Multi-Granularity History and Entity Similarity Learning for Temporal Knowledge Graph Reasoning

Anonymous ACL submission

Abstract

Temporal Knowledge Graph (TKG) reasoning, aiming to predict future unknown facts based on historical information, has attracted considerable attention due to its great practical value. Insight into history is the key to predict the future. However, most existing TKG reasoning models singly capture repetitive history, ignoring the entity's multi-hop neighbour history which can provide valuable background knowledge for TKG reasoning. In this paper, we propose **M**ulti-**G**ranularity History and **E**ntity **S**imilarity **L**earning (MGESL) model for Temporal Knowledge Graph Reasoning, which models historical information from both coarse-grained and fine-grained history. Since similar entities tend to exhibit similar behavioural patterns, we also design a hypergraph convolution aggregator to capture the similarity between entities. Furthermore, we introduce a more realistic setting for the TKG reasoning, where candidate entities are already known at the timestamp to be predicted. Extensive experiments on three benchmark datasets demonstrate the effectiveness of our proposed model.

1 Introduction

Temporal Knowledge Graphs (TKGs), served as a way to represent and store dynamic knowledge, have shown great value in many applications, such as event prediction (Deng et al., 2020), question answering (Mavromatis et al., 2022) and recommendation (Liu et al., 2023b). In TKGs, each fact is represented as a quadruple, e.g., (Obama, sanction, Russia, 2016-12-29) in Figure 1(a).

Reasoning over TKGs can be performed under two primary settings, i.e., interpolation and extrapolation (Jin et al., 2020). Given a TKG with timestamps from t_0 to t_n, interpolation mainly aims at inferring missing facts that occur at time t ($t_0 \leq t \leq t_n$), while extrapolation attempts to

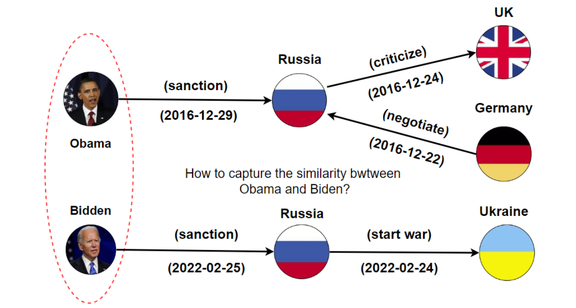

(a) An example of similarity learning problem

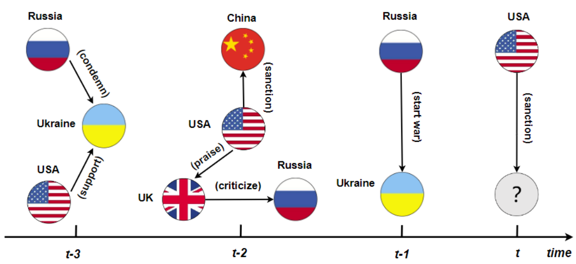

(b) An example of history of different granularities of entity

Figure 1: Illustration of the two problems of TKG reasoning task.

predict facts that occur at time t ($t > t_n$). In this paper, we mainly focus on TKG extrapolation. Most of existing extrapolation models (Jin et al., 2020; Li et al., 2021b, 2022b; Liu et al., 2023a) assume the candidate entities are unknown during the reasoning. However, there are cases that we already know the candidate entities, e.g., suspects are often identified beforehand in criminal investigations and candidates are usually already determined before the presidential election. In these cases, those extrapolation models (Jin et al., 2020; Li et al., 2021b, 2022b; Liu et al., 2023a) cannot effectively utilize the information of those candidate entities because they treat all entities equally during the reasoning. Therefore, we introduce a new setting called the candidate entity known setting, where all the entities at t are known in advance. In contrast, if the candidate entities at t are unknown during the reasoning, we call this the candidate entity unknown setting. In this paper, both candidate entity known and unknown settings will be discussed.

To predict what will happen in the future, we found that (1) searching for similar entities, observing and understanding the evolutionary pattern of the actions of similar entities, and (2) delving into the entity historical context from multi-granularity are crucial. Figure 1(a) shows an example of TKG similarity learning problem, where Obama and Biden both sanction Russia. However, since Obama and Biden are not connected in this example, vanilla graph convolution is unable to capture the interaction between them. To address this issue, we realize that both Obama and Biden share the same relation of sanction. Since hypergraph convolution can enable information interaction among entities under the same relation, we therefore design a hypergraph convolutional aggregator to capture similarity information between them. Additionally, existing models (Jin et al., 2020; Li et al., 2021b) mainly focus on utilising the available temporal and structural information in the TKG for inference, ignoring the history information. Even though some recent studies (Zhu et al., 2021; Li et al., 2022a; Xu et al., 2023) tried to find the correct answer from long-term global repeated history (i.e., fine-grained history), but they ignore the more generalised history. For instance, Figure 1(b) illustrates a temporal knowledge graph with several timestamps, where the task is to predict the answer to the query (USA, sanction, ?, t). Most models (Zhu et al., 2021; Xu et al., 2023) prioritize repeated history, and return China as the answer. However the correct answer to the question is Russia which is a multi-hop neighbour of USA. To overcome this limitations, we further consider multi-hop neighbour entities (i.e., coarse-grained history) in TKG reasoning.

To this end, we consider history at two levels of granularity (i.e., fine and coarse-grained history) and entity similarity learning simultaneously, and propose the **M**ulti-**G**ranularity History and **E**ntity **S**imilarity **L**earning (MGESL) model for Temporal Knowledge Graph Reasoning. Specifically, MGESL consists of three modules, i.e., (1) Entity Similarity Learning Module, which is used to capture the similarity between entities that share the same relation; (2) Temporal Evolution Module, which is used to aggregate and transfer the KG information from spatial and temporal views, respectively; (3) Multi-Granularity History Module, which is used to capture history from both coarse and fine granularities. Our main contributions are summarized as follows:

- We propose a TKG reasoning model MGESL, which can simultaneously consider entity similarity learning, coarse-grained and fine-grained history. To the best of our knowledge, we are the first to consider these features together.
- We design a novel hypergraph convolutional aggregator to capture similarities between entities, and utilize the coarse-grained history to capture multi-hop historical contextual information and fine-grained history to guide decoder for decoding to make full use of historical information.
- Besides the candidate entity unknown setting, we also propose another realistic TKG reasoning setting, i.e., the candidate entities are already known. Extensive experiments on three benchmark datasets show that our proposed MGESL model outperforms existing TKG reasoning methods under both settings.

2 Related Work

Since TKG interpolation is not the research object of this paper, we mainly review the existing TKG reasoning models under the extrapolation setting. Many extrapolation models utilise the available temporal and structural information in TKG for inference. RE-Net (Jin et al., 2020) utilizes heterogeneous graph convolution (RGCN) (Schlichtkrull et al., 2018) to capture the structural information within the same timestamp and employs a recurrent neural network (RNN) to model the temporal information between different timestamps. RE-GCN (Li et al., 2021b) further constrains the evolution of entities by incorporating additional static attributes. However, they do not consider the history information. CyGNet (Zhu et al., 2021) and CENET (Xu et al., 2023) propose a copy mechanism to find the correct answer among long-term global history, i.e., the fine-grained history. TiRGN (Li et al., 2022a) considers the sequential, repetitive and cyclical patterns of historical facts. However, they ignore the multi-hop neighbour history, i.e., the coarse-grained history. xERTE (Han et al., 2021) employs a subgraph sampling technique to construct interpretable reasoning graphs. CluSTeR

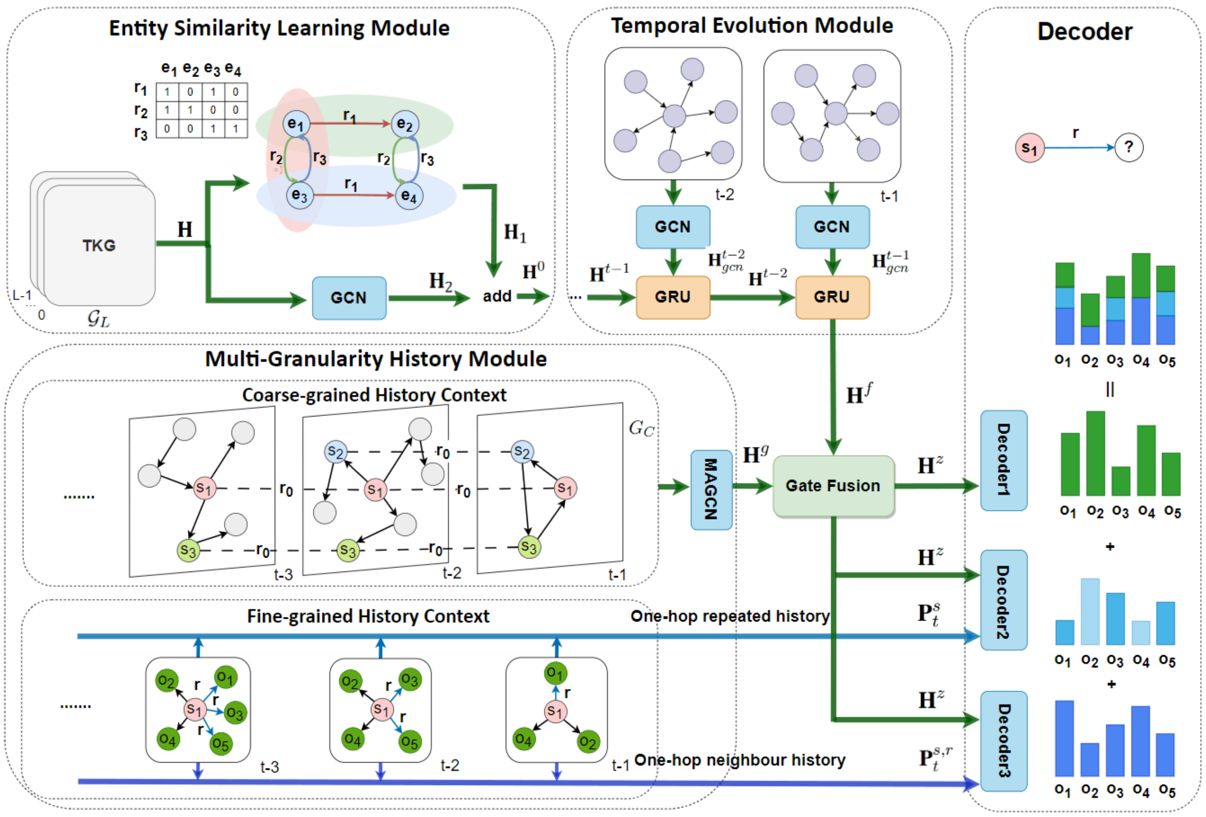

Figure 2: Illustration of the proposed MGESL model. Entity Similarity Learning Module captures the similarities between entities that share the same relation. Temporal Evolution Module aggregates and transfers the KG information from spatial and temporal views, respectively. Multi-Granularity History Module models history from both coarse and fine granularity.

(Li et al., 2021a) and TITer (Sun et al., 2021) both utilize reinforcement learning to search for a series of historical facts for reasoning. HGLS (Zhang et al., 2023) captures the long and short history of an entity by constructing global graphs. However, all the above models do not consider the importance of entity similarity learning in TKG reasoning.

3 Preliminaries

A temporal knowledge graph can be defined as $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, ..., \mathcal{G}_T\}$, and T is the number of timestamps. The subgraph $\mathcal{G}_t = (\mathcal{E}, \mathcal{R}, \mathcal{F}_t)$ at t is a directed multi-relational graph, where $\mathcal{E}$ is the set of entities, $\mathcal{R}$ is the set of relations, and $\mathcal{F}_t$ is the set of facts at t. A fact in $\mathcal{F}_t$ can be formalized as a quadruple (s, r, o, t), where $s, o \in \mathcal{E}$ and $r \in \mathcal{R}$. It describes that a fact of relation type r occurs between subject entity s and object entity o at time t.

The extrapolation reasoning task aims to predict the missing object entity o or subject s via answering query like $(s, r, ?, t_q)$ or $(?, r, o, t_q)$ based on the historical facts $\{(s, r, o, t_i)|t_i < t_q\}$. For each quadruple (s, r, o, t), an inverse relation quadruple (o, r^{-1}, s, t) is often added to the dataset (Vashishth et al., 2020). Therefore, when predicting the missing subject of a query $(?, r, o, t_q)$, we can convert it into predicting $(o, r^{-1}, ?, t_q)$. Based on this, the model in this paper only aims to predict the missing object entity. We use **bold** items to denote vector embeddings. For example, $\mathbf{H} \in \mathbb{R}^{|\mathcal{E}|\times d}$ and $\mathbf{R} \in \mathbb{R}^{|2\mathcal{R}|\times d}$ are used to represent the randomly initial embedding of entities and relations respectively, where d denotes the embedding dimension.

4 Methodology

4.1 Model Overview

The framework of MGESL is shown in Figure 2, comprising three modules: (1) the Entity Similarity Learning Module, (2) the Temporal Evolution Module, and (3) the Multi-Granularity History Module. First, the Entity Similarity Learning Module learns the representation of entity with similarity information. Next, the learned entity representation is fed to the Temporal Evolution Module, where it further learns about the structural and sequential characteristics of recent facts. Then, it combines with historical context information learnt from the coarse-grained history in the Multi-Granularity History Module. Finally, the entity representation is decoded under the guidance of the fine-grained history.

4.2 Entity Similarity Learning

4.2.1 Pre-Learning Graph

Inspired by the pre-training model (Devlin et al., 2019), we first construct a pre-learning graph and initially learn the representation of entities on the pre-learning graph. Entity similarity information is also learnt on this graph. For a TKG $\mathcal{G}$, we ignore the time factor to merge the subgraphs of the first L timestamps to form a pre-learning graph $\mathcal{G}_L$, i.e., $\mathcal{G}_L = (\mathcal{E}, \mathcal{R}, \mathcal{F}_L)$, where $\mathcal{F}_L = \{(s,r,o) \,|\, (s,r,o,t) \in \mathcal{F}_t, 0 < t < L\}$ is a set of facts.

4.2.2 Hypergraph Convolution

To effectively capture the similarity between entities in the pre-learning graph, we design a hypergraph convolutional network. First, we construct a hypergraph neighbourhood matrix $D \in \mathbb{R}^{|\mathcal{E}| \times 2|\mathcal{R}|}$, where $D_{i,j} = 1$ means the i^{th} entity is the subject entity of the j^{th} relation, otherwise it equals 0. As stated in Section 3, for each relation, we only aggregate messages from its subject entity through employing an inverse relation.

First, messages from the subject entity are passed into the relation:

$$\mathbf{X} = \frac{1}{2}\mathbf{W}_1 D^{-1}\mathbf{H} + \frac{1}{2}\mathbf{W}_2\mathbf{R} \tag{1}$$

where $\mathbf{W}_1$, $\mathbf{W}_2$ are the learnable weights. The result $\mathbf{X} \in \mathbb{R}^{|2\mathcal{R}| \times d}$ contains messages from the subject entities and the relation itself. Next, the relation message is passed into the subject entity:

$$\mathbf{H}_1 = \sigma(\frac{1}{2}\mathbf{W}_3 D\mathbf{X} + \frac{1}{2}\mathbf{W}_4\mathbf{H}) \tag{2}$$

where $\mathbf{W}_3$, $\mathbf{W}_4$ are the learnable weights and σ is the ReLu activation function. Through the above steps, we can initially learn the representation of entities $\mathbf{H}_1$, which incorporates the similarity information between entities.

4.2.3 Structural Encoder

Hypergraph convolution on the pre-learning graph mainly captures the similarity information between entities, but it cannot capture the inherent graph structure information of the pre-learning graph. Therefore, we utilize a heterogeneous graph convolution network (Vashishth et al., 2020) as a structural encoder to aggregate information from multiple relations and multi-hop neighbour entities on the pre-learning graph, which is defined as follows:

$$\mathbf{h}_s^{l+1} = \sigma\left(\sum\nolimits_{(s,r,o)\in\mathcal{F}_L} \frac{1}{c_s}\mathbf{W}_0^l(\mathbf{h}_o^l + \mathbf{r}) + \mathbf{W}_1^l\mathbf{h}_s^l\right) \tag{3}$$

where $\mathbf{h}_s^l$, $\mathbf{h}_o^l$ denote the l^{th} layer embeddings of entities s, o respectively, $\mathbf{r}$ denotes the embedding of relation r, c_s is a normalizing factor equal to the number of neighbours of s, $\mathbf{W}_0^l$ and $\mathbf{W}_1^l$ denote the learnable weights of the l^{th} layer, and σ is the ReLu activation function. We denote the entity embedding of the output of the last layer as $\mathbf{H}_2$.

Given that the meaning of relation r remains consistent over time, we do not update relation embedding in this paper to maintain its semantic stability. Finally, we combine $\mathbf{H}_1$ and $\mathbf{H}_2$ to get the entity representation $\mathbf{H}^0$,

$$\mathbf{H}^0 = \alpha\mathbf{H}_1 + (1-\alpha)\mathbf{H}_2 \tag{4}$$

$\mathbf{H}^0$ denotes entity embedding obtained by learning on the pre-learning graph, incorporating similarity and structural information between entities.

4.3 Temporal Evolution

Future facts are usually closely related to recent facts, and our temporal evolution module aims to model recent facts. KGs naturally have graph structure information, while TKGs have the additional dimension of time compared to KGs. Therefore, we aggregate and transfer the TKG information from both spatial and temporal views. To capture the structural information between entities, we also utilize the heterogeneous graph convolutional network in Equation (3) for each timestamp,

$$\mathbf{H}_{gcn}^t = \mathrm{GCN}(\mathbf{H}^t, \mathbf{R}) \tag{5}$$

where $\mathbf{H}^t$ denotes the entity embedding at time t, whose initial value is the output of the similarity learning module $\mathbf{H}^0$. $\mathbf{H}_{gcn}^t$ denotes the entity embedding after aggregation by GCN Encoder. In order to include the sequential dependencies of subgraphs at the previous timestamps, we utilize GRU to update the representations of entities,

$$\mathbf{H}^{t+1} = \mathrm{GRU}(\mathbf{H}_{gcn}^t, \mathbf{H}^t). \tag{6}$$

We denote the output of the last timestamp as $\mathbf{H}^f$.

4.4 Multi-Granularity History Learning

4.4.1 Background Graph

In order to more accurately model the representation of entities and the connections between them, we construct a background graph G_C based on the most recent C timestamps, similar to HGLS (Zhang et al., 2023). Specifically, when the candidate entities are known, the steps to construct the background graph are as follows: (1) identify the position where each candidate entity appears in the recent C timestamps. (2) conduct breadth-first search from each candidate entity to extract their n-hop neighbours. (3) merge the common neighbours of candidate entities and add temporal edge $\mathbf{r}_0$ (a randomly initial vector) between identical entities across different timestamps. With the steps above, we have established a background graph for more accurate entity representation learning. When the candidate entities are unknown, we take all entities in TKG as candidates and then execute the above three steps to construct the background graph.

4.4.2 Muti-head Attention GCN (MAGCN)

We employ a heterogeneous graph convolution network that incorporates the muti-head attention mechanism to effectively capture entity representation in the background graphs. First, all entities in the background graph are initialised by $\mathbf{H}$ for their initial embedding. Next, we combine the embeddings of the subject entity, the relation, and the object entity to calculate their attention scores,

$$\beta_{s,r,o} = LeakyRelu(\mathbf{W}_5[\mathbf{h}^s \oplus \mathbf{r} \oplus \mathbf{h}^o]) \quad (7)$$

where $\mathbf{h}^s$, $\mathbf{h}^o$ and $\mathbf{r}$ denote the embeddings of entities s, o and relation r, respectively, $\mathbf{W}_5$ denotes learnable weight, and $\oplus$ is the concatenation operation. After that, we further calculate their coefficients based on the scores of each triple,

$$\alpha_{s,r,o} = \frac{exp(\beta_{s,r,o})}{\sum_{(s,r^i,o^i)\in N_s} exp(\beta_{s,r^i,o^i})} \quad (8)$$

where N_s denotes the set of all triples with s as subject entity. After that, we can attentively aggregate message from all neighbours of entity s in the background graph. The utilization of the multi-head attention mechanism can enhance the stability of the convolution. Formally, the aggregator is defned as follows:

$$\mathbf{h}_s^{l+1} = \|_{m=1}^{M} \sigma \left(\sum_{(s,r,o)\in N_s} \alpha_{s,r,o} \mathbf{W}_6^l (\mathbf{h}_o^l + \mathbf{r}) + \mathbf{W}_7^l \mathbf{h}_s^l \right) \quad (9)$$

where M denotes the number of attention heads, $\|$ represents concatenation, $\mathbf{h}_s^l$ and $\mathbf{h}_s^o$ denote the embedding of entity s and o after the l^{th} layer aggregation, $\mathbf{r}$ denotes the embedding of relation r, $\mathbf{W}_6^l$ and $\mathbf{W}_7^l$ are learnable weights, and σ is the ReLu activation function. We denote the entity embedding of the last layer as $\mathbf{H}^g$.

Finally, we use a gate mechanism to fuse the entity embedding learnt from the temporal evolution module with the entity embedding learnt from the background graph,

$$\mathbf{H}^z = \sigma(\mathbf{U}) \odot \mathbf{H}^f + (1 - \sigma(\mathbf{U})) \odot \mathbf{H}^g \quad (10)$$

where $\mathbf{U} \in \mathbb{R}^{|\mathcal{E}|\times d}$ denotes the gate vector, $\odot$ denotes element-wise dot option, σ denotes $sigmoid$ function to map values to the range of 0 to 1. Finally, we obtain a representation of the entity $\mathbf{H}^z$, which incorporates the similarity information between entities, the entity's recent temporal information and contextual information.

4.4.3 Fine-grained History

Based on human experience in predicting future facts, the answer to a query is often an entity that is closely related to the current entity. Therefore, we extract two types of fine-grained histories, i.e., one-hop history neighbours and repeated history answers (Li et al., 2022a). Specifically, for a query $(s, r, ?, t)$ the set of one-hop history neighbours for the entity s at t can be defined as follows:

$$\mathbf{P}_t^s = \mathbf{p}_0^s \cup \mathbf{p}_1^s \cup \mathbf{p}_2^s \cup ... \cup \mathbf{p}_{t-1}^s \quad (11)$$

where $\mathbf{p}_t^s$ denotes a vector where each element represents an entity. If the corresponding element of an entity is 1, it means that the entity is a one-hop neighbour of s at t, otherwise it is 0. The symbol $\cup$ means union option. $\mathbf{P}_t^s$ is the union vector of all vectors before t. Similarly, we can calculate the repeated history answer vector $\mathbf{P}_t^{s,r}$.

4.5 Fine-grained History Guided Decoder

4.5.1 Scoring Function

We utilize ConvTransE (Shang et al., 2019) as decoder to fuse the semantic information of s and r in query $(s, r, ?, t)$. Since $\mathbf{H}^z$ already incorporates information of the coarse-grained history, the scores caculated based on coarse-grained history can be defined as follows:

$$\mathbf{p}^{coarse} = softmax(\mathrm{ConvTransE}(\mathbf{h}_t^s, \mathbf{r})\mathbf{H}^z) \quad (12)$$

where $\mathbf{h}_t^s$ and $\mathbf{r}$ denote the embedding of subject entity s and relation r, respectively. For the fine-grained history (i.e., one-hop neighbour history and repeated history), we use these two vectors ($\mathbf{P}_t^s$ and $\mathbf{P}_t^{s,r}$) generated in section 4.4.3 to guide the decoder in scoring, i.e.,

$$\mathbf{p}^{local} = softmax(\text{ConvTransE}(\mathbf{h}_t^s, \mathbf{r})\mathbf{H}^z\mathbf{P}_t^s) \tag{13}$$

$$\mathbf{p}^{history} = softmax(\text{ConvTransE}(\mathbf{h}_t^s, \mathbf{r})\mathbf{H}^z\mathbf{P}_t^{s,r}) \tag{14}$$

where $\mathbf{p}^{local}$ and $\mathbf{p}^{history}$ denote the scores guided by one-hop neighbour history and repeated history respectively. The final score is calculated as follows:

$$\mathbf{p} = \mu_1\mathbf{p}^{coarse} + \mu_2\mathbf{p}^{local} + \mu_3\mathbf{p}^{history} \tag{15}$$

where μ_1, μ_2, μ_3 are hyperparameters and $\mu_1+\mu_2+\mu_3 = 1$.

4.5.2 Training Objective

Predicting the object entity based on a given query $(s, r, ?, t)$ can be viewed as a multi-class classification task (Jin et al., 2020), where each class corresponds to one entity. The learning objective is to minimize the following cross-entropy loss $\mathcal{L}$ during training:

$$\mathcal{L} = -\sum\nolimits_{(s,r,o,t)\in\mathcal{G}} \mathbf{y}_t^e \log \mathbf{p}(o \,|\, s, r, t) \tag{16}$$

where $\mathbf{p}(o \,|\, s, r, t)$ is the final probability score of entity, $\mathbf{y}_t^e \in \mathbb{R}^{|\mathcal{E}|}$ is the label vector, of which the element is 1 if the fact occurs, otherwise is 0.

5 Experiments

5.1 Setup

5.1.1 Datasets

We use three typical TKG datasets in our experiments: ICEWS14 (Riloff et al., 2018), ICEWS18 (Jin et al., 2020), and ICEWS05-15 (Riloff et al., 2018). We divide them into training, validation, and test sets with a proportion of 80%, 10%, and 10% by timestamps following RE-GCN (Li et al., 2021b, 2022a; Xu et al., 2023).

5.1.2 Baselines

Under the candidate entity unknown setting, we compare our proposed MGESL model with three kinds of baselines, i.e., (1) Static KG reasoning models, i.e., DistMult (Yang et al., 2015), ConvE (Dettmers et al., 2018), ComplEx (Trouillon et al., 2016), Conv-TransE (Shang et al., 2019), RotatE (Sun et al., 2019) and R-GCN (Schlichtkrull et al., 2018). (2) Interpolated TKG reasoning models, i.e., TTransE (Jiang et al., 2016), HyTE (Dasgupta et al., 2018), and TA-DistMult (Riloff et al., 2018). (3) Current state-of-the-art extrapolated TKG reasoning models, i.e., RE-NET (Jin et al., 2020), CyGNet (Zhu et al., 2021), xERTE (Han et al., 2021), RE-GCN (Li et al., 2021b), TITER (Sun et al., 2021), TLogic (Liu et al., 2022), CEN (Li et al., 2022b), TiRGN (Li et al., 2022a), CENET (Xu et al., 2023), RETIA (Liu et al., 2023a) and DaeMon (Dong et al., 2023).

Under the candidate entity known setting, we mainly focus on comparing to the extrapolated TKG reasoning models, including RE-GCN (Li et al., 2021b), TiRGN (Li et al., 2022a) and HGLS (Zhang et al., 2023). As all previous TKG extrapolation models were conducted under the candidate entity unknown setting, we intentionally revealed all candidate entities of the timestamp to be predicted. This means that these models only need to score and find the correct answer from the revealed candidate entities, not from all entities in the TKG.

5.1.3 Training Settings and Evaluation Metrics

We report a widely used time-aware filtered version (Sun et al., 2021; Li et al., 2022a,b) of Mean Reciprocal Ranks (MRR) and Hits@1/3/10. As to model configurations, we set the embedding dimension (d) to 200, L is 50, α is 0.2, C is 20 for the candidate unknown setting and 10 for the candidate known setting, n is 2, the layer of structural encoder and muti-head attention GCN are both 2. μ_1, μ_2 and μ_3 are 0.3, 0.5 and 0.2, respectively. Adam is used for parameter learning, and the learning rate is set to 0.001.

5.2 Results

Table 1 presents the MRR and Hits@1/3/10 results of entity prediction on three TKGs under the candidate entity unknown setting. Specifically, our proposed MGESL significantly outperforms all the static models (i.e., the first block in Table 1) because they ignore the time dimension of the facts in TKGs. MGESL also performs much better than the temporal models for the interpolation setting (i.e., the second block in Table 1) because MGESL additionally captures temporally sequential patterns by temporal evolution module. In comparison

Model	ICEWS14				ICEWS18				ICEWS05-15			
	MRR	Hit@1	Hit@3	Hit@10	MRR	Hit@1	Hit@3	Hit@10	MRR	Hit@1	Hit@3	Hit@10
DistMult	20.32	6.13	27.59	46.61	13.86	5.61	15.22	31.26	19.91	5.63	27.22	47.33
ConvE	30.30	21.30	34.42	47.89	22.81	13.63	25.83	41.43	31.40	21.56	35.70	50.96
ComplEx	22.61	9.88	28.93	47.57	15.45	8.04	17.19	30.73	20.26	6.66	26.43	47.31
Conv-TransE	31.50	22.46	34.98	50.03	23.22	14.26	26.13	41.34	30.28	20.79	33.80	49.95
RotatE	25.71	16.41	29.01	45.16	14.53	6.47	15.78	31.86	19.01	10.42	21.35	36.92
R-GCN	28.03	19.42	31.95	44.83	15.05	8.13	16.49	29.00	27.13	18.83	30.41	43.16
TTransE	12.86	3.14	15.72	33.65	8.44	1.85	8.95	22.38	16.53	5.51	20.77	39.26
HyTE	16.78	2.13	24.84	43.94	7.41	3.10	7.33	16.01	16.05	6.53	20.20	34.72
TA-DistMult	26.22	16.83	29.72	45.23	16.42	8.60	18.13	32.51	27.51	17.57	31.46	47.32
RE-NET	39.86	30.11	44.02	58.21	29.78	19.73	32.55	48.46	43.67	33.55	48.83	62.72
GyGNet	37.65	27.43	42.63	57.90	27.12	17.21	30.97	46.85	40.42	29.44	46.06	61.60
xERTE	40.79	32.70	45.67	57.30	29.31	21.03	33.51	46.48	46.62	37.84	52.31	63.92
RE-GCN	39.42	30.13	43.80	57.08	27.51	17.82	31.17	46.55	38.27	27.43	43.06	59.93
TITER	41.73	32.74	—	58.44	29.98	22.05	—	44.83	47.60	38.29	—	64.86
TLogic	40.90	32.10	45.50	57.60	30.00	22.10	33.50	44.80	47.70	38.00	52.90	65.80
CEN	42.20	32.08	47.46	61.31	31.50	21.70	35.44	50.59	45.27	34.18	—	66.46
TiRGN	41.52	32.04	46.20	59.62	31.70	21.82	35.90	51.15	48.52	37.55	53.54	68.74
CENET	41.30	32.58	—	58.22	29.65	19.98	—	48.23	47.13	37.25	—	67.61
DaeMon	—	—	—	—	31.85	22.67	35.92	49.80	—	—	—	—
RETIA	41.61	31.66	46.36	60.61	31.23	21.55	35.07	50.17	>20Days	>20Days	>20Days	>20Days
MGESL	**45.65**	**35.28**	**51.12**	**65.48**	**34.18**	**23.66**	**38.64**	**54.89**	**53.78**	**42.52**	**60.40**	**75.04**

Table 1: Performance on three datasets in terms of MRR (%), Hit@1 (%), Hit@3 (%) and Hit@10 (%) under the candidate entity unknown setting. The best is highlighted in boldface, and the second is underlined.

Model	ICEWS14				ICEWS18				ICEWS05-15			
	MRR	Hit@1	Hit@3	Hit@10	MRR	Hit@1	Hit@3	Hit@10	MRR	Hit@1	Hit@3	Hit@10
RE-GCN	46.19	34.97	51.79	67.97	33.90	23.20	38.06	55.11	54.98	43.50	61.52	76.49
TIRGN	47.46	36.50	52.68	68.65	34.88	23.96	39.33	56.48	55.87	44.44	62.31	77.45
HGLS	47.00	35.06	—	70.41	29.32	19.21	—	49.83	46.21	35.32	—	67.12
MGESL	**51.86**	**40.49**	**58.26**	**73.41**	**37.57**	**26.10**	**42.63**	**60.16**	**58.06**	**46.84**	**64.47**	**79.63**

Table 2: Performance on three datasets in terms of MRR (%), Hit@1 (%), Hit@3 (%) and Hit@10 (%) under the candidate entity known setting. The best is highlighted in boldface, and the second is underlined.

to the current sate-of-the-art temporal models under the extrapolation setting (i.e., the third block in Table 1), our model also achieves notable improvements. Specifically, MGESL improves approximately 8.48%, 7.76%, 7.71%, and 8.79% on ICEWS14 for MRR, Hit@1, Hit@3, and Hit@10, respectively. This is because our model can effectively capture the similarity information between entities by hypergraph convolution and model the representation of entities more accurately from multiple granularities.

Table 2 shows that MGESL also significantly outperforms other TKG extrapolation models under the candidate known setting. Specifically, MGESL improves approximately 9.27%, 10.93%, 10.59%, and 4.26% on ICEWS14 for MRR, Hit@1, Hit@3, and Hit@10, respectively. These improvements arises from the background graph constructed by the candidate entities. This graph allows us to comprehensively understand and analyze the connections between these entities and effectively find the correct answer.

5.3 Ablation Study

The ablation studies are performed on ICEWS14 with all four evaluation metrics. Five submodels are compared, including (1) MGESL without similarity learning module (MGESL w/o SLM), (2) MGESL without temporal evolution module (MGESL w/o TEM), (3) MGESL without fine-grained history (MGESL w/o Fine), (4) MGESL without coarse-grained history (MGESL w/o Coarse), (5) the original MGESL model (MGESL).

Table 3 shows the ablation results under the candidate entity unknown setting. When the similarity learning module (SLM) and temporal evolution module (TEM) are removed, the performance of

Model	ICEWS14			
	MRR	Hit@1	Hit@3	Hit@10
MGESL w/o SLM	44.10	33.89	49.37	63.35
MGESL w/o TEM	43.71	33.14	49.17	64.32
MGESL w/o Fine	42.20	32.14	46.78	61.93
MGESL w/o Coarse	42.94	33.07	48.08	61.59
MGESL	**45.65**	**35.28**	**51.12**	**65.48**

Table 3: Ablation results under the candidate unknown setting. The best performance is highlighted in boldface.

Model	ICEWS14			
	MRR	Hit@1	Hit@3	Hit@10
MGESL w/o SLM	50.21	39.05	56.38	71.30
MGESL w/o TEM	49.69	38.57	55.91	70.58
MGESL w/o Fine	46.61	35.37	52.50	68.51
MGESL w/o Coarse	42.75	32.26	47.82	61.59
MGESL	**51.86**	**40.49**	**58.26**	**73.41**

Table 4: Ablation results under the candidate known setting. The best performance is highlighted in boldface.

the model decreased by 3.41% and 4.35% for MRR respectively, which indicates the effectiveness of these two module. We can notice that removing the fine-grained history module (Fine) degrades the performance of the model more severely compared to removing the coarse-grained history module (Coarse), which causes a 7.56% performance degradation for MRR compared with MGESL. This is because coarse-grained history may contain more noisy information compared to fine-grained history under candidate unknown setting.

Table 4 shows the ablation results under the candidate entity known setting. Performance declined when either the entity similarity module (SLM) or the temporal evolution module (TEM) is removed. In contrast to the candidate unknown setting, the candidate known setting demonstrates that removing coarse-grained history has a more significant impact on model performance compared to removing fine-grained history, causing a 17.2% performance degradation for MRR compared with MGESL. This is because when we have knowledge of the candidate entities, the background graph that we build using these entities can serve as an effective means to understand and learn the relationships between them.

5.4 Convergence Analysis

Figure 3 presents the convergence analysis results of our study on ICEWS14 dataset. Obviously, after the initial training epoch, "MGESL w/o Fine" falls noticeably behind the other models in terms of MRR metrics, and requires more epochs to attain the optimal performance compared to the other models as shown in Figure 3(a). This demonstrates that fine-grained history can serve as a good guide for the model to learn during the training process.

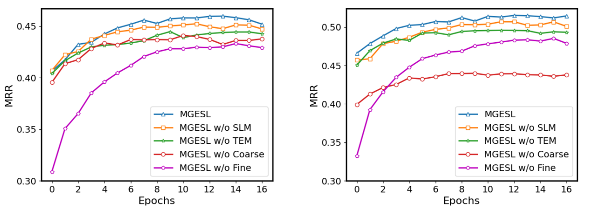

(a) candidate unknown setting (b) candidate known setting

Figure 3: Convergence analysis results on ICEWS14 in MRR.

Similarly, as shown in Figure 3(b), we notice that after the initial epoch of training, the results of "MGESL w/o Fine" are still the lowest. Besides, the results of "MGESL w/o Coarse" no longer remain almost the same with other models as in Figure 3(a). This observation indicates that both coarse-grained and fine-grained histories are crucial in facilitating the model's convergence during training, particularly when the candidate entities are known. The fine-grained history can make the model converges faster, while the coarse-grained history can improve the accuracy of the model to a great extent. These findings further validate the effectiveness of our capturing historical information from various granularities.

6 Conclusion

In this paper, we propose a new model named MGESL for temporal knowledge graph reasoning. This model considers entity similarity learning, coarse-grained history and fine-grained history simultaneously. To capture the similarities between entities, we design a hypergraph convolutional aggregator. We also construct a background graph to effectively capture the coarse-grained history and utilize the fine-grained history to guide the decoder during the decoding process. Moreover, we introduce a more realistic setting for TKG extrapolation, i.e., candidate entities are known in advance. Extensive experiments on three real datasets demonstrate the effectiveness of MGESL compared to the baseline models.

Limitations

In this section, we discuss the limitations of our proposed model. Specifically, under the candidate entity known setting, the candidate entities are deliberately leaked out in our experiments. Nevertheless, there are instances in which we lack knowledge about the candidate entities, making it necessary to make predictions under the candidate entity unknown setting. Unfortunately, this often leads to a significant decrease in the prediction accuracy as our experiments show. As part of our future research, we will focus on exploring methods to improve the accuracy of predicting candidate entities.

References

Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha P. Talukdar. 2018. Hyte: Hyperplane-based temporally aware knowledge graph embedding. In *Proceedings of the 22th Conference on Empirical Methods in Natural Language Processing*, pages 2001–2011, Brussels, Belgium. Association for Computational Linguistics.

Songgaojun Deng, Huzefa Rangwala, and Yue Ning. 2020. Dynamic knowledge graph based multi-event forecasting. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 1585–1595, New York, NY, USA. Association for Computing Machinery.

Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence*, pages 1811–1818, New Orleans, Louisiana, USA. AAAI Press.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 17th Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, Minneapolis, MN, USA. Association for Computational Linguistics.

Hao Dong, Zhiyuan Ning, Pengyang Wang, Ziyue Qiao, Pengfei Wang, Yuanchun Zhou, and Yanjie Fu. 2023. Adaptive path-memory network for temporal knowledge graph reasoning. In *Proceedings of the 32th International Joint Conference on Artificial Intelligence*, pages 2086–2094, Macao, SAR, China. ijcai.org.

Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. 2021. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *Proceedings of the 9th International Conference on Learning Representations*, Austria. OpenReview.net.

Tingsong Jiang, Tianyu Liu, Tao Ge, Lei Sha, Baobao Chang, Sujian Li, and Zhifang Sui. 2016. Towards time-aware knowledge graph completion. In *Proceedings of the 26th International Conference on Computational Linguistics*, pages 1715–1724, Osaka, Japan. ACL.

Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. 2020. Recurrent event network: Autoregressive structure inference over temporal knowledge graphs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6669–6683, online. Association for Computational Linguistics.

Yujia Li, Shiliang Sun, and Jing Zhao. 2022a. Tirgn: Time-guided recurrent graph network with local-global historical patterns for temporal knowledge graph reasoning. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence*, pages 2152–2158, Vienna, Austria. ijcai.org.

Zixuan Li, Saiping Guan, Xiaolong Jin, Weihua Peng, Yajuan Lyu, Yong Zhu, Long Bai, Wei Li, Jiafeng Guo, and Xueqi Cheng. 2022b. Complex evolutional pattern learning for temporal knowledge graph reasoning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 290–296, Dublin, Ireland. Association for Computational Linguistics.

Zixuan Li, Xiaolong Jin, Saiping Guan, Wei Li, Jiafeng Guo, Yuanzhuo Wang, and Xueqi Cheng. 2021a. Search from history and reason for future: Two-stage reasoning on temporal knowledge graphs. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, pages 4732–4743. Association for Computational Linguistics.

Zixuan Li, Xiaolong Jin, Wei Li, Saiping Guan, Jiafeng Guo, Huawei Shen, Yuanzhuo Wang, and Xueqi Cheng. 2021b. Temporal knowledge graph reasoning based on evolutional representation learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 408–417, Canada. ACM.

Kangzheng Liu, Feng Zhao, Guandong Xu, Xianzhi Wang, and Hai Jin. 2023a. RETIA: relation-entity twin-interact aggregation for temporal knowledge graph extrapolation. In *Proceedings of the 39th IEEE International Conference on Data Engineering*, pages 1761–1774, Anaheim, CA, USA. IEEE.

Xiaoqian Liu, Xiuyun Li, Yuan Cao, Fan Zhang, Xiongnan Jin, and Jinpeng Chen. 2023b. Mandari: Multi-modal temporal knowledge graph-aware sub-graph embedding for next-poi recommendation. In *IEEE International Conference on Multimedia and Expo, ICME 2023, Brisbane, Australia, July 10-14, 2023*, pages 1529–1534. IEEE.

Yushan Liu, Yunpu Ma, Marcel Hildebrandt, Mitchell Joblin, and Volker Tresp. 2022. Tlogic: Temporal logical rules for explainable link forecasting on temporal knowledge graphs. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pages 4120–4127. AAAI Press.

Costas Mavromatis, Prasanna Lakkur Subramanyam, Vassilis N. Ioannidis, Adesoji Adeshina, Phillip R. Howard, Tetiana Grinberg, Nagib Hakim, and George Karypis. 2022. Tempoqr: Temporal question reasoning over knowledge graphs. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, pages 5825–5833. AAAI Press.

Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii. 2018. Learning sequence encoders for temporal knowledge graph completion. In *Proceedings of the 22th Conference on Empirical Methods in Natural Language Processing*, pages 4816–4821, Brussels, Belgium. Association for Computational Linguistics.

Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *Proceedings of the 15th Semantic Web International Conference*, volume 10843, pages 593–607, Heraklion, Crete, Greece. Springer.

Chao Shang, Yun Tang, Jing Huang, Jinbo Bi, Xiaodong He, and Bowen Zhou. 2019. End-to-end structure-aware convolutional networks for knowledge base completion. In *Proceedings of the 33th AAAI Conference on Artificial Intelligence*, pages 3060–3067, Honolulu, Hawaii, USA. AAAI Press.

Haohai Sun, Jialun Zhong, Yunpu Ma, Zhen Han, and Kun He. 2021. Timetraveler: Reinforcement learning for temporal knowledge graph forecasting. In *Proceedings of the 25th Conference on Empirical Methods in Natural Language Processing*, pages 8306–8319, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *Proceedings of the 7th International Conference on Learning Representations*, New Orleans, LA, USA. OpenReview.net.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of the 33nd International Conference on Machine Learning*, pages 2071–2080, New York City, NY, USA. JMLR.org.

Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. 2020. Composition-based multi-relational graph convolutional networks. In *Proceedings of the 8th International Conference on Learning Representations*, Addis Ababa, Ethiopia. OpenReview.net.

Yi Xu, Junjie Ou, Hui Xu, and Luoyi Fu. 2023. Temporal knowledge graph reasoning with historical contrastive learning. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, pages 4765–4773, Washington, DC, USA. AAAI Press.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, CA, USA.

Mengqi Zhang, Yuwei Xia, Qiang Liu, Shu Wu, and Liang Wang. 2023. Learning long- and short-term representations for temporal knowledge graph reasoning. In *Proceedings of the ACM Web Conference 2023*, pages 2412–2422, Austin, TX, USA. ACM.

Cunchao Zhu, Muhao Chen, Changjun Fan, Guangquan Cheng, and Yan Zhang. 2021. Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 4732–4740. AAAI Press.