

DPCL-Diff: Temporal Knowledge Graph Reasoning Based on Graph Node Diffusion Model with Dual-Domain Periodic Contrastive Learning

Yukun Cao, Lisheng Wang *, Luobin Huang

Shanghai University of Electric Power, Shanghai, China
caoyukun@shiep.edu.cn, 2207992575@mail.shiep.edu.cn, 1877098981@mail.shiep.edu.cn

Abstract

Temporal knowledge graph (TKG) reasoning that infers future missing facts is an essential and challenging task. Predicting future events typically relies on closely related historical facts, yielding more accurate results for repetitive or periodic events. However, for future events with sparse historical interactions, the effectiveness of this method, which focuses on leveraging high-frequency historical information, diminishes. Recently, the capabilities of diffusion models in image generation have opened new opportunities for TKG reasoning. Therefore, we propose a graph node diffusion model with dual-domain periodic contrastive learning (DPCL-Diff). Graph node diffusion model (GNDiff) introduces noise into sparsely related events to simulate new events, generating high-quality data that better conforms to the actual distribution. This generative mechanism significantly enhances the model’s ability to reason about new events. Additionally, the dual-domain periodic contrastive learning (DPCL) maps periodic and non-periodic event entities to Poincaré and Euclidean spaces, leveraging their characteristics to distinguish similar periodic events effectively. Experimental results on four public datasets demonstrate that DPCL-Diff significantly outperforms state-of-the-art TKG models in event prediction, demonstrating our approach’s effectiveness. This study also investigates the combined effectiveness of GNDiff and DPCL in TKG tasks.

Introduction

Knowledge graphs (KGs) store a large amount of information about the natural world and have shown great success in many downstream applications, such as natural language processing (Hu et al. 2022; Wang et al. 2020), recommendation system (Xuan et al. 2023; Yu et al. 2022), and information retrieval (Zhou et al. 2022). Traditional KGs typically only contain static snapshots of facts, integrating facts (also known as events) in the form of static relational triples (s, r, o) , where s and o represent the subject and object entities, respectively, and r represents the relationship type. However, in the real world, knowledge evolves and constantly exhibits complex temporal dynamics (Wang et al. 2017; Yin and Cui 2016), which has inspired the construction and application of Temporal Knowledge Graph (TKG). TKG extends the previous static relational triples (s, r, o) to quaternions (s, r, o, t) with

timestamps. Thus, the TKG consists of multiple snapshots where facts in the same snapshot appear simultaneously.

Existing studies have identified two main types of new events in TKG: periodic events that will recur (Liang et al. 2023; Li et al. 2021), and events that have not occurred before but may occur in the future (Xu et al. 2023). For these new events, TKG reasoning offers fresh perspectives and insights for many downstream applications, such as event prediction (Zhou et al. 2023), political decision making (Zhang and Zhou 2022), dialogue generation (Hou et al. 2024), and text generation (Li et al. 2022b). These applications have greatly stimulated intense interest in TKG. In this work, we focus on predicting new facts in future time, a task known as graph extrapolation (Chen et al. 2023; Li et al. 2021). Our goal is to predict the missing entities in the query $(s, p, ?, t)$ for future timestamps t that still need to be observed in the training set.

Most current work models the structural and temporal characteristics of TKG to capture the specific relationships and temporal dependencies between different events for future event prediction. Many studies (Liang et al. 2023; Xu et al. 2023; Li et al. 2021) can predict repetitive or periodic events by referring to the known historical events and distinguishing the different effects of periodic and non-periodic events on reasoning tasks through contrastive learning. However, in actual reasoning tasks, some periodic events share the same head entities and relations, differing only in their tail entities. This results in overly similar representations, making it difficult to distinguish them during reasoning. Additionally, in event-based TKG, new events that have never occurred account for about 40% (Jäger 2018). Due to the sparse traces of temporal interactions throughout the timeline of these new events, it is impossible to use reasoning methods based on high-frequency periodic events, resulting in poor reasoning performance for these types of events.

To address the above issue, we propose a diffusion model-based dual-domain periodic contrast learning (DPCL-Diff) for temporal knowledge graph reasoning (TKG), using a Graph Node Diffusion model (GNDiff) to infer new events. GNDiff solves the sparse interaction traces for new events by adding noise to correlated sparse events of new entities on the timeline. This diffusion mechanism simulates the real-world occurrence of new events, generating high-quality data that closely align with the actual distribution through gradual correction during the denoising process. It also captures the

*Lisheng Wang is the corresponding author.

event graph structure, enriching semantic representations and improving the model’s ability to reflect real-world dynamics. The diffusion model (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020) is a state-of-the-art generative framework with proven strength in both image (Ramesh et al. 2022; Rombach et al. 2022; Saharia et al. 2022) and text generation (Ou and Jian 2024; Gong et al. 2022; Strudel et al. 2022). This enhances new event inference and reasoning in TKGs. For tasks with large temporal interaction data, we propose a novel Dual-Domain Periodic Contrastive Learning (DPCL) method. Here, periodic and non-periodic event entities are mapped into Poincaré and Euclidean spaces for contrastive learning, respectively. The Poincaré space is used to better distinguish similar periodic entities, capturing entity relationships more effectively than Euclidean space, leading to more accurate link prediction (Balazevic, Allen, and Hospedales 2019; Zeb et al. 2024).

The main contributions of this thesis are summarized below:

- We propose a TKG model called DPCL-Diff. DPCL-Diff can reason not only about periodic events but also about new events through a diffusion generation mechanism.
- GNDiff is introduced in the field of TKG. GNDiff can perform non-Markov decision chain-based graph node diffusion for new events, which generates a large amount of high-quality graph data related to new events.
- We propose the DPCL method, where periodic and non-periodic event entities are mapped into Poincaré and Euclidean spaces, respectively, to more accurately capture the relationships between periodic event entities and distinguish similar ones.
- We conducted experiments on four public datasets. The results show that DPCL-Diff outperforms the state-of-the-art TKG model in event prediction. Additionally, we purposefully explore the effectiveness of GNDiff and DPCL in the TKG task.

Related Work

Diffusion Model

Diffusion models (Ho, Jain, and Abbeel 2020) are latent variable models for continuous data, achieving state-of-the-art image and audio generation quality. For discrete data, prior work explored text diffusion in discrete spaces (Li et al. 2024; Lezama et al. 2022; Austin et al. 2021). Li et al. (2022a) enables gradient-based control of text diffusion with continuous latent representations. Ou and Jian (2024) combines discrete diffusion and denoising to improve text generation quality and diversity in both unconditional and conditional tasks.

Temporal Knowledge Graph Reasoning

TKG reasoning involves interpolation and extrapolation (Chen et al. 2023; Li et al. 2021). Interpolation fills missing events within t_0 to t_n , while extrapolation predicts events beyond t_n . Interpolation models like TeAST (Li, Su, and Gao 2023), Re-Temp (Wang, Han, and Poon 2023), and TempCaps (Fu et al. 2022) address missing relationships but cannot predict future events. Extrapolation methods focus on future

events: Know-Evolve (Trivedi et al. 2017) captures nonlinear entity evolution but struggles with long-term dependencies, while CyGNet (Zhu et al. 2021) and CENET (Xu et al. 2023) leverage repetitive patterns and contrastive learning. TAR-GAN (Xie et al. 2023) improves long-term modeling with attention mechanisms. Models like DynaQ (Liu et al. 2024a), EvoGNN (Pareja et al. 2020), and TempoR (Wu et al. 2020) enhance prediction but face challenges in capturing temporal evolution, semantic dependencies, and scalability.

Method

As shown in Fig. 1, for new events, the DPCL-Diff model leverages GNDiff to inject noise into the sparse correlated events of the new event entity. This approach simulates the real-world mechanism of new events, generating numerous high-quality data samples. For other events, the DPCL method identifies highly correlated entities through contrastive learning and more accurately captures the relationships between periodic event entities using the properties of Poincaré space, thereby distinguishing similar periodic entities. In the following sections, we describe our proposed approach in detail.

Preliminaries

TKG $G = \{G_0, G_1, \dots, G_t, \dots, G_T\}$ is a sequence of KGs, each $G_t = \{\mathcal{E}, \mathcal{R}, \mathcal{T}\}$ contains facts that occur at timestamp t . Let \mathcal{E} , \mathcal{R} and \mathcal{T} denote a finite set of entities, relation types, and timestamps, respectively. Each fact is a quadruple (s, r, o, t) , where $s \in \mathcal{E}$ is the subject entity, $o \in \mathcal{E}$ is the object entity, and $r \in \mathcal{R}$ is the relation at timestamp t between s and o that occurs, the dimension is d . $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times d}$ is the embedding of all entities, and its rows represent the embedding vectors of an entity, such as s and o . Similarly, $\mathbf{R} \in \mathbb{R}^{|\mathcal{R}| \times d}$ is the embedding of all relation types. The TKG inference task aims to predict missing objects by answering a query $q = (s, r, ?, t)$ using a given history of TKG.

For a given query $q = (s, r, ?, t)$, we define the set of periodic events $Q_t^{s,r}$ as the set of events that have the same head entity relationship with the query event, and the corresponding set of periodic event entities as $\mathcal{H}_t^{s,r}$ in the following equations:

$$Q_t^{s,r} = \bigcup_{k < t} \{(s, r, o, k) \in \mathcal{G}_k\}, \quad (1)$$

$$\mathcal{H}_t^{s,r} = \{o | (s, r, o, k) \in Q_t^{s,r}\}. \quad (2)$$

Entities outside $\mathcal{H}_t^{s,r}$ are non-periodic. If an event (s, r, o, t) does not exist in $Q_t^{s,r}$, it is a new event that has never occurred and is handled by the GNDiff. For periodic event reasoning, DPCL is primarily used to predict target entities. The following sections detail these modules.

Graph Node Diffusion Model

In this section, we introduce GNDiff, demonstrating the effectiveness of diffusion models in TKGs. While diffusion models have been applied to discrete textual data, they are equally effective for discrete graph nodes in TKGs. Traditional methods, reliant on historical data, struggle with new events due to sparse data. In contrast, GNDiff uses a diffusion

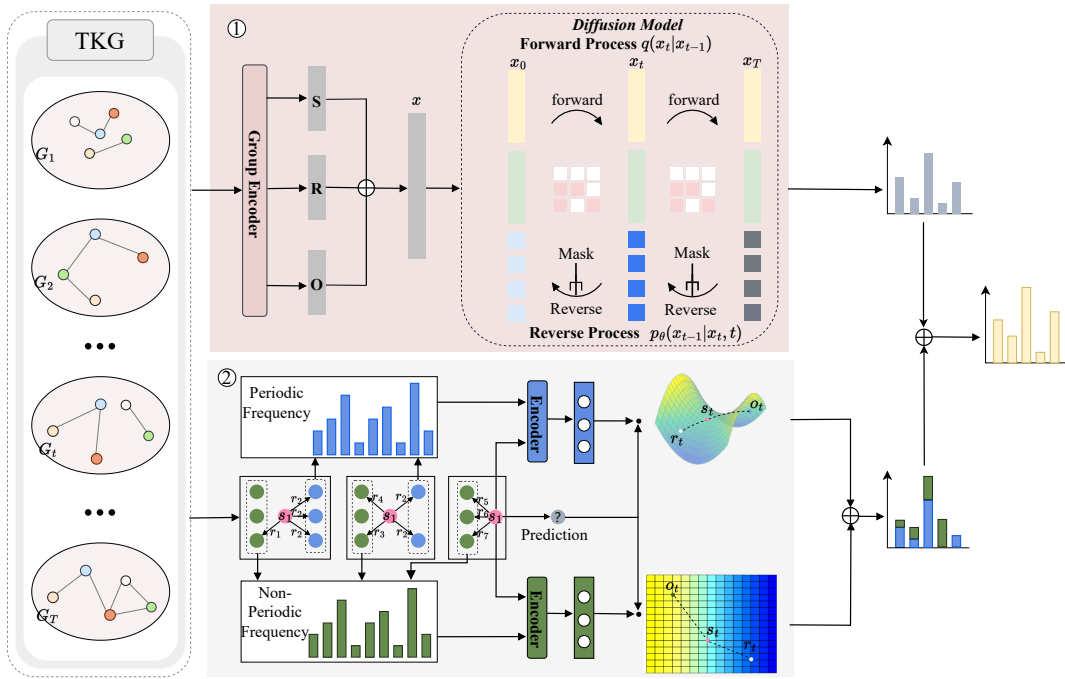


Figure 1: Overall architecture of DPCL-Diff. The first part is GNDiff, and the second part is DPCL.

model to simulate event generation, producing high-quality samples for new event inference.

GNDiff includes a forward and a reverse diffusion process. For each quadruple (s, r, o, t) , the vector representations of s , r , and o are concatenated, forming a single vector as input to the graph node diffusion model. The sample data $x_0 \sim q(x_0)$ is then obtained. Since TKGs are discrete, each element of x_t is a discrete random variable with K categories, where $K = |\mathcal{V}|$ is the number of node types. The noise addition process is represented by a stack of heat vectors as shown in Equation 3:

$$q(x_t | x_{t-1}) = \text{Cat}(x_t; p = x_{t-1}Q_t), \quad (3)$$

where $\text{Cat}(\cdot)$ denotes that the noise addition process is based on the category distribution, and Q_t is used to control the transition matrix of the node from one state to another:

$$[Q_t]_{i,j} = q(x_t = j | x_{t-1} = i). \quad (4)$$

The conversion relationship is shown in Equation 5:

$$\begin{aligned} q(x_{t-1} | x_t, x_0) &= \frac{q(x_t | x_{t-1}, x_0)q(x_{t-1} | x_0)}{q(x_t | x_0)} \\ &= \text{Cat}\left(x_{t-1}; p = \frac{x_t Q_t^\top \odot x_0 \bar{Q}_{t-1}}{x_0 \bar{Q}_t x_t^\top}\right), \end{aligned} \quad (5)$$

where $\bar{Q}_t = Q_1 Q_2 \cdots Q_t$. Note that \odot denotes element-wise multiplication, and division is performed rowwise. To incorporate the pre-trained language model into the denoising process, we introduce a specific "inactive" state as the absorbing state, allowing for the generation of samples during denoising. We use the mask mechanism to simulate the possible inactive state of the node. The transition matrix adjustment process is shown in Equation 6:

$$[Q_t]_{i,j} = \begin{cases} 1 & \text{if } i = j = [M], \\ \beta_t & \text{if } j = [M], i \neq [M], \\ 1 - \beta_t & \text{if } i = j \neq [M]. \end{cases} \quad (6)$$

Here, $[M]$ denotes a node in a masked state (an inactive node), allowing nodes to transition naturally in and out of this state as the graph evolves, eventually converging to a smooth distribution $q(x_T)$. This approach concentrates all probability mass on a sequence labelled with a mask. By using $\bar{Q}_t = Q_1 Q_2 \cdots Q_t$, the transfer relation for multiple time steps can be derived. Samples are then generated through a backward diffusion process:

$$p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1} | x_t, t). \quad (7)$$

We optimize the inverse process by minimizing the variational lower bound to maximize the log-likelihood of the original data. This approach derives the optimization objective $p_\theta(x_{0:T})$ and controls the optimization through $\mathcal{L}_{\text{diff}}$:

$$\begin{aligned} \mathcal{L}_{\text{diff}} &= \mathbb{E}_q[\text{D}_{\text{KL}}(q(x_T | x_0) \| p_\theta(x_T))] + \\ &\mathbb{E}_q \left[\sum_{t=2}^T \text{D}_{\text{KL}}(q(x_{t-1} | x_t, x_0) \| p_\theta(x_{t-1} | x_t, t)) \right] \\ &- \log p_\theta(x_0 | x_1), \end{aligned} \quad (8)$$

where E_q denotes the expectation under a noisy data distribution. D_{KL} stands for Kullback-Leibler divergence.

To better control noise addition, GNDiff employs a noise schedule (He et al. 2022) that measures noise at each step by masking and evenly distributing information. Tokens are assigned varying mask probabilities, ordered by informativeness into T stages. The forward process begins by masking the most informative tokens and ends with the least informative, simplifying the reverse process. Specifically, mask

information is evenly distributed during the forward step, as shown in Equation 9:

$$\sum_{i=1}^n \bar{\alpha}_t^i H(x_0^i) = \left(1 - \frac{t}{T}\right) \sum_{i=1}^n H(x_0^i), \quad (9)$$

where H represents entropy, quantifying the information content of a node. x_0^i denotes the i th token in the sequence and n is the sequence length. $\bar{\alpha}_t^i = \prod_{j=1}^t (1 - \beta_j^i)$ is the probability that the node retains its original state at time step t , i.e. $x_t^i = x_0^i$. We expect that if $H(x_t^i) < H(x_t^j)$, then $\bar{\alpha}_t^i > \bar{\alpha}_t^j$, so low-information tokens appear earlier in the reverse process than high-information ones. In practice, a token's entropy $H(x)$ is computed as the negative logarithm of its frequency in the training corpus. Based on these properties, we construct $\bar{\alpha}_t^i$ as follows:

$$\bar{\alpha}_t^i = 1 - \frac{t}{T} - G(t) \cdot \tilde{H}(x_0^i), \quad (10)$$

$$G(t) = \mu \sin \frac{t\pi}{T}, \quad (11)$$

$$\tilde{H}(x_0^i) = 1 - \frac{\sum_{j=1}^n H(x_0^j)}{nH(x_0^i)}, \quad (12)$$

where $G(t)$ controls the influence of information at time step t , with its effect modulated by the hyperparameter μ .

Dual-Domain Periodic Contrastive Learning

Inspired by Xu et al. (2023), we propose a dual-domain periodic contrastive learning method. Unlike the original approach, we use a dual-domain mapping strategy because we observed that specific queries involve multiple periodic events that are challenging to distinguish. To address this issue, we map entities from periodic and non-periodic events to Poincaré and Euclidean spaces, respectively, to identify highly correlated entities within each domain. By leveraging the spatial characteristics of Poincaré space, we more accurately capture the relationships between periodic entities and distinguish similar ones.

During data preprocessing, we analyze entity frequencies for a given query $q = (s, r, ?, t)$, as shown in Equation 13:

$$\mathbf{Z}_t^{s,r}(o) = \lambda \cdot (\Phi_{\beta>0} - \Phi_{\beta=0}), \quad (13)$$

$$\beta = \sum_{(s,r,o_i,k) \in G_k} (o_i = o), \quad (14)$$

where the value of each entity frequency is constrained by the hyperparameter λ . Φ_β is an indicator function that returns 1 if β is true and 0 otherwise. $\mathbf{Z}_t^{s,r}(o) > 0$ indicates that quadruple (s, r, o, t_k) is a periodic event bound to s and r ; otherwise, it is a non-periodic event. DPCL then learns dependencies in periodic and non-periodic events based on the input $\mathbf{Z}_t^{s,r}(o)$. The distance between entities in Poincaré space is used as the final similarity score between an entity $o_i \in \mathcal{E}$ and a query q . The formula for calculating the Poincaré distance is provided as follows:

$$\mathbf{d}_p(s, o) = \cosh^{-1} \left(1 + 2 \frac{\|s - o\|_2^2}{(1 - \|s\|_2^2)(1 - \|o\|_2^2)} \right), \quad (15)$$

$$\mathbf{S}_p^{s,r}(o_i) = \tanh(\mathbf{W}_p s \oplus r + \mathbf{b}_p) \mathbf{E}^T + \mathbf{Z}_t^{s,r} + \mathbf{d}_p(s, o_i), \quad (16)$$

where $\mathbf{S}_p^{s,r}$ denotes the periodic dependency score for different object entities, \tanh is the activation function, \oplus denotes the connection operator, and $\mathbf{W}_p \in \mathbb{R}^{d \times 2d}$, $\mathbf{b}_p \in \mathbb{R}^{d \times 2d}$

are the trainable parameters. We compute the initial similarity score between each entity $o_i \in \mathcal{E}$ and the query q using a linear layer with activation.

For non-periodic dependencies, we calculate the similarity scores $d_E(s, o) = \|s - o\|$ using the following Euclidean distance formula to obtain the non-periodic dependency score $\mathbf{S}_{np}^{s,r}$:

$$\mathbf{S}_{np}^{s,r}(o_i) = \tanh(\mathbf{W}_{np} s \oplus r + \mathbf{b}_{np}) \mathbf{E}^T - \mathbf{Z}_t^{s,r} + d_E(s, o_i). \quad (17)$$

The training objective of learning from periodic and non-periodic events is to minimize the loss \mathcal{L}_{ce} :

$$\mathcal{L}_{ce} = - \sum_{o_i \in \mathcal{I}} \log \left\{ \frac{\exp(\mathbf{S}_p^{s,r}(o_i))}{\sum_{o_j \in \mathcal{E}} \exp(\mathbf{S}_p^{s,r}(o_j))} + \frac{\exp(\mathbf{S}_{np}^{s,r}(o_i))}{\sum_{o_j \in \mathcal{E}} \exp(\mathbf{S}_{np}^{s,r}(o_j))} \right\}. \quad (18)$$

Here, o_i denotes the ground truth object entity for a given query q . The goal of \mathcal{L}_{ce} is to distinguish the ground truth from other values by comparing each scalar in $\mathbf{S}_p^{s,r}$ and $\mathbf{S}_{np}^{s,r}$.

The model employs supervised contrastive learning to enhance its discriminative power by enforcing consistency within the same class and differentiation between different classes, controlled by the supervised contrastive loss \mathcal{L}_{sup} :

$$\mathcal{L}_{sup} = \sum_{q \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_q \cdot z_p / \tau)}{\sum \exp(z_q \cdot z_a / \tau)}, \quad (19)$$

where I is the set of positive samples, $P(i)$ is the set of queries with the same label in the small batch, z_q denotes the embedding vector of query q , z_p represents the embedding vector of queries sharing the same label as q , z_a denotes the embedding vector of all non- q queries in the small batch.

Training Details and Inference

DPCL-Diff minimizes the loss function throughout model training:

$$\mathcal{L}_{dpcl} = \alpha \cdot \mathcal{L}_{diff} + (1 - \alpha) \cdot (\mathcal{L}_{ce} + \mathcal{L}_{sup}), \quad (20)$$

where α is the equilibrium coefficient.

In GNDiff, the inference process closely mirrors the generation process, so candidate entity probability is computed via reverse diffusion.

$$\mathbf{P}_{diff}(o_i | q) = \prod_{t=1}^T p_\theta(x_{t-1} | x_t), \quad (21)$$

where o_i is the tail entity representation in the generated sample x_t , and x_t represents the sample corresponding to the query q .

For DPCL the probability distribution of entities is calculated based on the predicted missing objects:

$$\mathbf{P}_{dpcl}(o_i | q) = \frac{\exp(s(o_i | q))}{\sum_{o_j \in \mathcal{E}} \exp(s(o_j | q))}. \quad (22)$$

The final probability $\mathbf{P}(o_i)$ is obtained by averaging the probabilities from the diffusion process and the DPCL prediction:

$$\mathbf{P}(o_i | q) = \frac{1}{2} (\mathbf{P}_{diff}(o_i | q) + \mathbf{P}_{dpcl}(o_i | q)). \quad (23)$$

DPCL-Diff will choose the object with the highest probability as the final prediction o_i .

Method	ICEWS18				ICEWS14				WIKI			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
ComplEx(2016)	30.09	21.88	34.15	45.96	24.47	16.13	27.49	41.09	47.84	38.15	50.08	58.46
R-GCN(2018)	23.19	16.36	25.34	36.48	26.31	18.23	30.43	45.34	37.57	28.15	39.66	—
ConvE(2018)	36.67	28.51	39.80	50.69	40.73	33.20	43.92	54.35	47.57	38.76	50.1	63.74
RE-Net(2019)	42.93	36.19	45.47	55.80	45.71	38.42	49.06	59.12	51.97	48.01	52.07	53.91
xERTE(2020)	36.95	30.71	40.38	49.76	32.92	26.44	36.58	46.05	58.75	58.46	58.85	59.34
CyGNet(2021)	<u>46.69</u>	40.58	<u>49.82</u>	57.14	48.63	41.77	52.50	60.29	45.5	50.48	50.79	68.95
EvoKG(2022)	29.67	12.92	33.08	<u>58.32</u>	18.30	6.30	19.43	39.37	50.66	12.21	63.84	64.73
RPC(2023)	34.91	24.34	38.74	55.89	44.55	34.87	49.80	65.08	—	—	—	—
CENET(2023)	43.72	37.91	45.65	54.79	<u>51.93</u>	<u>48.24</u>	52.81	59.33	<u>66.87</u>	<u>66.77</u>	<u>66.93</u>	66.98
DHE-TKG(2024)	29.23	19.15	33.31	—	40.02	30.13	44.99	—	51.20	57.47	69.25	—
RLGNet(2024)	29.90	20.18	33.64	49.08	39.06	29.34	42.03	58.12	64.34	61.03	66.71	69.51
HisRES(2024)	37.69	26.46	42.75	59.70	50.48	39.57	<u>56.65</u>	<u>71.09</u>	—	—	—	—
DPCL-Diff	47.02	40.03	49.93	60.66	66.59	62.89	67.79	73.64	68.44	68.41	68.45	68.55
Improve(%)	7.55 ↑	5.59 ↑	9.38 ↑	10.71 ↑	28.23 ↑	30.37 ↑	28.37 ↑	24.12 ↑	2.35 ↑	2.46 ↑	2.26 ↑	2.34 ↑

Table 1: Experimental results of temporal link prediction on three datasets TKGs. The time-filtered MRR, H@1, H@3, and H@10 metrics are multiplied by 100. The best results are in bold, and the previous best results are underlined. **Improve** refers to the improvement of DPCL-Diff compared to our baseline CENET.

Dataset	#Ent	#Re	#Train	#Valid	#Test
ICEWS18	23,033	256	373,018	45,995	49,545
ICEWS14	12,498	260	535,654	63,788	65,861
WIKI	12,554	24	539,286	67,538	63,110
YAGO	10,623	10	161,540	19,523	20,026

Table 2: Statistics of the datasets.

Experiments

Experimental Setup

Datasets and Evaluation Metrics. To evaluate the effectiveness of the proposed DPCL-Diff, we report the reasoning performance on four public datasets: ICEWS14 (Trivedi et al. 2017), ICEWS18 (Jin et al. 2019), WIKI (Leblay and Chekol 2018), and YAGO (Mahdisoltani, Biega, and Suchanek 2013). We divided all datasets into three groups, roughly partitioning them into train (80%), validation (10%), and test (10%) sets by timestamp. The partitioning of the ICEWS14 dataset differs slightly from the original dataset. Table 2 shows more detailed information about the datasets. In our experiments, we chose the mean reversed rank (MRR) and the hit rate @1,3,10 (Liang et al. 2023) as the metrics.

Baselines. DPCL-Diff is compared with both static and temporal KG reasoning models. Static KG reasoning models, which do not consider temporal information, include ComplEx (Trouillon et al. 2016), R-GCN (Schlichtkrull et al. 2018) and ConvE (Dettmers et al. 2018). In contrast, temporal KG reasoning models are designed to process temporal information and capture evolving patterns in TKG, such as RE-Net (Jin et al. 2019), xERTE (Z et al. 2020), CyGNet (Zhu et al. 2021), EvoKG (Park et al. 2022), RPC (Liang et al. 2023), CENET (Xu et al. 2023), DHE-TKG (Liu et al. 2024b), RLGNet (Lv et al. 2024), and HisRES (Zhang et al. 2024).

Implementation Details. For all datasets, DPCL uses an

Method	YAGO			
	MRR	Hits@1	Hits@3	Hits@10
ComplEx(2016)	61.29	54.88	62.28	69.42
R-GCN(2018)	41.3	32.56	44.44	—
ConvE(2018)	62.32	56.19	63.97	70.44
RE-Net(2019)	65.16	63.29	65.63	68.08
xERTE(2020)	58.75	58.46	58.85	60.48
CyGNet(2021)	63.47	64.26	65.71	68.95
EvoKG(2022)	55.11	54.37	81.38	92.73
CENET(2023)	79.64	<u>78.65</u>	79.91	81.42
DHE-TKG(2024)	62.93	71.00	82.72	—
RLGNet(2024)	<u>80.17</u>	76.52	<u>83.57</u>	84.96
DPCL-Diff	84.45	84.23	84.65	84.75
Improve(%)	6.04 ↑	7.09 ↑	5.86 ↑	4.09 ↑

Table 3: Experimental results of temporal link prediction on YAGO.

embedding dimension of 200, while GNDiff uses 128. We set the batch size to 64, the learning rate to 0.001, and used the Adam optimizer. The training calendar element limit for L is 30, and the comparison learning calendar element limit for the second stage is 20. The value of the hyperparameter α is set to 0.2, and λ is set to 2. Additionally, τ is the temperature parameter set to 0.1 in the supervised contrastive loss \mathcal{L}_{sup} . For baseline settings, we use their recommended configuration. All experiments were implemented on servers with RTX 3090 GPUs equipped with Intel Xeon Gold 6330 CPUs.

Main Results

We compared the experimental results in Table 1 and Table 3, presenting the MRR and Hits@1/3/10 metrics of DPCL-Diff across four datasets. DPCL-Diff consistently outperforms 12

Method	ICEWS18				YAGO			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
w/o- GNDif	44.93	38.22	47.53	57.08	81.30	80.64	81.42	82.26
w/o- DPCL	46.51	39.92	48.92	59.67	84.21	84.19	84.23	84.34
DPCL-Diff	47.02	40.03	49.93	60.66	84.45	84.23	84.65	84.75

Table 4: DPCL-Diff ablation studies on YAGO and ICEWS18 datasets.

Method	ICEWS14(new events)				ICEWS18(new events)				WIKI(new events)			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
GNDiff	47.36	41.56	49.04	58.72	37.46	31.64	39.71	48.42	62.51	61.97	62.77	63.37
DPCL	42.37	36.41	44.78	53.03	33.93	29.47	35.21	42.18	57.16	56.81	57.28	57.58

Table 5: Performance of Graph Diffusion Models on New Events in ICEWS14 and ICEWS18 and WIKI.

Method	YAGO(new events)			
	MRR	Hits@1	Hits@3	Hits@10
GNDiff	79.49	79.44	79.48	79.54
DPCL	71.88	71.78	71.89	71.98

Table 6: Performance of Graph Diffusion Models on New Events in YAGO.

New Events Dataset	#Train	#Valid	#Test
ICEWS18	221,000	31,407	33,211
ICEWS14	275,762	35,744	34,641
WIKI	113,629	30,317	28,811
YAGO	44,949	10,704	10,985

Table 7: Statistics of New Events Datasets.

baseline models, with a notable **30.37% (Hits@1)** improvement over CENET on ICEWS14. This substantial gain is not solely attributable to the effective data generation of GNDiff for handling unseen events but also to its ability to generalize on periodic events. ICEWS14 high proportion of new events and large data volume amplify GNDiff impact, while the pre-trained language model (PLM) further enhances contextual representation. Additionally, DPCL-Diff’s ability to model structural and semantic relationships across events improves performance on both new and periodic events, with the sensitivity of the Hits@1 metric to ranking adjustments further emphasizing the improvement.

On the YAGO and WIKI datasets, DPCL-Diff also achieves strong overall performance, with significantly higher Hits@1 compared to other models, even though Hits@10 on YAGO is slightly lower than RLGNet. The model’s ability to capture relationships between periodic events enhances prediction accuracy despite the lower proportion of new events and a higher proportion of periodic events, which limit GNDiff’s direct impact. Furthermore, the experiments were conducted with a relatively small batch size of 64, potentially constraining the model’s performance. Increasing the batch size to 512 or 1024 could further enhance its capacity to

capture contextual and relational information, especially in datasets with more complex event patterns.¹²

Ablation Study

To investigate the effectiveness of the modules in DPCL-Diff, we verified the effectiveness of GNDiff and DPCL through ablation experiments. The ablation results are shown in Table 4. Here, "w/o-GN Diff" denotes the removal of the graph node diffusion model, and "w/o-DPCL" denotes the removal of the Poincaré mapping, which maps both positive and negative samples into the Euclidean space for contrast learning. These results illustrate the effectiveness of our proposed GNDiff and DPCL. The performance degradation is particularly noticeable after the removal of GNDiff. This demonstrates that generating a large number of high-quality data related to the predicted event through GNDiff significantly affects event prediction. DPCL leverages the properties of Poincaré space to better capture the relationships between periodic entities and distinguish similar periodic events, thereby enhancing inference accuracy.

Effectiveness Analysis of Graph Diffusion in New Events

To evaluate the effectiveness of graph diffusion models on new events, we curated datasets from four sources by selecting the first occurrences of distinct events, thereby focusing exclusively on new data. We extracted datasets based on the first appearance of all events that occurred, ensuring that only new and unique occurrences were included, while subsequent appearances of the same events were excluded. Detailed dataset statistics are provided in Table 7, which supports our methodological approach. The performance comparisons, detailed in Table 5 and Table 6, show that the GNDiff model consistently outperforms the DPCL model across all metrics and datasets. This robust performance highlights the model’s capability to adeptly handle dynamic, unseen data, proving GNDiff’s exceptional utility in environments that require the continuous integration and processing of new information.

¹<https://github.com/Wangls123/DPCL-Diff>

²<https://doi.org/10.48550/arXiv.2411.01477>

Method	YAGO			
	MRR	Hits@1	Hits@5	Hits@10
Euc/Hyp	81.20	80.60	81.26	82.08
Hyp/Hyp	81.23	80.62	81.30	82.14
Euc/Euc	79.64	78.65	79.91	81.42
Hyp/Euc	81.30	80.64	81.42	82.26

Table 8: Experimental results of DPCL under different spatial mapping strategies

Dual Domain Mapping Strategy Analysis

In this study, we examined the effects of different spatial mapping strategies on DPCL’s learning performance. Table 8 presents the results on the YAGO dataset without GN-Diff. Euc/Hyp denotes that periodic and non-periodic event entities are mapped to Euclidean and Poincaré spaces, respectively, for comparative learning, Hyp/ Euc denotes that the two types of entities are mapped to Poincaré and Euclidean space, Hyp/Hyp denotes that both types of entities are mapped to Poincaré space. Euc / Euc denotes that the two types of entities are mapped to Euclidean space, respectively. The experimental results show that the Hyp/Euc method is the most effective, while the Euc/Euc method is the least effective. This suggests that mapping periodic and non-periodic event entities to separate spaces can improve the modelling of comparative learning. Specifically, compared to Euclidean space, the geometric properties of the Poincaré space allow for a more accurate capture of the relationships between periodic entities and distinguish similar ones, leading to better link prediction results.

Dual Domain Embedded Visual Analytics

In this section, we visualize entity pairs—(Germany, Berlin) and (Germany, Munich)—with the same relation, embedded in both Euclidean and Poincaré spaces, as shown in Fig. 2. In Euclidean space, the relationships between Germany and the two cities are not well differentiated. In contrast, Poincaré space provides a clearer distinction, better reflecting the relationships between the entities. By leveraging the geometric properties of Poincaré space, we improve the model’s ability to capture entity relationships, better distinguish similar entities, thereby enhancing prediction accuracy.

Hyper-parameter Analysis

DPCL-Diff involves two key hyperparameters, α and λ , whose impact on performance is analyzed for the ICEWS14 and YAGO datasets in Fig. 3. α balances L_{diff} and L_{dpcl} . On ICEWS14, $\alpha = 0.3$ enhances temporal interaction modeling with minimal sensitivity to changes. For YAGO, a lower α is more effective, highlighting the importance of dual-domain contrastive learning for periodic events. For λ , after fixing α , $\lambda = 4$ optimizes performance on ICEWS14 by balancing temporal variations and preventing overfitting. On YAGO, progressively increasing λ improves performance, handling the dataset’s complex structure. These results demonstrate




Figure 2: Illustration of 2D entity embedding learned by Euc (left) and Hyp (right) on YAGO.

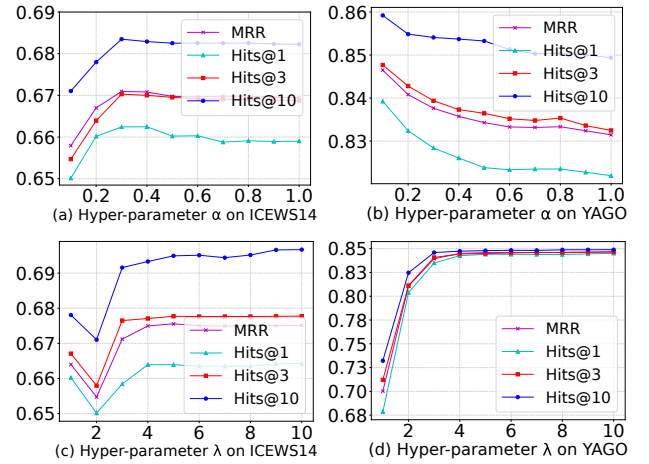


Figure 3: Results of hyper-parameters α and λ of DPCL-Diff on ICEWS14 and YAGO.

DPCL-Diff’s robustness and ability to maintain high performance with stable hyperparameter settings, effectively balancing diffusion and contrastive learning across datasets.

Conclusion

In this paper, we propose a temporal knowledge graph inference method based on a graph node diffusion model with dual-domain periodic contrastive learning. This method aims to alleviate the prediction difficulties caused by the limited interactions of new events on the timeline. We enhance new event reasoning by generating graph data through graph node diffusion, where noise is introduced and adjusted to simulate real-world event mechanisms, producing samples that better reflect actual distributions. Additionally, the performance of the temporal knowledge reasoning task is enhanced by mapping event entities to Poincaré space and Euclidean space through dual-domain periodic contrastive learning, effectively distinguishing between periodic and non-periodic events. Extensive experiments demonstrate the reasoning advantages and effectiveness of DPCL-Diff on public temporal knowledge graph datasets.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62402307, No. U23B2021).

References

- Austin, J.; Johnson, D. D.; Ho, J.; Tarlow, D.; and Van Den Berg, R. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34: 17981–17993.
- Chen, Z.; Xu, C.; Su, F.; Huang, Z.; and Dou, Y. 2023. Temporal Extrapolation and Knowledge Transfer for Lifelong Temporal Knowledge Graph Reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 6736–6746.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Fu, G.; Meng, Z.; Han, Z.; Ding, Z.; Ma, Y.; Schubert, M.; Tresp, V.; and Wattenhofer, R. 2022. TempCaps: a capsule network-based embedding model for temporal knowledge graph completion. In *Proceedings of the Sixth Workshop on Structured Prediction for NLP*, 22–31. Association for Computational Linguistics.
- Gong, S.; Li, M.; Feng, J.; Wu, Z.; and Kong, L. 2022. Difuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*.
- He, Z.; Sun, T.; Wang, K.; Huang, X.; and Qiu, X. 2022. Diffusionbert: Improving generative masked language models with diffusion models. *arXiv preprint arXiv:2211.15029*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hou, Z.; Jin, X.; Li, Z.; Bai, L.; Guo, J.; and Cheng, X. 2024. Selective Temporal Knowledge Graph Reasoning. *arXiv preprint arXiv:2404.01695*.
- Hu, Z.; Xu, Y.; Yu, W.; Wang, S.; Yang, Z.; Zhu, C.; Chang, K.-W.; and Sun, Y. 2022. Empowering language models with knowledge graph reasoning for question answering. *arXiv preprint arXiv:2211.08380*.
- Jäger, K. 2018. The limits of studying networks via event data: Evidence from the ICEWS dataset. *Journal of Global Security Studies*, 3(4): 498–511.
- Jin, W.; Qu, M.; Jin, X.; and Ren, X. 2019. Recurrent event network: Autoregressive structure inference over temporal knowledge graphs. *arXiv preprint arXiv:1904.05530*.
- Leblay, J.; and Chekol, M. W. 2018. Deriving validity time in knowledge graph. In *Companion proceedings of the the web conference 2018*, 1771–1776.
- Lezama, J.; Salimans, T.; Jiang, L.; Chang, H.; Ho, J.; and Essa, I. 2022. Discrete predictor-corrector diffusion models for image synthesis. In *The Eleventh International Conference on Learning Representations*.
- Li, B.; Gao, Z.; Zhu, Y.; Yin, K.; Cao, H.; Jiang, D.; and Xu, L. 2024. Few-shot Temporal Pruning Accelerates Diffusion Models for Text Generation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 7259–7269.
- Li, J.; Su, X.; and Gao, G. 2023. Teast: Temporal knowledge graph embedding via archimedean spiral timeline. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15460–15474.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022a. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35: 4328–4343.
- Li, Z.; Guan, S.; Jin, X.; Peng, W.; Lyu, Y.; Zhu, Y.; Bai, L.; Li, W.; Guo, J.; and Cheng, X. 2022b. Complex evolutionary pattern learning for temporal knowledge graph reasoning. *arXiv preprint arXiv:2203.07782*.
- Li, Z.; Jin, X.; Li, W.; Guan, S.; Guo, J.; Shen, H.; Wang, Y.; and Cheng, X. 2021. Temporal knowledge graph reasoning based on evolutionary representation learning. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 408–417.
- Liang, K.; Meng, L.; Liu, M.; Liu, Y.; Tu, W.; Wang, S.; Zhou, S.; and Liu, X. 2023. Learn from relational correlations and periodic events for temporal knowledge graph reasoning. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, 1559–1568.
- Liu, R.; Yin, G.; Liu, Z.; and Tian, Y. 2024a. Reinforcement learning with time intervals for temporal knowledge graph reasoning. *Information Systems*, 120: 102292.
- Liu, X.; Zhang, J.; Ma, C.; Liang, W.; Xu, B.; and Zong, L. 2024b. Temporal Knowledge Graph Reasoning with Dynamic Hypergraph Embedding. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 15742–15751.
- Lv, A.; Huang, Y.; Ouyang, G.; Chen, Y.; and Xie, H. 2024. RLGNet: Repeating-Local-Global History Network for Temporal Knowledge Graph Reasoning. *arXiv preprint arXiv:2404.00586*.
- Mahdisoltani, F.; Biega, J.; and Suchanek, F. M. 2013. Yago3: A knowledge base from multilingual wikipedias. In *CIDR*.
- Ou, Y.; and Jian, P. 2024. Effective Integration of Text Diffusion and Pre-Trained Language Models with Linguistic Easy-First Schedule. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 5551–5561.
- Pareja, A.; Domeniconi, G.; Chen, J.; Ma, T.; Suzumura, T.; Kanezashi, H.; Kaler, T.; Schardl, T.; and Leiserson, C. 2020. Evolvegn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 5363–5370.
- Park, N.; Liu, F.; Mehta, P.; Cristofor, D.; Faloutsos, C.; and Dong, Y. 2022. Evokg: Jointly modeling event time and

- network structure for reasoning over temporal knowledge graphs. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 794–803.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Van Den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15*, 593–607. Springer.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Strudel, R.; Tallec, C.; Altché, F.; Du, Y.; Ganin, Y.; Mensch, A.; Grathwohl, W.; Savinov, N.; Dieleman, S.; Sifre, L.; et al. 2022. Self-conditioned embedding diffusion for text generation. *arXiv preprint arXiv:2211.04236*.
- Trivedi, R.; Dai, H.; Wang, Y.; and Song, L. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *international conference on machine learning*, 3462–3471. PMLR.
- Trouillon, T.; Welbl, J.; Riedel, S.; Gaussier, É.; and Bouchard, G. 2016. Complex embeddings for simple link prediction. In *International conference on machine learning*, 2071–2080. PMLR.
- Wang, K.; Han, S. C.; and Poon, J. 2023. Re-Temp: Relation-Aware Temporal Representation Learning for Temporal Knowledge Graph Completion. *arXiv preprint arXiv:2310.15722*.
- Wang, Q.; Li, M.; Wang, X.; Parulian, N.; Han, G.; Ma, J.; Tu, J.; Lin, Y.; Zhang, H.; Liu, W.; et al. 2020. COVID-19 literature knowledge graph construction and drug repurposing report generation. *arXiv preprint arXiv:2007.00576*.
- Wang, Q.; Mao, Z.; Wang, B.; and Guo, L. 2017. Knowledge graph embedding: A survey of approaches and applications. *IEEE transactions on knowledge and data engineering*, 29(12): 2724–2743.
- Wu, J.; Cao, M.; Cheung, J. C. K.; and Hamilton, W. L. 2020. Temp: Temporal message passing for temporal knowledge graph completion. *arXiv preprint arXiv:2010.03526*.
- Xie, Z.; Zhu, R.; Liu, J.; Zhou, G.; and Huang, J. X. 2023. TARGAT: A time-aware relational graph attention model for temporal knowledge graph embedding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31: 2246–2258.
- Xu, Y.; Ou, J.; Xu, H.; and Fu, L. 2023. Temporal knowledge graph reasoning with historical contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4765–4773.
- Xuan, H.; Liu, Y.; Li, B.; and Yin, H. 2023. Knowledge enhancement for contrastive multi-behavior recommendation. In *Proceedings of the sixteenth ACM international conference on web search and data mining*, 195–203.
- Yin, H.; and Cui, B. 2016. *Spatio-temporal recommendation in social media*. Springer.
- Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 1294–1303.
- Z, H.; P., C.; Y., M.; and V., T. 2020. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *International Conference on Learning Representations*.
- Zhang, J.; Hui, B.; Mu, C.; Sun, M.; and Tian, L. 2024. Historically Relevant Event Structuring for Temporal Knowledge Graph Reasoning. *arXiv preprint arXiv:2405.10621*.
- Zhang, L.; and Zhou, D. 2022. Temporal knowledge graph completion with approximated Gaussian process embedding. In *Proceedings of the 29th International Conference on Computational Linguistics*, 4697–4706.
- Zhou, W.-T.; Kang, Z.; Tian, L.; and Su, Y. 2023. Intensity-free convolutional temporal point process: Incorporating local and global event contexts. *Information Sciences*, 646: 119318.
- Zhou, Y.; Chen, X.; He, B.; Ye, Z.; and Sun, L. 2022. Rethinking knowledge graph completion evaluation from an information retrieval perspective. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 916–926.
- 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 *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4732–4740.