between the messages.

Additionally, we construct confusion matrices for both RPLM and FinEvent, comparing their predicted labels with the true event labels, thereby observing the accuracy of event detection for each. Given the discrepancy between the number of events obtained from density-based clustering methods and the actual number of events, we uniformly adopt the K-Means method to obtain predicted labels. As shown in Figure 10, we observe that the clustering results of $RPLM_{SED}$ (Figure 10(a)) generally have higher values on the diagonal compared to FinEvent (Figure 10(b)), and its prediction results are more concentrated. This indicates that the message representations obtained by $RPLM_{SED}$ have higher accuracy, whereas FinEvent is more disturbed by noise during the feature aggregation process, leading to less clear feature boundaries. It is noteworthy that FinEvent is almost unable to identify events with very few messages (events numbered 17 to 27).




Fig. 10. Confusion matrices of M_{19} within Events2012.

5.10 Clustering Algorithm Study

This section investigates the performance of three density-based clustering algorithms, **OPTICS** [4], **DBSCAN** [19], and **HDBSCAN** [10], on message representation obtained by $RPLM_{SED}$. These algorithms are widely used in real-world applications due to their ability to cluster data without requiring a predefined number of event clusters. Specifically, the core of OPTICS is to generate an ordered list of points by analyzing the density

and mutual relationships of data points, thus identifying clustering structures at various scales. OPTICS is particularly sensitive to the parameters of *Minimum Points* and *Cluster Reachability*. The parameter of *Minimum Points* determines the minimum number of points in a cluster. Setting this parameter too low may lead to the formation of numerous small clusters, whereas setting it too high may cause small clusters to be merged. *Cluster Reachability* evaluates the density changes at cluster boundaries to decide whether to split regions into separate clusters. Setting this parameter too low may blur cluster boundaries, causing regions with small density differences to be merged; setting it too high could result in over-segmentation of regions with minor density variations into multiple clusters.

DBSCAN forms clusters by defining core, border, and noise points. It is highly sensitive to the parameters of *Epsilon Neighbourhood* and *Minimum Points*. A point is considered a core point if the number of points in its neighborhood is greater than or equal to *Minimum Points*; otherwise, it is classified as a border or noise point. The algorithm starts from an unvisited core point and identifies all points within its neighborhood to form a new cluster. It recursively applies the same process to the core points within the neighborhood until the cluster can no longer be expanded. When the *Epsilon Neighbourhood* is set too small, a single cluster may be divided into multiple sub-clusters, instead, set too large, different clusters may be merged. Setting the *Minimum Points* too low may result in too many core points and ignoring noise points; setting it too high may result in only high-density regions being recognized as clusters, leading to more points being classified as noise.

Compared to DBSCAN, HDBSCAN automatically adjusts the density threshold to recognize clustering structures at varying densities. It constructs a density hierarchy tree to represent data structures at different densities and eventually prunes the tree to obtain the optimal clustering result. The key parameters for HDBSCAN are the *Minimum Cluster Size* and the *Minimum Sample Number*. The *Minimum Cluster Size* specifies the minimum number of points required for a cluster; setting it too low increases the number of noise points, while setting it too high may ignore small clusters. The minimum number of samples determines core points and influences cluster shapes. Typically, the *Minimum Cluster Size* and the *Minimum Sample Number* are set to be equal.

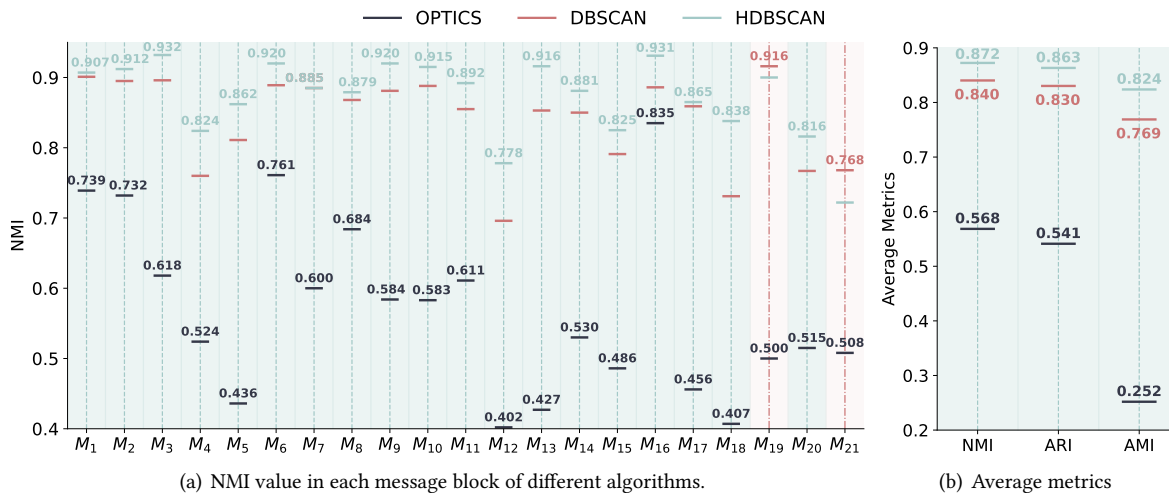


Fig. 11. Performance of different clustering algorithms on Events2012. (a) The horizontal axis represents the message blocks, while the vertical axis denotes each message block's corresponding NMI value. (b) The horizontal axis represents the metrics, and the vertical axis denotes each metric's corresponding average value. Each message block (metric)'s background color corresponds to the algorithm's color with the highest score with that block (metric).

Experiments are conducted on the Events2012 dataset in the online detection scenario, in which metrics (NMI, AMI, ARI) for each message block under different algorithms are calculated. The *Minimum Points* and *Cluster Reachability* parameters for OPTICS are set to 8 and 0.2, respectively. For DBSCAN, the *Epsilon Neighbourhood* and *Minimum Points* are set to 5 and 8, respectively. Additionally, the *Minimum Cluster Size* for HDBSCAN is set to 8. Figure 11 demonstrates that HDBSCAN is most suitable for the message representations learned by RPLM_{SED}. This is due to HDBSCAN requiring adjustment of only one parameter, the *Minimum Cluster Size*. In contrast, OPTICS and DBSCAN are more sensitive to multiple parameters. The distribution of data points across different message blocks may vary. Therefore, setting the same *Cluster Reachability* (OPTICS) or *Epsilon Neighborhood* (DBSCAN) parameters is not applicable to all message blocks.

6 RELATED WORK

This section provides an overview of the existing literature and methodologies relevant to the proposed RPLM_{SED}. The discussion is structured into three categories: Social Event Detection, Pre-trained Language Models, and Prompt Learning.

6.1 Social Event Detection

Social event detection represents a pivotal research domain, extensively categorized based on its objectives and application scenarios. In terms of goal-oriented classification, social event detection methods are primarily divided into document-pivot (DP) [1, 11, 29, 39, 57, 60, 65, 85, 90] and feature-pivot (FP) [21, 23] approaches. DP methods focus on clustering social messages by analyzing their interrelations, emphasizing the grouping of relevant messages for a deeper understanding of the development and evolution of societal events. These methods are particularly crucial in analyzing specific domain events such as disasters, stock trading, major sports events, and political activities, as they aid in organizing and comprehending the historical information of these events. The proposed RPLM_{SED} method in this paper is an exemplar of this category. In contrast, FP methods cluster messages based on the distribution characteristics of social message elements, such as vocabulary and named entities. They focus on the co-occurrence and distribution patterns of message elements to reveal inter-element correlations. Furthermore, social event detection methods can be categorized based on their application scenarios into offline [57] and online methods [11, 21, 24, 29, 39, 59, 60, 65, 84]. Offline methods are primarily used for analyzing historical events, whereas online methods continuously monitor social data streams, tracking the real-time development of events. Online methods are crucial for rapid response and real-time event monitoring, especially in situations requiring swift action. Additionally, based on the employed techniques and mechanisms, these methods can be further subdivided into incremental clustering methods [1, 29, 56, 85], community detection methods [21, 39–41, 44, 72, 83], and topic modeling methods [7, 8, 13, 14, 58, 90]

Recently, methods based on Graph Neural Networks (GNNs) [11, 16, 57, 59, 60, 65–67, 67, 73] have demonstrated potential in the domain of social event detection. By modeling social messages as graphs, these approaches effectively integrate messages' semantic and structural information, overcoming the limitations of insufficient information utilization within traditional methods. For instance, KPGNN [11] converts messages with common attributes into homogeneous message graphs, which achieves incremental social event detection by leveraging the inductive learning capabilities of GNNs. FinEvent [60] reduces the impact of noisy edges between message nodes by constructing weighted multi-relational graphs and utilizes reinforcement learning to determine optimal relation weights. Additionally, QSGNN [65] employs utilizing the finest known samples and dependable knowledge transfer to extend knowledge from the known to the unknown, markedly diminishing the requirement for manual annotation in model training. However, GNN-based methods learn representations from explicit message graphs. These graphs may contain noisy edges, and messages from the same event without common attributes lack connections. This leads to GNNs being impacted when learning message representations, an

effect that is exacerbated in the presence of sparse structural information. Moreover, these methods initialize message embeddings statically before training, preventing the model from dynamically obtaining better message embeddings during the training phases. In recent research, Cao et al. [12] propose HISEvent, which explores the neighborhood of a message graph by minimizing 1-dimensional Structural Entropy (1D SE) and supplementing edges in the existing message graph between semantically close messages. Subsequently, by minimizing 2-dimensional structural entropy (2D SE) for partitioning the message graph, the approach eliminates the cost of manual annotation, achieving unsupervised social event detection. It is imperative to acknowledge that although HISEvent employs a hierarchical approach to mitigate the computational complexity of conventional vanilla greedy 2D SE minimization algorithm [36] from $O(N^3)$ to $O(n^3)$ (where N and n respectively denote the number of nodes within a message graph and that within a manually selected sub-graph), methods based on GNNs or PLMs exhibit a computational complexity of merely $O(N)$. Given that HISEvent is an unsupervised approach, it has not been compared in this paper. In addition, Our proposed 'pairwise' approach significantly differs from that introduced in PPGCN [57]. In PPGCN, after obtaining representations, the pairwise sampling method selects a positive and a negative sample for each message, addressing the challenge of numerous event types with few messages per event. Unlike PPGCN, our pairwise message modeling strategy is implemented during the data preprocessing stage to preserve structural relations rather than only focusing on inter or intra-event relations.

6.2 Pre-trained Language Models

In recent years, Pre-trained Language Models (PLMs) [30, 35, 43, 47, 64, 87] have made significant strides in the field of Natural Language Processing (NLP). These models have learned rich language representations by undergoing pre-training on large-scale text corpora. When applied to downstream tasks, fine-tuned PLMs can adapt to a variety of application scenarios, including, but not limited to, text classification [30, 74], question answering [71, 78], language generation [64], and text summarization [35]. In processing social data, PLMs have demonstrated high adaptability to informal texts and multiple languages [47, 87], effectively handling texts containing internet-specific abbreviations, emojis, or slang. This has extended the application of PLMs to social event detection tasks. However, previous social event detection approaches often treated PLMs merely as a tool for word embedding, extracting message embeddings through PLMs and then applying clustering algorithms for event detection. This overlooks the capabilities of PLMs to learn, retain, and expand knowledge, leading to suboptimal performance in the complex environment of social media data streams. In contrast to prior studies, we leverage PLM as the backbone architecture and fine-tune it during the initial training and maintenance phases, thereby achieving high-quality detection of social events.

6.3 Prompt Learning

Prompt learning [34, 42], as an emerging technology in NLP, has garnered widespread attention and has been applied to a variety of tasks, including recommendation systems [88], aspect sentiment triplet extraction [61], text generation [37], question answering [63], etc. The core of prompt learning lies in the design of prompt templates, which are typically categorized into discrete templates, continuous templates, and hybrid templates. Discrete templates employ explicit natural language terms, making them direct, easy to understand, and implement, but they also rely on human experience. For instance, Petroni et al. [63] design specific sentence structures as prompts to query the PLMs about their knowledge of certain facts. Contrary to discrete templates, continuous templates do not directly carry semantic information but consist of a set of learnable embedding vectors [18, 25]. By continuously optimizing these prompts throughout the training process, the model is better adapted to the task. In our proposed RPLM_{SED}, the multi-relational prompt embeddings are a type of continuous prompt template. Hybrid templates combine discrete and continuous templates, aiming to leverage human expertise and model auto-optimization to achieve superior performance and broader adaptability [88]. Unlike traditional prompt

learning, to ensure that $RPLM_{SED}$ can adapt to dynamic social streams, we utilize multi-relation prompts and fine-tune the model's parameters during the training or maintenance phases.

7 CONCLUSION

This paper studies $RPLM_{SED}$ relational prompt-based pre-trained language models for Social Event Detection. A pairwise message modeling strategy is proposed to address the issues of missing and noisy edges in social message graphs. A multi-relational prompt-based pairwise message learning mechanism is presented to simultaneously leverage the content and structural information of messages, learning more robust and stable message representations. A clustering constraint is designed to optimize the training process, enhancing message representations' distinguishability. The conducted experiments on three real-world social media datasets demonstrate that $RPLM_{SED}$ achieves SOTA performance in offline, online, low-resource, and long-tail distribution social event detection tasks. In the future, we aim to explore how supervised $RPLM_{SED}$ can be extended to semi-supervised or unsupervised social event detection tasks while also investigating ways to optimize the model architecture further to achieve more efficient detection.

ACKNOWLEDGMENTS

This work is supported by the NSFC through grants 62266028, 62322202, U21B2027, and 62432006, Yunnan Provincial Major Science and Technology Special Plan Projects through grants 202402AD080002, 202302AD080003 and 202202AD080003, Key Projects of Basic Research in Yunnan Province through grants 202301AS070047, 202301AT070471.

REFERENCES

- [1] Charu C Aggarwal and Karthik Subbian. 2012. Event detection in social streams. In *Proceedings of the 2012 SIAM international conference on data mining*. 624–635.
- [2] Alaa Alharbi and Mark Lee. 2021. Kawarith: an Arabic Twitter corpus for crisis events. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*. Association for Computational Linguistics, 42–52.
- [3] Hadi Amiri and Hal Daume III. 2016. Short text representation for detecting churn in microblogs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30. 1–7.
- [4] Mihael Ankerst, Markus M Breunig, Hans-Peter Kriegel, and Jörg Sander. 1999. OPTICS: Ordering points to identify the clustering structure. *ACM Sigmod record* 28, 2 (1999), 49–60.
- [5] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 789–798.
- [6] Farzindar Atefeh and Wael Khreich. 2015. A survey of techniques for event detection in twitter. *Computational Intelligence* 31, 1 (2015), 132–164.
- [7] Hila Becker, Mor Naaman, and Luis Gravano. 2011. Beyond trending topics: Real-world event identification on twitter. In *Proceedings of the international AAAI conference on web and social media*, Vol. 5. 438–441.
- [8] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.
- [9] Piotr Bojanowski, Édouard Grave, Armand Joulin, and Tomáš Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [10] Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. 2013. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*. Springer, 160–172.
- [11] Yuwei Cao, Hao Peng, Jia Wu, Yingdong Dou, Jianxin Li, and Philip S. Yu. 2021. Knowledge-preserving incremental social event detection via heterogeneous gnns. In *Proceedings of the Web Conference 2021*. 3383–3395.
- [12] Yuwei Cao, Hao Peng, Zhengtao Yu, and Yu Philip S. 2024. Hierarchical and incremental structural entropy minimization for unsupervised social event detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 8255–8264.
- [13] Xueqi Cheng, Xiaohui Yan, Yanyan Lan, and Jiafeng Guo. 2014. Btm: Topic modeling over short texts. *IEEE Transactions on Knowledge and Data Engineering* 26, 12 (2014), 2928–2941.

- [14] Mário Cordeiro. 2012. Twitter event detection: combining wavelet analysis and topic inference summarization. In *Doctoral symposium on informatics engineering*, Vol. 1, 11–16.
- [15] Mário Cordeiro and João Gama. 2016. Online social networks event detection: a survey. *Solving Large Scale Learning Tasks. Challenges and Algorithms: Essays Dedicated to Katharina Morik on the Occasion of Her 60th Birthday* (2016), 1–41.
- [16] Wanqiu Cui, Junping Du, Dawei Wang, Feifei Kou, and Zhe Xue. 2021. MVGAN: Multi-view graph attention network for social event detection. *ACM Transactions on Intelligent Systems and Technology (TIST)* 12, 3 (2021), 1–24.
- [17] Kaize Ding, Yancheng Wang, Yingzhen Yang, and Huan Liu. 2023. Eliciting structural and semantic global knowledge in unsupervised graph contrastive learning. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*. 7378–7386.
- [18] Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, and Maosong Sun. 2022. OpenPrompt: An Open-source Framework for Prompt-learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, 105–113.
- [19] Martin Ester, Hans-Peter Kriegel, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, Vol. 96. AAAI Press, 226–231.
- [20] Pablo A Estévez, Michel Tesmer, Claudio A Perez, and Jacek M Zurada. 2009. Normalized mutual information feature selection. *IEEE Transactions on neural networks* 20, 2 (2009), 189–201.
- [21] Mateusz Fedoryszak, Brent Frederick, Vijay Rajaram, and Changtao Zhong. 2019. Real-time event detection on social data streams. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2774–2782.
- [22] Wei Feng, Chao Zhang, Wei Zhang, Jiawei Han, Jianyong Wang, Charu Aggarwal, and Jianbin Huang. 2015. STREAMCUBE: Hierarchical spatio-temporal hashtag clustering for event exploration over the Twitter stream. In *2015 IEEE 31st international conference on data engineering*. IEEE, 1561–1572.
- [23] Gabriel Pui Cheong Fung, Jeffrey Xu Yu, Philip S. Yu, and Hongjun Lu. 2005. Parameter free bursty events detection in text streams. In *Proceedings of the 31st international conference on Very large data bases*. ACM, 181–192.
- [24] Anuradha Goswami and Ajeay Kumar. 2016. A survey of event detection techniques in online social networks. *Social Network Analysis and Mining* 6 (2016), 1–25.
- [25] Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. PPT: Pre-trained Prompt Tuning for Few-shot Learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 8410–8423.
- [26] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. 2016. Cross modal distillation for supervision transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2827–2836.
- [27] Alexander Hermans, Lucas Beyer, and Bastian Leibe. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017), 1–15.
- [28] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. 1–9.
- [29] Linmei Hu, Bin Zhang, Lei Hou, and Juanzi Li. 2017. Adaptive online event detection in news streams. *Knowledge-Based Systems* 138 (2017), 105–112.
- [30] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 4171–4186.
- [31] Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. 2016. Structured Attention Networks. In *International Conference on Learning Representations*. 1–21.
- [32] Thomas N Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*. 1–12.
- [33] Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations*. 1–14.
- [34] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 3045–3059.
- [35] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 1–10.
- [36] Angsheng Li, Xianchen Yin, and Yicheng Pan. 2016. Three-dimensional gene map of cancer cell types: Structural entropy minimisation principle for defining tumour subtypes. *Scientific reports* 6, 1 (2016), 20412.
- [37] Junyi Li, Tianyi Tang, Jian-Yun Nie, Ji-Rong Wen, and Wayne Xin Zhao. 2022. Learning to Transfer Prompts for Text Generation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 3506–3518.

- [38] Qian Li, Jianxin Li, Jiawei Sheng, Shiyao Cui, Jia Wu, Yiming Hei, Hao Peng, Shu Guo, Lihong Wang, and Amin Beheshti. 2022. A survey on deep learning event extraction: Approaches and applications. *IEEE Transactions on Neural Networks and Learning Systems* (2022), 1–22.
- [39] Bang Liu, Fred X Han, Di Niu, Linglong Kong, Kunfeng Lai, and Yu Xu. 2020. Story forest: Extracting events and telling stories from breaking news. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 14, 3 (2020), 1–28.
- [40] Fanzhen Liu, Shan Xue, Jia Wu, Chuan Zhou, Wenbin Hu, Cecile Paris, Surya Nepal, Jian Yang, and Philip S. Yu. 2021. Deep learning for community detection: progress, challenges and opportunities. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 4981–4987.
- [41] Jian Liu, Yubo Chen, Kang Liu, and Jun Zhao. 2018. Event detection via gated multilingual attention mechanism. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32. 4865–4872.
- [42] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *Comput. Surveys* 55, 9 (2023), 1–35.
- [43] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019), 1–13.
- [44] Yaopeng Liu, Hao Peng, Jianxin Li, Yangqiu Song, and Xiong Li. 2020. Event detection and evolution in multi-lingual social streams. *Frontiers of Computer Science* 14 (2020), 1–15.
- [45] Yixin Liu, Yu Zheng, Daokun Zhang, Hongxu Chen, Hao Peng, and Shirui Pan. 2022. Towards unsupervised deep graph structure learning. In *Proceedings of the ACM Web Conference 2022*. 1392–1403.
- [46] Zhiwei Liu, Yang Yang, Zi Huang, Fumin Shen, Dongxiang Zhang, and Heng Tao Shen. 2017. Event early embedding: Predicting event volume dynamics at early stage. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 997–1000.
- [47] Daniel Loureiro, Kiamehr Rezaee, Talayeh Riahi, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-Collados. 2023. Tweet insights: a visualization platform to extract temporal insights from twitter. *arXiv preprint arXiv:2308.02142* (2023), 1–10.
- [48] James MacQueen. 1967. Some methods for classification and analysis of multivariate observations. 1, 14 (1967), 281–297.
- [49] Fabrizio Marozzo and Alessandro Bessi. 2018. Analyzing polarization of social media users and news sites during political campaigns. *Social Network Analysis and Mining* 8 (2018), 1–13.
- [50] Béatrice Mazoyer, Julia Cagé, Nicolas Hervé, and Céline Hudelot. 2020. A French Corpus for Event Detection on Twitter. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. 6220–6227.
- [51] Andrew J McMinn, Yashar Moshfeghi, and Joemon M Jose. 2013. Building a large-scale corpus for evaluating event detection on twitter. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 409–418.
- [52] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*. 1–12.
- [53] Muhammad Tasnim Mohiuddin, M Saiful Bari, and Shafiq Joty. 2020. LNMap: Departures from Isomorphic Assumption in Bilingual Lexicon Induction Through Non-Linear Mapping in Latent Space. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2712–2723.
- [54] Keval Morabia, Neti Lalita Bhanu Murthy, Aruna Malapati, and Surender Samant. 2019. SEDTWik: segmentation-based event detection from tweets using Wikipedia. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*. Association for Computational Linguistics, 77–85.
- [55] Tahir M Nisar and Man Yeung. 2018. Twitter as a tool for forecasting stock market movements: A short-window event study. *The journal of finance and data science* 4, 2 (2018), 101–119.
- [56] Ozer Ozdakis, Pinar Karagoz, and Halit Oğuztüzün. 2017. Incremental clustering with vector expansion for online event detection in microblogs. *Social Network Analysis and Mining* 7 (2017), 1–17.
- [57] Hao Peng, Jianxin Li, Qiran Gong, Yangqiu Song, Yuanxing Ning, Kunfeng Lai, and Philip S. Yu. 2019. Fine-grained event categorization with heterogeneous graph convolutional networks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 3238–3245.
- [58] Hao Peng, Jianxin Li, Yu He, Yaopeng Liu, Mengjiao Bao, Lihong Wang, Yangqiu Song, and Qiang Yang. 2018. Large-scale hierarchical text classification with recursively regularized deep graph-cnn. In *Proceedings of the 2018 world wide web conference*. 1063–1072.
- [59] Hao Peng, Jianxin Li, Yangqiu Song, Renyu Yang, Rajiv Ranjan, Philip S. Yu, and Lifang He. 2021. Streaming social event detection and evolution discovery in heterogeneous information networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 5 (2021), 1–33.
- [60] Hao Peng, Ruitong Zhang, Shaoning Li, Yuwei Cao, Shirui Pan, and Philip S. Yu. 2022. Reinforced, incremental and cross-lingual event detection from social messages. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 1 (2022), 980–998.
- [61] Kun Peng, Lei Jiang, Hao Peng, Rui Liu, Zhengtao Yu, Jiaqian Ren, Zhifeng Hao, and Philip S. Yu. 2024. Prompt Based Tri-Channel Graph Convolution Neural Network for Aspect Sentiment Triplet Extraction. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*. 1–9.

- [62] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [63] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 2463–2473.
- [64] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018), 1–12.
- [65] Jiaqian Ren, Lei Jiang, Hao Peng, Yuwei Cao, Jia Wu, Philip S. Yu, and Lifang He. 2022. From known to unknown: quality-aware self-improving graph neural network for open set social event detection. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1696–1705.
- [66] Jiaqian Ren, Lei Jiang, Hao Peng, Zhiwei Liu, Jia Wu, and Philip S. Yu. 2022. Evidential temporal-aware graph-based social event detection via dempster-shafer theory. In *2022 IEEE International Conference on Web Services (ICWS)*. IEEE, 331–336.
- [67] Jiaqian Ren, Hao Peng, Lei Jiang, Zhifeng Hao, Jia Wu, Shengxiang Gao, Zhengtao Yu, and Qiang Yang. 2024. Towards Cross-lingual Social Event Detection with Hybrid Knowledge Distillation. *ACM Transactions on Knowledge Discovery from Data (TKDD)* (2024), 1–37.
- [68] Jiaqian Ren, Hao Peng, Lei Jiang, Zhiwei Liu, Jia Wu, Zhengtao Yu, and Philip S. Yu. 2023. Uncertainty-guided Boundary Learning for Imbalanced Social Event Detection. *IEEE Transactions on Knowledge and Data Engineering* (2023), 1–15.
- [69] Sihem Sahnoun, Samir Elloumi, and Sadok Ben Yahia. 2020. Event detection based on open information extraction and ontology. *Journal of Information and Telecommunication* 4, 3 (2020), 383–403.
- [70] Kari Sentz and Scott Person. 2002. Combination of evidence in Dempster-Shafer theory. (2002), 1–96.
- [71] Yiheng Shu, Zhiwei Yu, Yuhan Li, Börje Karlsson, Tingting Ma, Yuzhong Qu, and Chin-Yew Lin. 2022. TIARA: Multi-grained Retrieval for Robust Question Answering over Large Knowledge Base. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 8108–8121.
- [72] Xing Su, Shan Xue, Fanzhen Liu, Jia Wu, Jian Yang, Chuan Zhou, Wenbin Hu, Cecile Paris, Surya Nepal, Di Jin, et al. 2022. A comprehensive survey on community detection with deep learning. *IEEE Transactions on Neural Networks and Learning Systems* (2022), 1–29.
- [73] Bingyi Sun, Xiaohui Wei, Jiaxu Cui, and Yan Wu. 2023. Social activity matching with graph neural network in event-based social networks. *International Journal of Machine Learning and Cybernetics* 14, 6 (2023), 1989–2005.
- [74] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification?. In *Chinese computational linguistics: 18th China national conference, CCL 2019, Kunming, China, October 18–20, 2019, proceedings 18*. Springer, 194–206.
- [75] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30, 5998–6008.
- [76] Nguyen Xuan Vinh, Julien Epps, and James Bailey. 2009. Information theoretic measures for clusterings comparison: is a correction for chance necessary? (2009), 1073–1080.
- [77] Zhongqing Wang and Yue Zhang. 2017. A neural model for joint event detection and summarization. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 4158–4164.
- [78] Yifan Wei, Yisong Su, Huanhuan Ma, Xiaoyan Yu, Fangyu Lei, Yuanzhe Zhang, Jun Zhao, and Kang Liu. 2023. MenatQA: A New Dataset for Testing the Temporal Comprehension and Reasoning Abilities of Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 1434–1447.
- [79] Wei Xie, Feida Zhu, Jing Jiang, Ee-Peng Lim, and Ke Wang. 2016. Topicsketch: Real-time bursty topic detection from twitter. *IEEE Transactions on Knowledge and Data Engineering* 28, 8 (2016), 2216–2229.
- [80] Chen Xing, Yuan Wang, Jie Liu, Yalou Huang, and Wei-Ying Ma. 2016. Hashtag-based sub-event discovery using mutually generative lda in twitter. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 30. 2666–2672.
- [81] Xiaohui Yan, Jiafeng Guo, Yanyan Lan, Jun Xu, and Xueqi Cheng. 2015. A probabilistic model for bursty topic discovery in microblogs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29. AAAI Press, 353–359.
- [82] Junhan Yang, Zheng Liu, Shitao Xiao, Chaozhuo Li, Defu Lian, Sanjay Agrawal, Amit Singh, Guangzhong Sun, and Xing Xie. 2021. GraphFormers: GNN-nested Transformers for Representation Learning on Textual Graph. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 28798–28810.
- [83] Weiren Yu, Jianxin Li, Md Zakirul Alam Bhuiyan, Richong Zhang, and Jinpeng Huai. 2017. Ring: Real-time emerging anomaly monitoring system over text streams. *IEEE Transactions on Big Data* 5, 4 (2017), 506–519.
- [84] Chao Zhang, Liyuan Liu, Dongming Lei, Quan Yuan, Honglei Zhuang, Tim Hanratty, and Jiawei Han. 2017. Triovecevent: Embedding-based online local event detection in geo-tagged tweet streams. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 595–604.
- [85] Kuo Zhang, Juan Zi, and Li Gang Wu. 2007. New event detection based on indexing-tree and named entity. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. 215–222.

- [86] Ruitong Zhang, Hao Peng, Yingtong Dou, Jia Wu, Qingyun Sun, Yangyang Li, Jingyi Zhang, and Philip S Yu. 2022. Automating dbscan via deep reinforcement learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2620–2630.
- [87] Xinyang Zhang, Yury Malkov, Omar Florez, Serim Park, Brian McWilliams, Jiawei Han, and Ahmed El-Kishky. 2023. TwHIN-BERT: A socially-enriched pre-trained language model for multilingual tweet representations at twitter. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*. ACM, 5597–5607.
- [88] Zizhuo Zhang and Bang Wang. 2023. Prompt learning for news recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 227–237.
- [89] Wayne Xin Zhao, Jing Jiang, Jianshu Weng, Jing He, Ee-Peng Lim, Hongfei Yan, and Xiaoming Li. 2011. Comparing twitter and traditional media using topic models. In *Advances in Information Retrieval: 33rd European Conference on IR Research, ECIR 2011, Dublin, Ireland, April 18-21, 2011. Proceedings 33*. Springer, 338–349.
- [90] Xiangmin Zhou and Lei Chen. 2014. Event detection over twitter social media streams. *The VLDB journal* 23, 3 (2014), 381–400.

Received April 5, 2024; Revised June 29, 2024; Accepted September 7, 2024