Interpreting Temporal Knowledge Graph Reasoning (Student Abstract)

Bin Chen[1], **Kai Yang**[1], **Wenxin Tai**[1,2], **Zhangtao Cheng**[1,2], **Leyuan Liu**[1], **Ting Zhong**[1,2], **Fan Zhou**[1,2,*]

[1]University of Electronic Science and Technology of China
[2]Kashi Institute of Electronics and Information Industry
binchen4110@gmail.com, kaiyang.cs@outlook.com, wxtai@std.uestc.edu.cn, zhangtao.cheng@outlook.com, leyuanliu@uestc.edu.cn, zhongting@uestc.edu.cn, fan.zhou@uestc.edu.cn

Abstract

Temporal knowledge graph reasoning is an essential task that holds immense value in diverse real-world applications. Existing studies mainly focus on leveraging structural and sequential dependencies, excelling in tasks like entity and link prediction. However, they confront a notable interpretability gap in their predictions, a pivotal facet for comprehending model behavior. In this study, we propose an innovative method, LSGAT, which not only exhibits remarkable precision in entity predictions but also enhances interpretability by identifying pivotal historical events influencing event predictions. LSGAT enables concise explanations for prediction outcomes, offering valuable insights into the otherwise enigmatic "black box" reasoning process. Through an exploration of the implications of the most influential events, it facilitates a deeper understanding of the underlying mechanisms governing predictions.

Introduction

Temporal Knowledge Graph Reasoning problem aims to predict missing facts on Temporal Knowledge Graph (TKGs). For example, given a specific query (s, r, ?, t) or (?, r, o, t), we want to predict the missing entity according to a sequence of TKGs $\mathcal{G}_{0:T}$. In this work, we mainly focus on the extrapolation scenario where t $>$ T. Although substantial research endeavors have been dedicated to advancing representation learning techniques for TKGs, it is imperative to acknowledge that there still exist several significant challenges that need to be solved: (1) *Structural dependencies*: how to comprehend and model the complex structural dependencies among concurrent facts, while concurrently accounting for the inherent heterogeneity within the knowledge graph; (2) *Sequential patterns*: how to effectively capture the evolving nature of specific events that reflect the behavioral trends and preferences of entities and relations; (3) *Interpretability*: despite the commendable performance achieved by existing methods in the task, they remain enigmatic "black box".

To solve challenges (1) and (2), existing methods like RE-GCN(Li et al. 2021) employ RNN to capture *sequential patterns* among adjacent timestamps and R-GCN for aggregating neighboring information, encompassing *structural dependencies*. However, they failed to account for the varying significance of different neighbors within their aggregation mechanism, consequently hindering their ability to accurately capture the intricacies of *structural dependencies*. Regarding challenge (3), there is a limited body of research addressing the aspect of *interpretability* within GNN-based TKG Reasoning.

To address the aforementioned three challenges, we present our solution, LSGAT, for TKG Reasoning. To effectively capture *structural dependencies* and *sequential patterns*, we propose an attention-based relation-aware GNN to learn evolving representations of entities. Furthermore, to enhance *interpretability*, inspired by (Ying et al. 2019), we devise an interpretable module to identify the most influential events during the prediction process.

Methodology

Our proposed LSGAT consists of two components: *Reasoning Module* and *Interpreting Module*, where the former seeks to reason for future events stemming from historical events while the latter provides reasonable interpretability for such predictions.

Reasoning Module. We aim to predict future events by extracting the expressive *structural dependencies* and *sequential patterns* from historical events. To capture the evolution pattern of events in the short term, we design a hierarchical relational graph attention network, called HRGAT, which considers both nodes and edge, as our semantic encoder to obtain the embedding of each entity h_t at timestamp t. For a KG at timestamp t, an object entity at layer l could get information from its neighbor entities and corresponding relations, and then update its representations at the next $l+1$ layer, i.e.,

$$\alpha^l_{s,o} = \frac{\exp\left(\mathbf{a}^{\mathrm{T}}\mathrm{g}\left(\mathrm{W}^l_1\left[\mathrm{h}^l_{s,t}, \mathrm{h}^l_{o,t}, \mathrm{r}^l_t\right]\right)\right)}{\sum\limits_{(s^*,r^*,o)\in\mathcal{F}_t}\exp\left(\mathbf{a}^{\mathrm{T}}\mathrm{g}\left(\mathrm{W}^l_1\left[\mathrm{h}^l_{s^*,t}, \mathrm{h}^l_{o,t}, \mathrm{r^*}^l_t\right]\right)\right)},$$

$$\mathrm{h}^{l+1}_{o,t} = \sum_{(s,r,o)\in\mathcal{F}_t}\alpha^l_{s,o}(\mathrm{h}^l_{s,t} + \mathrm{r}^l_t), \tag{1}$$

where $\mathrm{h}^l_{s,t}$, $\mathrm{h}^l_{o,t}$, r^l_t denote the l^{th} layer embedding of entities s, o and relation r at t; $\mathrm{h}^l_{s,t} + \mathrm{r}^l_t$ implies the translational property between the subject entity and the corresponding

*Corresponding author

object entity via the relation r at t; $\alpha^l_{s,o}$ denotes the influence coefficient between entity s and o in relation r, $\mathcal{F}_t$ denotes all facts that happened at t; $\mathbf{W}^l_1$ and $\mathbf{a}$ denote learnable weights, and $g(\cdot)$ is the LeakyReLU function. In addition, we use GRU to maintain the continuity of the development of events. For the long-term spatial and temporal dependencies of entities, we utilize R-GCN as a semantic aggregator to obtain the embeddings of all entities and relations based on the constructed global graph Φ_t, which contains M subgraphs:$\{\mathcal{G}_{t-M}, \mathcal{G}_{t-M+1}, \cdots, \mathcal{G}_{t-1}\}$. After that, we utilize ConvTransE as decoder to calculate the probability of interaction between subject s and object o under the relation r at time t. Formally,

$$p(o|s, r, \mathrm{H}_t, \mathrm{R}_t) = \sigma(\mathrm{H}_t \mathrm{ConvTransE}(\mathrm{h}_{s,t}, \mathrm{r}_t)), \quad (2)$$

where $\sigma(\cdot)$ is the Softmax function; $\mathrm{H}_t, \mathrm{R}_t$ denote the obtained embedding matrix for entities and relations at timestamp t; $\mathrm{h}_{s,t}, \mathrm{h}_{o,t}, \mathrm{r}_t$ are the embeddings of s, o, r in H_t and R_t, respectively. After obtaining the probability $\boldsymbol{p}^{short}$, $\boldsymbol{p}^{long}$ from short term and long term, the final probability is obtained by weighted summation:

$$\boldsymbol{p}^{final} = \alpha \times \boldsymbol{p}^{long} + (1 - \alpha) \times \boldsymbol{p}^{short}. \quad (3)$$

Interpreting Module. We choose to provide explanations for the predicted events by identifying pertinent influential historical events. We observe that a node o's computation graph tells the GNN how to generate o's embedding h_o, which determines the final prediction. Intuitively, our goal is to elucidate the model's prediction by deriving $\mathcal{G}_S$, where $\mathcal{G}_S$ is a small subgraph of the computation graph. Then, in such $\mathcal{G}_S$, we can identify the most influential historical events supporting the predicted future events. In practice, a mask of adjacency matrix $M \in \mathbb{R}^{n \times n}$ is learned to obtain $\mathcal{G}_S$. Formally, to generate an explanation in terms of $\mathcal{G}_S$, our optimization target is:

$$\min_M - \sum_{c=1}^{C} \mathbb{1}\left[y = c\right] \log P(Y = y | \mathcal{G} = A_c \odot \psi(M)), \quad (4)$$

where $\mathbb{1}$ is an indicator function, A_c denotes the adjacent matrix of computation graph and $\psi(\cdot)$ is the Sigmoid function.

Experiments

To Evaluate the effectiveness of our proposed method, we conduct experiments on two typical TKG datasets: ICEWS18, and ICEWS05-15. We compare our method with the following baselines: RE-GCN(Li et al. 2021), CyGNet(Zhu et al. 2021), CENET(Xu et al. 2023), and HGLS(Zhang et al. 2023). Besides, we also compare LSGAT with different variants: LSGAT-L, which only considers the long-term evolution dependency; LSGAT-S, which only considers the short-term dependency. We evaluate the LSGAT and baselines using three widely employed metrics in TKG Reasoning: MRR and Hits@$\{1, 10\}$ under the raw setting.

Main results. Table 1 shows the experimental results of all models on these two datasets and we have the following remarks. LSGAT outperforms all baselines in non-trivial margins. GNN-based models achieve better performances than copy-mechanism based models such as CyGNet and CENET. Moreover, our LSGAT outperforms all other GNN based methods which present the superiority of the proposed LSGAT. Furthermore, it is evident that LSGAT outperforms both LSGAT-L and LSGAT-S across all evaluation metrics. This observation validates that the amalgamation of long-term and short-term dependencies effectively and adequately captures the evolving characteristics of events.

Method	**ICEWS18**			**ICEWS05-15**		
	MRR	H@1	H@10	MRR	H@1	H@10
RE-GCN*	29.11	19.10	48.90	45.55	34.34	66.57
CyGNet	26.46	16.45	46.43	39.18	27.92	60.52
CENET	26.45	17.57	44.25	39.00	28.71	58.82
HGLS	<u>29.27</u>	<u>19.20</u>	<u>49.72</u>	<u>46.19</u>	<u>35.21</u>	<u>67.15</u>
LSGAT-L	27.01	17.45	45.83	39.44	28.45	61.28
LSGAT-S	29.24	19.10	49.35	45.34	34.05	66.62
LSGAT	**30.06**	**19.86**	**50.24**	**46.92**	**35.71**	**68.04**

Table 1: Performance Comparison. Noted that, * indicates that we remove the static information from the model to ensure the fairness of comparisons between all baselines.

influential events	weights
(Russia, Accuse, Israel, 9.18)	0.3
(Putin, Make statement, Israel, 9.18)	0.35
(Konashenkov, Make statement, Israel, 9.17)	0.1
(Israel, attack, Syria, 9.12)	0.1

Table 2: Case study: the most influential events and their corresponding weights for the predicted future event: (Government(Russia), Accuse, Israel, 9.19) in ICEWS18.

Case study. Table 2 demonstrated the interpretability of our model. LSGAT can provide the most influential historical events and their corresponding weights for the prediction, e.g., (Government(Russia), Accuse, Israel, 9.19).

Acknowledgements

This work was supported in part by National Natural Science Foundation of China (Grant No.62176043 and No.62072077).

References

Li, Z.; Jin, X.; Li, W.; Guan, S.; Guo, J.; Shen, H.; Wang, Y.; and Cheng, X. 2021. Temporal knowledge graph reasoning based on evolutional representation learning. In *SIGIR*, 408–417.

Xu, Y.; Ou, J.; Xu, H.; and Fu, L. 2023. Temporal knowledge graph reasoning with historical contrastive learning. In *AAAI*, volume 37, 4765–4773.

Ying, Z.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. Gnnexplainer: Generating explanations for graph neural networks. *NeurIPS*, 32.

Zhang, M.; Xia, Y.; Liu, Q.; Wu, S.; and Wang, L. 2023.

Learning Long-and Short-term Representations for Temporal Knowledge Graph Reasoning. In *TheWebConf*, 2412–2422.

Zhu, C.; Chen, M.; Fan, C.; Cheng, G.; and Zhang, Y. 2021. Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks. In *AAAI*, volume 35, 4732–4740.