Social Event Prediction via Temporal Knowledge Graph Learning

Wentao Shi
School of Computer Science
Nanjing University
Nanjing, China
wentao.shi@smail.nju.edu.cn

Jingjing Zhang
College of Computing, Data Science, and Society
UC Berkeley
Berkeley, USA
jingjing.z@berkeley.edu

Xiaobin Hong
School of Computer Science
Nanjing University
Nanjing, China
xiaobinhong@smail.nju.edu.cn

Mingkai Lin
School of Computer Science
Nanjing University
Nanjing, China
mingkai@nju.edu.cn

Sanglu Lu
School of Computer Science
Nanjing University
Nanjing, China
sanglu@nju.edu.cn

Wenzhong Li
School of Computer Science
Nanjing University
Nanjing, China
lwz@nju.edu.cn

***Abstract*—Social event prediction aims to infer future interactions among social actors from historical observations. It is a fundamental task for understanding the evolution of complex social systems, with potential applications in international relations analysis, conflict early warning, public security, and crisis monitoring. Unlike ordinary temporal forecasting, social event prediction is relational, temporal, and context-dependent: each event is shaped by the participating actors, relation semantics, historical dependencies, and latent social backgrounds. Temporal knowledge graphs provide an effective formulation by representing social events as time-stamped relational facts. However, existing methods mainly focus on structural and temporal patterns while overlooking latent contexts that shape event dynamics. In real-world event streams, the same relation may exhibit different temporal patterns under different political, military, or diplomatic contexts, while context labels are often coarse, noisy, and overlapping. To address these issues, we propose a context-adaptive temporal knowledge graph learning model for social event prediction. The framework refines hard context labels into soft context distributions, performs context-conditional temporal encoding, and injects query-relevant temporal evidence into prediction. Experiments on three GDELT-based [1] benchmarks show that our method improves prediction performance over strong baselines.**

I. Introduction

Social event prediction focuses on forecasting future interactions among social actors, such as cooperation, conflict, diplomatic actions, protests, or public security incidents. This task is important because social events are not isolated observations; they form evolving relational systems in which actors continuously influence one another over time. Accurately predicting future events can help reveal emerging social trends, identify potential risks, and support decision-making in domains such as international relations, regional security, and crisis response.

We formulate social event prediction as a temporal knowledge graph reasoning problem. Let $\mathcal{G}_{t-L:t-1} = \mathcal{G}_{t-L}, \ldots, \mathcal{G}_{t-2}, \mathcal{G}_{t-1}$, denote the historical event graphs between timestamp $t-L$ and $t-1$. Each graph $\mathcal{G}\tau$ contains social events observed at time τ. A standard temporal event can be represented as (h, r, o, τ), where h and o are the head and tail actors, r is the event relation, and τ is the timestamp. In this work, we further consider a context-aware formulation (h, r, o, τ, c), where c denotes the associated event context. Given a future query (h, r, t) and the historical graphs $\mathcal{G}_{t-L:t-1}$, the goal is to rank all candidate tail entities and predict the most likely actor o that will participate in the event.

As shown in Fig. 1, social event prediction involves three successive stages. Raw event sources, such as news reports, are first converted into structured event tuples, where actors and relations form the basic relational facts. Meanwhile, associated news texts provide contextual signals, from which topic distributions, statistical priors, and hard context labels can be constructed. These context-aware events are then organized into historical temporal graphs, which serve as the input for temporal graph learning. Given a future query, the model ranks candidate tail entities and outputs the most likely future event.

This task is challenging in two aspects. First, the same relation may indicate different future outcomes under different social backgrounds while contexts in real-world event datasets are usually coarse and ambiguous. Second, existing temporal knowledge graph methods often encode historical snapshots into temporal representations and then use a relatively static decoder for prediction. Such a design may overlook the fact that different queries depend on different historical evidence.

II. Method

To address these challenges, we propose a context-adaptive temporal knowledge graph learning model. The core idea is to model context as a flexible signal and make temporal evidence directly participate in query-level prediction.

First, we refine the original hard context label into a soft context distribution by combining the observed context label,

Corresponding Author: Wenzhong Li.

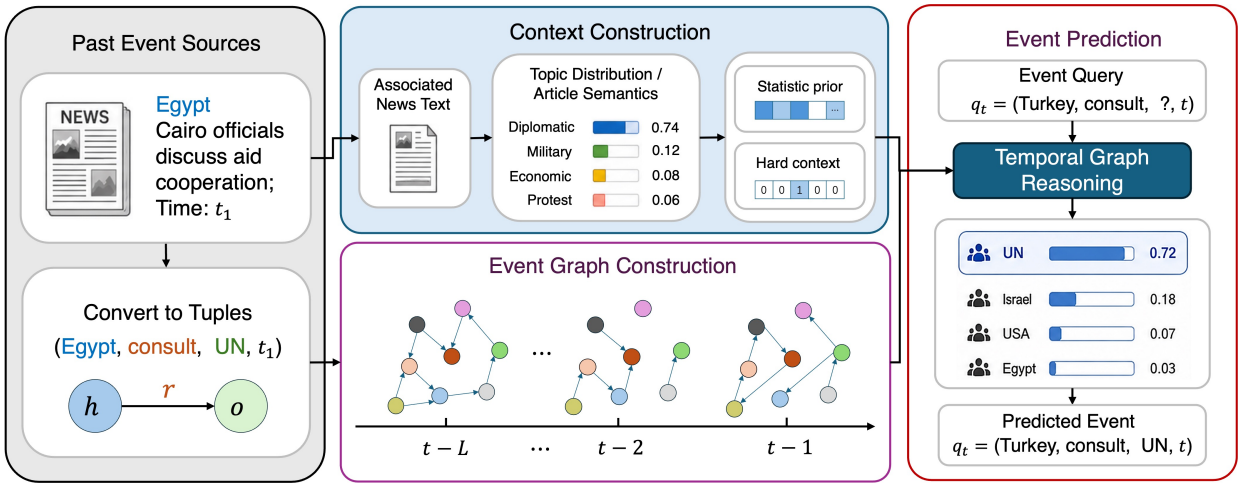

Fig. 1. General Process of Social Event Prediction with Context

statistical context priors, and learnable context mixtures. This design preserves useful context information while allowing each event to be associated with multiple latent contexts. The refined distribution is then mapped to a context embedding through trainable context prototypes.

Second, we perform context-conditional temporal encoding on the original event graph. Instead of splitting the graph into sparse context-specific subgraphs, relation representations are modulated by refined context embeddings during message passing. The context-conditioned messages are pooled into timestamp-level graph representations, which are then fed into sequence models to capture long-range temporal evolution and produce historical temporal memory.

Third, we conduct query-conditioned temporal retrieval. Given a query (h, r, t), the model constructs a query representation from the head entity, relation, and context embedding, and uses it to retrieve relevant evidence from temporal memory and entity memory. The combination of global temporal dynamics and actor-level information forms a query-aware temporal representation.

Finally, the retrieved temporal representation is injected into the decoder by modulating the head entity and relation representations. A ConvTransE-style decoder [2] is used to rank candidate tail entities. In this way, prediction is guided by both context-adaptive representations and query-relevant temporal evidence.

III. Experiments

As shown in Table I, our model consistently achieves the best MRR and Hits@3 on all three datasets based on GDELT, demonstrating the effectiveness of the proposed framework for social event prediction.

In addition, we observe from the experiments that context dynamics reveal meaningful social patterns across different regional event systems. Specifically, stable or centralized context structures usually correspond to more context-specialized interactions, while highly mobile multi-context structures indicate that actors and relations frequently shift across different social backgrounds. This suggests that context is not only useful for improving prediction performance, but also provides an interpretable perspective for understanding the evolution of regional social dynamics.

TABLE I
MRR and Hits@3 results on EG, IR, and IS.

Model	EG		IR		IS	
	MRR	H@3	MRR	H@3	MRR	H@3
RE-GCN [3]	0.1245	0.1366	0.1647	0.1796	0.1301	0.1281
SeCoGD [4]	0.1464	0.1605	0.1757	0.1902	0.1552	0.1588
Ours	**0.1553**	**0.1692**	**0.1862**	**0.2021**	**0.1616**	**0.1684**

IV. Conclusion

We propose a context-adaptive temporal knowledge graph learning model for social event prediction. The framework addresses two key characteristics of social event data: ambiguous latent contexts and query-dependent temporal evidence. By refining hard contexts into soft distributions, encoding context-conditioned temporal dynamics, and injecting retrieved temporal evidence into prediction, our method better captures the relational, temporal, and contextual nature of social events.

References

[1] K. Leetaru and P. A. Schrodt, "GDELT: Global data on events, location, and tone, 1979–2012," in ISA Annual Convention, vol. 2, 2013, pp. 1–49.
[2] C. Shang, Y. Tang, J. Huang, J. Bi, X. He, and B. Zhou, "End-to-end structure-aware convolutional networks for knowledge base completion," in Proceedings of the AAAI Conference on Artificial Intelligence. AAAI Press, 2019, pp. 3060–3067.
[3] Z. Li, X. Jin, W. Li, S. Guan, J. Guo, H. Shen, Y. Wang, and X. Cheng, "Temporal knowledge graph reasoning based on evolutional representation learning," in Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM, 2021, pp. 408–417.
[4] Y. Ma, C. Ye, Z. Wu, X. Wang, Y. Cao, and T.-S. Chua, "Context-aware event forecasting via graph disentanglement," in Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 1643–1652.