

Chain-of-History Reasoning for Temporal Knowledge Graph Forecasting

Yuwei Xia^{1,2*}, Ding Wang^{3,4*}, Qiang Liu^{3,4†}, Liang Wang^{3,4}, Shu Wu^{3,4}, Xiaoyu Zhang^{1,2†}

¹Institute of Information Engineering, Chinese Academy of Sciences

²School of Cyber Security, University of Chinese Academy of Sciences

³New Laboratory of Pattern Recognition (NLPR),

State Key Laboratory of Multimodal Artificial Intelligence Systems,
Institute of Automation, Chinese Academy of Sciences

⁴School of Artificial Intelligence, University of Chinese Academy of Sciences

xiayuwei@iie.ac.cn, wangding2024@ia.ac.cn, qiang.liu@nlpr.ia.ac.cn

liang.wang@cripac.ia.ac.cn, shu.wu@nlpr.ia.ac.cn, zhangxiaoyu@iie.ac.cn

Abstract

Temporal Knowledge Graph (TKG) forecasting aims to predict future facts based on given histories. Most recent graph-based models excel at capturing structural information within TKGs but lack semantic comprehension abilities. Nowadays, with the surge of LLMs, the LLM-based TKG prediction model has emerged. However, the existing LLM-based model exhibits three shortcomings: (1) It only focuses on the first-order history for prediction while ignoring high-order historical information, resulting in the provided information for LLMs being extremely limited. (2) LLMs struggle with optimal reasoning performance under heavy historical information loads. (3) For TKG prediction, the temporal reasoning capability of LLM alone is limited. To address the first two challenges, we propose Chain-of-History (CoH) reasoning which explores high-order histories step-by-step, achieving effective utilization of high-order historical information for LLMs on TKG prediction. To address the third issue, we design CoH as a plug-and-play module to enhance the performance of graph-based models for TKG prediction. Extensive experiments on three datasets and backbones demonstrate the effectiveness of CoH.

1 Introduction

As a carrier of facts with temporal information, Temporal Knowledge Graphs (TKGs) hold significant practical value across various applications (Xiang et al., 2022; Chen et al., 2023). Most advanced research on TKGs mainly focuses on predicting future facts occur at time t_n based on given historical facts occur at time t with $t < t_n$.

Recent supervised methods (Jin et al., 2020; Li et al., 2021b, 2022) primarily rely on Graph Neural

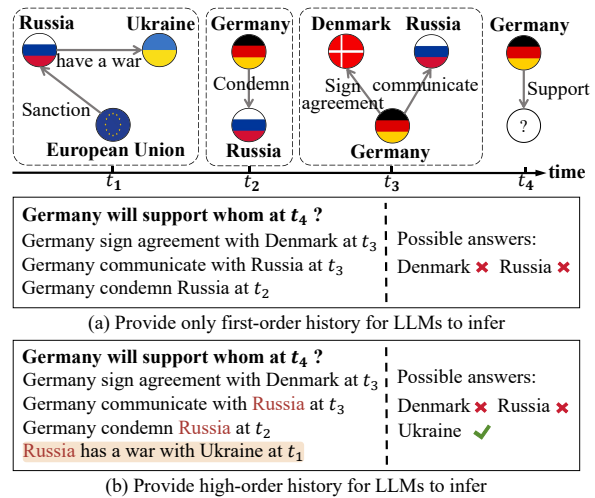


Figure 1: An example of reasoning over TKG with LLMs. In Figures (a) and (b), we provide LLMs with different histories, which prompt LLMs to reason different answers for the predicted fact.

Networks (GNNs) to capture structural dependencies within TKGs, yet they often fall short in effectively modeling semantic information. With the advent of Large Language Models (LLMs), the potential for enhanced temporal reasoning across various tasks is becoming increasingly evident (Jain et al., 2023; Yuan et al., 2023). Lee et al. made the first attempt at TKG reasoning using LLMs (Lee et al., 2023), presenting histories to LLMs in textual form. Despite these advancements, we contend that significant challenges remain unaddressed.

Firstly, the existing TKG prediction model with LLMs only focuses on the first-order histories, ignoring important high-order historical information. Taking Figure 1 as an example, LLMs aim to infer “Germany will support whom at t_4 ” with provided histories. The existing model provides LLMs with only first-order histories. In this case, LLMs are constrained to infer wrong answers to “Denmark”

*These authors contributed to the work equally.

†To whom correspondence should be addressed.

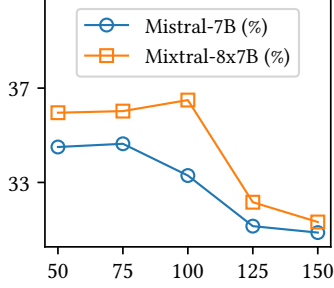


Figure 2: The performance (MRR (%)) of LLMs of two sizes based on different history lengths on TKG prediction. The provided histories contain both first- and second-order histories. The y-axis represents the MRR (%) value, and the x-axis denotes the total length of provided first- and second-order histories. The results are based on the commonly used TKG dataset ICEWS14.

and “Russia” because given histories fail to encompass the correct answer. When supplied with more high-order histories, LLMs can utilize the history chain “Germany→Russia→Ukraine” to reason the correct answer “Ukraine” more possibly.

Secondly, LLMs struggle to maintain reasoning performance under heavy historical information loads. Intuitively, we provide more comprehensive high-order histories for LLMs to infer. However, as shown in Figure 2, the performance of LLMs does not necessarily improve or remain stable with the increase in history length, instead experiencing a steep decline beyond a certain threshold of history length regardless of the model size. This indicates that over-complicated historical information may confuse LLMs (Shi et al., 2023), making LLMs hard to reason correct answers. Thus, exploring ways to offer higher-order histories for LLMs effectively is a worthwhile investigation.

Thirdly, relying solely on the reasoning capabilities of LLMs still remains limited on TKG prediction. Though LLMs possess unique semantic comprehension advantages in reasoning, they still struggle to achieve the same level of ability in capturing complex structural information as graph-based models. However, this unique advantage of LLMs precisely compensates for the shortcomings of graph-based models in modeling semantic information, thereby enhancing the performance of graph-based models on TKG prediction.

To address the above issues, we propose a Chain-of-History (CoH) reasoning method for TKG prediction. Instead of providing LLMs with all histories at once, CoH provides LLMs with high-

order histories step-by-step. Specifically, CoH adopts LLMs to explore important high-order history chains step-by-step, and reason the answers to the query only based on inferred history chains in the last step. A two-step CoH reasoning procedure is shown in Figure 3. This also can extend to multiple-step reasoning for complex situations. LLMs can continue inferring important second-order history chains in Step 2, then infer answers with third-order history chains in Step 3, and so on. In this way, LLMs only need to process a limited quantity of histories at each step, preventing an overwhelming influx of complex information while effectively leveraging a more comprehensive set of high-order information. This approach enables LLMs to perform more accurate reasoning with higher-order information for TKG prediction. Moreover, we design CoH as a plug-and-play module for TKG reasoning. As shown in Figure 3, we fuse the predicted results obtained by LLMs and graph-based TKG models to make the final prediction more comprehensively. To summarize, the contributions of this paper can be listed as follows:

- We are the first to explore the necessity and difficulty of providing numerous high-order histories for LLMs on TKG prediction. And we propose the Chain-of-History reasoning method which adopts LLMs to explore the history chains step-by-step.
- We are the first to propose enhancing the performance of graph-based TKG models with LLMs, utilizing the semantic understanding advantage of LLM to compensate for the shortcoming of graph-based models.
- We conduct extensive experiments on three commonly used TKG datasets and three graph-based TKG models, the results demonstrate the effectiveness of CoH.

2 Problem Formulation

Temporal Knowledge Graph Prediction. Let \mathcal{E} and \mathcal{R} represent a set of entities and relations. A Temporal Knowledge Graph (TKG) \mathcal{G} can be defined as $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$. Each $\mathcal{G}_t \in \mathcal{G}$ contains facts that occur at time t . Each fact is represented as a quadruple (s, r, o, t) , in which $s, o \in \mathcal{E}$ and $r \in \mathcal{R}$. Given a query $q = (s^q, r^q, ?, t^q)$ or $q = (?, r^q, o^q, t^q)$, TKG prediction task aims to predict the missing object entity or subject entity with historical KG sequence $\mathcal{G}_{<t^q} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_{t^q-1}\}$. The candidate answers for q are all entities in \mathcal{E} ,

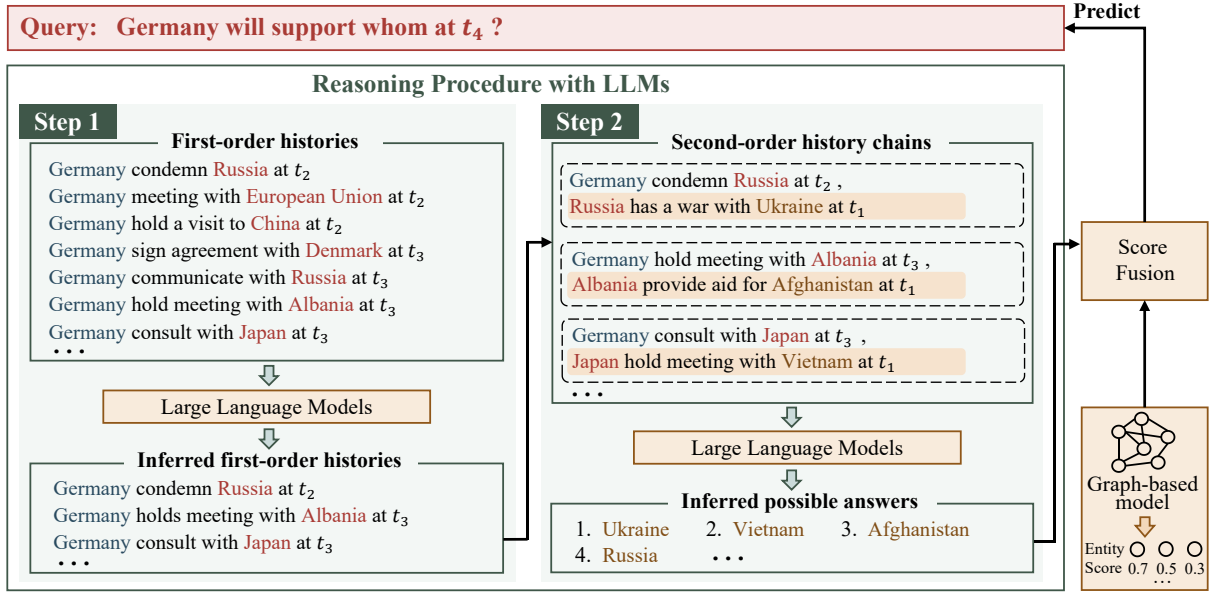


Figure 3: An illustration of a two-step CoH reasoning procedure. In the first step, LLMs are provided with only first-order histories and asked to infer the most important histories. In the second step, LLMs are provided with second-order history chains based on the inferred first-order histories and asked to infer possible answers to the given query. Then the answers inferred by LLMs and graph-based models are adaptively fused to make the final prediction. **Note that this only serves as a two-step reasoning example, more steps can be executed with CoH.**

each candidate $e_i \in \mathcal{E}$ will be estimated with a score by TKG prediction models.

High-order History Chains in TKGs. For each query $q = (s^q, r^q, ?, t^q)$ to be predicted, we denote $\{(s^q, r, o, t) | (s^q, r, o, t) \in \mathcal{G}_{<t^q}\}$ as the first-order histories of q . If (s^q, r, o, t) is the first-order history of q , we denote histories in the form of (o, r, o', t) as a set of second-order histories of q . The higher-order histories of q are deduced in this way. And we denote $[(s^q, r, o, t), (o, r, o', t^j)]$ as a second-order history chain of q , which consists of a first-order history of q and its associated second-order history of q . The higher-order history chains of q can be deduced in this way.

3 Chain-of-History Reasoning over Temporal Knowledge Graph

For a given query $q = (s^q, r^q, ?, t^q)$, CoH predicts the answers by exploring history chains related to q step-by-step, then answers are fused with predicted results by graph-based TKG models to make the final prediction for q . In this section, we illustrate CoH shown in Figure 3 in detail. Section 3.1 explains how to convert quadruples in TKGs into text formats suitable for LLMs. Section 3.2 demonstrates how to properly provide histories from $\mathcal{G}_{<t^q}$ for LLMs and how to instruct LLMs to reason in each step. Section 3.3 shows how to transform an-

swers predicted by LLMs into scores and fuse them with the results from graph-based TKG models.

3.1 History Processing

Each fact in TKGs is presented as a quadruple (s, r, o, t) , such as “(Germany, Sign agreement, Denmark, 2023-06-02)”. To make each quadruple more linguistically comprehensible for LLMs, we introduce prepositions to transform the quadruple into a more fluent sentence. Most importantly, considering the prior knowledge of LLMs, we process the time “2023-06-02” in each quadruple into a more abstract form like “153rd day” to prevent LLMs from directly exploiting prior knowledge for predictions. Consequently, we provide LLMs with each quadruple in the form of “(Germany Sign agreement with Denmark on the 153rd day)”.

3.2 Reasoning Steps

In a k -step reasoning procedure of CoH, LLMs are instructed to explore the most significant history chains related to the given query q from Step 1 to Step $k-1$. Subsequently, LLMs reason the possible answers to q in Step k with k -order history chains.

Step 1 to Step $k-1$ Reasoning. In Step 1, LLMs are provided with only first-order histories of q , and are instructed to reason n first-order histories that mostly contribute to answering q . From Step 2 to

Step i	Instruction
Step 1	There is a given text consisting of multiple historical events in the form of “{id}:{subject} {relation} {object} {time}”;. And there is a query in the form of: “{subject} {relation} {whom} time?” If you must infer several {object} that you think may be the answer to the given query based on the given historical events, what important historical events do you base your predictions on? Please list the top n most important histories and output their {id}.
Step 2 to Step $k-1$	There is a given text consisting of multiple history chains in the form of “{id}:{subject} {relation} {object} {time}, {subject} {relation} {object} {time}, ...”;. And there is a query in the form of: “{subject} {relation} {whom} time?” If you must infer several {object} that you think may be the answer to the given query based on the given historical events, what important history chains do you base your predictions on? Please list the top n most important history chains and output their {id}.
Step k	You must be able to correctly predict the {whom} of the given query from a given text consisting of multiple historical events in the form of “{subject} {relation} {object} {time}” and the query in the form of “{subject} {relation} {whom} {time}?” You must output several {object} that you think may be the answer to the given query based on the given historical events. Please list all possible {object} which may be answers to the query. Please assign each answer a serial number to represent its probability of being the correct answer. Note that answers with a high probability of being correct should be listed first.

Table 1: Instruction design for each step in a k -step CoH reasoning procedure.

Step $k-1$, LLMs are provided with i -order history chains in Step i ($i = \{2, 3, \dots, k-1\}$), and are instructed to infer n most significant history chains. The instruction design is shown in Table 1. Within this sequence of steps, the outputs of LLMs in Step $i-1$ are the inferred ($i-1$)-order history chains, then each of which is supplied with corresponding i -order histories to consist of i -order history chains. These history chains subsequently serve as input for the next Step i . As shown in Figure 3, “*Germany condemn Russia at t_2* ” is one of the outputs in Step 1. Then it is supplied with corresponding second-order history “*Russia has a war with Ukraine at t_1* ” to consist of the second-order history chain, which serves as the input for Step 2. The prompt example for Step 1 and Step i ($i = \{2, 3, \dots, k-1\}$) can refer to Appendix A.

Step k Reasoning. In Step k , LLMs are provided with k -order history chains and instructed to reason possible answers for the given query q . The instruction design for Step k is shown in Table 1. Especially, we instruct LLMs to prioritize outputting the entity with a higher probability of being the correct answer. As shown in Figure 3, the output of Step 2 includes several possible answers to q . Each answer is assigned a numerical index (1,2,3,...), with a lower index indicating a higher probability of the answer being correct. The prompt example for Step k can refer to Appendix A.

3.3 Results Processing and Fusion

In graph-based TKG models, each entity and relation in a quadruple are denoted with an id like

“(30, 13, 8, 2023-06-02)” instead of “(Germany, Sign agreement, Denmark, 2023-06-02)”. The lack of semantic modeling of histories makes graph-based models mainly depend on structural information within TKGs for prediction. However, entities and relations inherently carry semantic information, which also constitutes a significant part of TKGs. Considering the importance of both structural and semantic information within TKGs, we propose to fuse the predicted results of LLMs and graph-based models to obtain more comprehensive results for predicting over TKGs more accurately.

Firstly, for a given q , we need to obtain the score of each entity e_i in the LLMs’ predicted answer set $\mathcal{A}_{\text{LLM}}^q$. As we mentioned each answer predicted by LLMs in Step k is assigned an index, which represents the probability of the answer being correct. We convert the index of each answer $e_i \in \mathcal{A}_{\text{LLM}}^q$ into its corresponding score with an exponential decay function as follows:

$$\mathcal{S}_{\text{LLM}}^{e_i} = \frac{1}{1 + e^{\alpha \cdot \text{idx}^{e_i}}}, \quad (1)$$

where $\mathcal{S}_{\text{LLM}}^{e_i}$ denotes the score of the entity e_i obtained with LLMs for being the answer to q , idx^{e_i} represents the numerical index of the answer e_i in the outputs of LLMs, and α is a hyper-parameter to control the score disparity among answers with different indexes. Note that since the outputs of LLMs can not include all candidate entities in \mathcal{E} like graph-based models, we assign the score of e_i as 0 where $e_i \in \mathcal{E}$ but $e_i \notin \mathcal{A}_{\text{LLM}}^q$.

Then, we can fuse the score of each candidate entity $e_i \in \mathcal{E}$ obtained with LLMs and graph-based

models as follows:

$$\mathcal{S}^{e_i} = w \cdot \mathcal{S}_{\text{Graph}}^{e_i} + (1 - w) \cdot \mathcal{S}_{\text{LLM}}^{e_i}, \quad (2)$$

where $\mathcal{S}_{\text{Graph}}^{e_i}$ denotes the score of e_i obtained with graph-based models, and w is a hyper-parameter to determine the weight of different scores. \mathcal{S}^{e_i} represents the comprehensive score of the candidate e_i . Finally, the ranked candidate list based on comprehensive scores is used for predicting q .

4 Experiments

In this section, we conduct extensive experiments to evaluate and analyze CoH on three typical datasets and three backbones for TKG prediction. Details of datasets and backbones can be referred to Appendix B and C, respectively.

4.1 Experimental Settings

4.1.1 Evaluation

For evaluation, we adopt widely used metrics MRR and Hits@{1, 3, 10} in experiments. Without loss of generality (Li et al., 2021b), we only report the experimental results under the raw setting. Note that different from (Lee et al., 2023), we fully align the evaluation mechanism for LLMs in TKG prediction with those used in graph-based models to ensure a more fair comparison. Specifically, during the testing phase of graph-based models, the test set is typically augmented by doubling its size through reversing (s, r, o, t) into (o, r^{-1}, s, t) , to assess the model’s performance more comprehensively. Correspondingly, we also evaluate LLMs on TKG prediction with reversed test sets.

4.1.2 CoH Implementation Details

In this paper, we implement CoH with two-step reasoning based on an open-sourced language model Mixtral-8x7B (Jiang et al., 2024). In Step 1, we provide the LLM with 100 first-order histories and set n to 30, allowing the LLM to infer the most important 30 first-order histories from the given ones. In Step 2, we do not strictly limit the number of answers output from the LLM. For more details on implementation please refer to Appendix D.

4.2 Performance Comparison

In this section, we present a comprehensive evaluation of the proposed CoH. We first evaluate the performance of only utilizing LLMs on TKG prediction with CoH reasoning. Then we plug CoH on three existing state-of-the-art graph-based TKG

prediction models to see the potential gains it can yield. The results are shown in Table 2, from which we have the following observations.

On the one hand, from the results of LLMs, the two-step CoH reasoning outperforms ICL (Lee et al., 2023) which solely provides LLMs with first-order histories under all evaluation metrics on three datasets. This indicates the usefulness of the higher-order histories provided step-by-step. And we observe that the relative improvements of CoH over the existing method are more obvious on ICEWS18 than other datasets, which implies that ICEWS18 may contain more and complex information of history chains. Despite the progress achieved by CoH, the performance of only utilizing LLMs on TKG prediction is still pretty limited compared with graph-based models.

On the other hand, though the temporal reasoning capability of LLMs on TKG prediction is relatively limited, they can be flexibly used as a plug-and-play module to enhance the performance of graph-based models. From the results of plugging CoH and the ICL-based model into existing graph-based models, we can see that the two LLM-based models can effectively improve their performance. The gains introduced by CoH to the performance of graph-based models far surpass those achieved by the ICL-based model, which further demonstrates the effectiveness of our proposed model. Moreover, we analyze the rationale behind these gains may be attributed to the distinct reasoning mechanisms of LLMs and graph-based models, each of which possesses unique strengths. In this case, the powerful semantic understanding ability of LLMs may be capable of compensating to some extent for the inherent limitations in semantic information modeling of graph-based models.

4.3 Ablation Study

In this section, we conduct experiments to investigate the effectiveness of the high-order historical information, step-by-step reasoning mechanism, and score ranking procedure in CoH. The results are shown in Table 3.

Analysis of high-order historical information.

To verify the usefulness of the high-order historical information for TKG prediction with LLMs, we use the inferred first-order histories in Step 1 and the second-order history chains in Step 2 to predict answers, respectively. From the results of CoH shown in Table 3 we can observe that, the predicted

Model Type	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
LLMs	ICL (Lee et al., 2023)*	31.79	22.38	37.67	47.70	21.51	14.77	26.08	40.57	35.34	25.18	43.92	56.24
	CoH	34.51	24.20	39.67	51.21	23.94	16.81	28.15	42.68	37.51	27.72	47.17	59.58
	$\Delta Improve$	8.56%	8.13%	5.31%	7.36%	11.30%	13.81%	7.94%	5.20%	6.14%	10.09%	7.40%	5.94%
LLMs + Graph	RE-NET	38.75	28.96	43.64	57.61	28.72	18.84	32.66	48.18	44.05	33.22	51.23	65.02
	RE-NET + ICL*	39.39	29.12	44.37	58.25	29.01	18.98	33.11	48.78	45.12	33.98	52.09	66.23
	RE-NET + CoH	40.43	30.34	45.78	60.42	29.77	19.96	34.14	49.59	46.37	34.99	53.13	67.71
	$\Delta Improve^*$	1.65%	0.56%	1.67%	1.11%	1.00%	0.73%	1.39%	1.25%	2.42%	2.29%	1.67%	1.87%
	$\Delta Improve$	4.34%	4.77%	4.90%	4.87%	3.66%	5.94%	4.53%	2.93%	5.26%	5.33%	3.71%	4.14%
	RE-GCN	41.33	30.61	46.66	62.31	31.08	20.44	35.39	52.06	46.89	35.5	53.33	68.4
	RE-GCN + ICL*	41.84	30.84	47.27	62.97	31.31	20.65	35.79	52.61	47.87	36.16	54.24	69.43
	RE-GCN + CoH	42.41	31.77	47.85	63.80	32.10	21.75	36.51	53.37	47.98	37.53	54.94	70.68
	$\Delta Improve^*$	1.23%	0.75%	1.31%	1.06%	0.74%	1.03%	1.13%	1.06%	2.09%	1.86%	1.71%	1.51%
	$\Delta Improve$	2.61%	3.79%	2.55%	2.39%	3.28%	6.41%	3.16%	2.52%	2.32%	5.72%	3.02%	3.33%
	TiRGN	42.93	32.1	48.53	63.6	31.97	20.95	36.67	53.66	48.5	36.87	55.19	70.27
	TiRGN + ICL*	43.27	32.28	49.04	64.17	32.18	21.07	36.98	54.04	49.15	37.25	55.73	70.9
TiRGN + CoH	43.94	33.07	49.64	64.90	32.98	21.83	37.79	54.92	49.71	38.01	56.40	71.25	
$\Delta Improve^*$	0.79%	0.56%	1.05%	0.90%	0.66%	0.57%	0.85%	0.71%	1.34%	1.03%	0.98%	0.90%	
$\Delta Improve$	2.35%	3.02%	2.29%	2.04%	3.16%	4.20%	3.05%	2.35%	2.49%	3.09%	2.19%	1.39%	

Table 2: Performance comparison of CoH on TKG prediction on three datasets in terms of MRR (%), Hit@1 (%), Hit@3 (%), and Hit@10 (%). All results are obtained under raw metrics. The highest performance is highlighted in bold. And * represents the reproduced model with the same evaluation and LLM used in this paper. $\Delta Improve$ and $\Delta Improve^*$ indicate the relative improvements of CoH and the ICL-based model plugged into the graph-based models over the original graph-based backbones in percentage, respectively.

Model	Step i	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
CoH w/o LR	Step 1	32.31	23.18	37.92	49.15	20.97	13.88	24.83	36.52	34.5	22.24	42.01	56.18
	Step 2	32.68	23.74	38.12	50.09	21.89	14.17	25.32	37.98	34.89	22.48	43.26	57.34
CoH	Step 1	33.97	23.86	39.03	49.96	22.03	15.34	26.82	40.57	36.48	24.84	46.44	58.51
	Step 2	34.51	24.20	39.67	51.21	23.94	16.81	28.15	42.68	37.51	27.72	47.17	59.58
CoH w/o IS	Step 2	24.57	12.31	34.75	51.09	13.8	9.14	20.57	35.67	29.61	18.55	37.64	56.87

Table 3: Ablation studies to investigate the effectiveness of the high-order historical information, step-by-step mechanism, and score ranking procedure of CoH in terms of MRR (%), Hit@1 (%), Hit@3 (%), and Hit@10 (%). And all results are obtained under raw metrics.

results of LLMs based on the second-order history chains are notably superior to the results derived solely from first-order histories. This indicates the usefulness of second-order histories.

Analysis of step-by-step reasoning mechanism. Specifically, “LR” in Table 3 denotes the step of reasoning important first-order histories by LLMs in two-step CoH reasoning. And we implement “CoH w/o LR” by replacing the inferred n first-order histories by LLMs with n first-order histories in the latest timestamps. In this way, we can find out whether LLMs can deduce meaningful historical information within the step-by-step reasoning mechanism. From the results, we can see that CoH outperforms the one without LR under all evaluation metrics on three datasets, which illustrates the effectiveness of achieving the step-by-step reasoning mechanism with LLMs.

Analysis of the score ranking procedure. In the last step of CoH reasoning, we instruct LLMs to output possible answers in order based on their possibilities of being correct. To verify whether the output index of each answer is related to its correctness, we shuffle the index order of the answers, which is denoted as “CoH w/o IS” in Table 3. Compared the results of CoH w/o IS with CoH, we can see that the shuffled index order leads to a huge drop in performance. This indicates that the indexes outputted by LLMs can be helpful for score ranking on TKG prediction.

4.4 Case Study

In this section, we visualize the reasoning process of two queries to understand the reasoning mechanism of CoH. Moreover, to further understand how CoH benefits the graph-based TKG prediction models, we show the difference between the results predicted by CoH and those inferred by the graph-

Model	Query 1: Military_(Myanmar) Express_intent_to_meet_or_negotiate to whom on the 351th day? (GT: Thailand)
CoH	<p>[Military_(Myanmar) Fight_with_small_arms_and_light_weapons Myanmar 338] → [Myanmar Make_a_visit to Thailand 328]</p> <p>[Military_(Myanmar) Charge_with_legal_action Media_Personnel_(Myanmar) 305] → [Media_Personnel_(Myanmar) Make_statement Detainee_(Myanmar) 302]</p> <p>[Military_(Myanmar) Make_an_appeal_or_request Citizen_(Thailand) 272] → [Citizen_(Thailand) Use_violence to Thailand 271]</p> <p>[Military_(Myanmar) Make_an_appeal_or_request Citizen_(Thailand) 271] → [Citizen_(Thailand) Release_person(s) Activist_(Thailand) 267]</p> <p>[Military_(Myanmar) Praise_or_endorse Military_(Thailand) 185]</p> <p>[Military_(Myanmar) Use_military_force to Rebel_Group_(Myanmar) 174]</p> <p>... ..</p> <p>[Military_(Myanmar) Use_military_force to National_Liberation_Army 62] → [National_Liberation_Army Make_statement to Guerrilla_(Colombia) 58]</p> <p>Answer: 1. Thailand 2. Citizen_Thailand 3. Activist_Thailand 4. National_Liberation_Army 5. Media_Personnel_Myanmar</p>
Graph	Answer: 1. Myanmar 2. Malaysia 3. Cambodia 4. Citizen_(Thailand) 5. Thailand
Model	Query 2: Saudi_Army Use_military_force to whom on the 3744th day? (GT: Armed_Rebel_(Yemen))
CoH	<p>[Saudi_Army Make_an_appeal_or_request Zillur_Rahman 2581] → [Zillur_Rahman Make_empathetic_comment Citizen_(North_Korea) 2543]</p> <p>[Saudi_Army Consult Zillur_Rahman 2581] → [Zillur_Rahman Make_empathetic_comment citizen_(Nerth_Kerea) 2543]</p> <p>[Saudi_Army Express_intent_to_meet_or_negotiate with Thailand 2581] → [Thailand Use_military_force Military_(Cambodia) 2580]</p> <p>[Saudi_Army Investigate Armed_Gang_(Saudi_Arabia) 1798] → [Armed_Gang_(Saudi_Arabia) Use_unconventional_violence Citizen_(Saudi_Arabia) 1773]</p> <p>[Saudi_Army Employ_aerial_weapons Yemen 1769] → [Yemen Charge_with_legal_action Armed_Rebel_(Yemen) 1764]</p> <p>Answer: 1. Military_(Cambodia) 2. Armed_Rebel_(Yemen) 3. Armed_Gang_(Saudi_Arabia) 4. Citizen_(Saudi_Arabia) 5. Citizen_(North_Korea)</p>
Graph	Answer: 1. Yemen 2. Citizen_(Saudi_Arabia) 3. Saudi_Arabian_Defence_Forces 4. Police_(Saudi_Arabia) 5. Armed_Rebel_(Yemen)

Table 4: Case studies with two queries for showing the procedure of CoH reasoning. The bold histories denote the inferred first-order histories by LLMs in Step 1, and the histories in blue color denote corresponding second-order histories. These two kinds of histories consist of second-order history chains for LLMs to infer answers in Step 2. And the answers in red color represent the ground truth of the given query.

based model RE-NET (Jin et al., 2020). Details of the two cases are shown in Table 4.

From the reasoning process of the two cases, we can see that LLMs possess the capability of inferring important histories related to the given query. And in the way of CoH reasoning, LLMs can accurately infer the answers like “Thailand” in the second-order history chains. Furthermore, comparing the answers of CoH with the graph-based model, we can see that CoH can infer the correct answer more accurately in some scenarios. We analyze the potential reason is that the semantic reasoning capability can allow LLMs to identify crucial historical information.

For example, from the numerous histories in query 1, LLMs can more precisely infer facts involved with relations like “Fight with small arms”, “Charge with legal action”, or “Make an appeal” that are more likely to result in the occurrence of “Express intent to meet or negotiate” based on semantic comprehension. For query 2, understanding the semantic meaning differences and correlations between entity “Yemen” and “Armed_Rebel_(Yemen)” is a very crucial clue. However, the two entities are just regarded as two different IDs in graph-based models, and available histories for query 2 are too limited for them to aggregate abundant information. In this case, the semantic reasoning characteristics of CoH may be more advantageous. Consequently, LLMs exhibit semantic comprehension capability, while the graph-based models possess a powerful ability to capture structural information. Thus, the reasoning

capability of CoH based on LLMs may potentially complement the reasoning performed by graph-based models in certain scenarios.

4.5 Analysis of Explainability

In this section, we aim to explore the explainability of the reasoning procedure of CoH. To achieve this, we instruct the LLM to explain its inferred answers. An example is shown in Table 5, from which we can observe that the LLM possesses the ability to capture relationships between multi-hop histories. For the inferred answer “John_Kerry”, we can see from the given explanation that the LLM captures the semantic information of the high-order history chain “Federica_Mogherini → Mohammad_Javad_Zarif → John_Kerry”. That is, leveraging the LLM’s robust semantic understanding, it can provide predictions that are reasonably interpretable to a certain extent.

4.6 Analysis of Data Leakage

Due to the absence of publicly disclosed temporal horizons for the pre-trained data of Mixtral-8x7B (Jiang et al., 2024), we consider whether this LLM may have covered some data within ICEWS14, ICEWS18, and ICEWS05-15. To verify this problem, we have a conversation with Mixtral-8x7B to check whether it knows about the facts within the three datasets following (Shi et al., 2023). The ratio of known facts in each dataset is shown in Table 6, which indicates that Mixtral-8x7B only knows a very limited subset of facts. And the conversations are presented in Appendix E. To avoid the problem

Query	Federica_Mogherini Express_intent_to_meet_or_negotiate with whom on the 336th day? (GT: John_Kerry)
History	<p>Federica_Mogherini Make_an_appeal_or_request Iran on the 324th day]</p> <p>[Iran Consult Representatives_(United_States) on the 323 day]</p> <p>[Federica_Mogherini Discuss_by_telephone Mohammad_Javad_Zarif on the 289 day]</p> <p>[Mohammad_Javad_Zarif Consult John_Kerry on the 288 day]</p> <p>[Federica_Mogherini Express_intent_to_engage_in_diplomatic_cooperation_(such_as_policy_support) Fumio_Kishida on the 265 day]</p> <p>[Federica_Mogherini Consult Fumio_Kishida on the 265 day]</p> <p>[Fumio_Kishida Express_intent_to_meet_or_negotiate John_Kerry on the 264 day]</p> <p>... ..</p>
Answer	1. Iran 2. John_Kerry 3. Fumio_Kishida ...
Explanation	<p>1. Iran: Federica_Mogherini has previously expressed intent to meet or negotiate with Iran on the 324th day. It is possible that Federica_Mogherini may do so again on the 336th day.</p> <p>2. John_Kerry: Federica_Mogherini Discuss_by_telephone with Mohammad_Javad_Zarif on the 289th day. While Mohammad_Javad_Zarif have consulted with John_Kerry on the 288th day. It is possible that Federica_Mogherini may express intent to meet or negotiate with John_Kerry on the 336th day.</p> <p>3. Fumio_Kishida: Federica_Mogherini has expressed intent to engage in diplomatic cooperation with and consulted Fumio_Kishida on the 265th day. Federica_Mogherini may express intent to meet or negotiate with Fumio_Kishida on the 336th day.</p>

Table 5: An example to show the explainability of the reasoning procedure of CoH. The explanation of each answer is given by the LLM itself, detailing how the model inferred that answer.

Datasets	ICEWS14	ICEWS05-15	ICEWS18
Ratio of known facts	3.89%	5.67%	5.87%

Table 6: Known facts checking results.

Model	ICEWS14		ICEWS18		ICEWS05-15	
	Hit@1	Hit@10	Hit@1	Hit@10	Hit@1	Hit@10
Anon-CoH	17.00	48.18	13.81	38.39	23.21	51.62
CoH	24.2	51.21	16.81	42.68	25.61	59.58

Table 7: Anonymization experimental results.

of data leakage, we exclude this subset of known facts from the testing set for CoH reasoning.

4.7 Analysis on the effect of Prior Knowledge within LLMs

In this section, we conduct experiments to investigate how the prior knowledge within LLMs affects the performance of CoH on TKG prediction. Specifically, we anonymize the TKG data by representing each entity and relation with numerical IDs. The anonymized results of CoH are presented in Table 7, which is denoted as Anon-CoH. From Table 7 we can see that the anonymization leads to a certain decline in the performance of CoH, which indicates that the prior knowledge provides a certain degree of assistance for reasoning. Note that since the possible leakage data is filtered (Section 4.6), we analyze that the usefulness of prior knowledge is primarily attributed to certain static semantic knowledge.

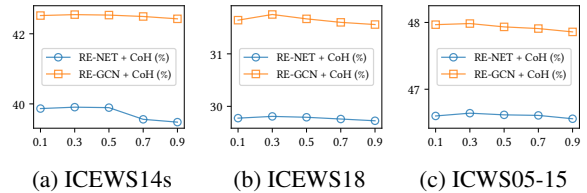


Figure 4: Performance of graph-based models plugged with CoH under different α -values in terms of MRR (%). The x-axis denotes different α -values, and the y-axis shows MRR (%) values.

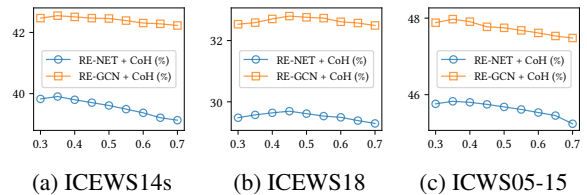


Figure 5: Performance of graph-based models plugged with CoH under different w -values in terms of MRR (%). The x-axis denotes different w -values, and the y-axis shows MRR (%) values.

4.8 Sensitivity Analysis

For converting the indexes outputted by LLM into corresponding scores, α determines the score gap for different indexes. we conduct two graph-based models plugged with CoH when α is in the range of $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. The results are shown in Figure 4, from which we can see that variations in the value of α within a very narrow range have a minimal impact on the model’s performance.

Moreover, w determines the score weight for fusing the predicted results of graph-based models and CoH. We conduct two graph-based models plugged with CoH under various w -values. The results are shown in Figure 5, from which we can see that in the optimal outcomes, the results of CoH contribute slightly more to the final score. We analyze the underlying reason leading to the aforementioned observations may be related to the score distribution of the graph-based models.

5 Related works

Temporal Knowledge Graph Forecasting with Supervised Models. As a carrier of real-world events extracted from news and documents (Liu et al., 2023; Sun et al., 2023; Liu et al., 2024), TKGs hold significant practical research value. Previous classic methods include GHNN (Han et al., 2020) and Know-Evolve (Trivedi et al., 2017), which model the temporal information within TKGs by temporal point process (TTP). And CyGNet (Zhu et al., 2021) proposes a copy-generation mechanism to explore patterns among repetitive histories. As GNNs have shown promise in sequential modeling (Liu et al., 2017), recently, most supervised models (Jin et al., 2020; Li et al., 2021b, 2022; Zhang et al., 2023b; Liang et al., 2023; Zhang et al., 2023a) for TKG prediction adopt GNNs to capture the structural information within TKGs. In particular, Zhang et al. consider both long-term (Wang et al., 2024) and short-term information. Based on these, TANGO (Han et al., 2021b) employs Neural Ordinary Differential Equations to build up continuous temporal information, CENET (Xu et al., 2023b) adopts contrastive learning to identify important non-historical entities, MetaTKG (Xia et al., 2022) and MetaTKG++ (Xia et al., 2024) explore the evolution patterns of events with meta-learning, and xERTE (Han et al., 2021a) proposes an explainable model by searching sub-graph in TKGs. Besides, some works (Sun et al., 2021; Li et al., 2021a) search significant paths with reinforcement learning, and Tlogic (Liu et al., 2022) extracts paths via temporal logic rules for TKG prediction.

Temporal Knowledge Graph with Large Language Models (LLMs). Recently, several works (Han et al., 2022; Gao et al., 2023; Xu et al., 2023a) have attempted to leverage Pre-trained Language Models (PLMs) on TKG reasoning, which mainly input histories in textual form into PLMs to obtain

contextualized knowledge embeddings. Nowadays, with the surge of LLMs, their reasoning capabilities on structural and temporal data are progressively under exploration (Jiang et al., 2023; Jain et al., 2023; Yuan et al., 2023; Aghzal et al., 2023; Wang and Zhao, 2023; Tan et al., 2023; Xiong et al., 2024). In the area of TKGs, Ding et al. input relations in textual form into LLMs to generate corresponding descriptions, which are then introduced into embed-based models as a supplement to the semantic information of zero-shot relations (Ding et al., 2023). And Lee et al. take the first attempt on TKG prediction using LLMs, which is the most closely related work to our paper. They convert TKG prediction into an In-context Learning (ICL) problem, providing LLMs with the first-order histories of the query in textual form to predict the possible answers (Lee et al., 2023).

6 Conclusion

In this paper, we first analyze the shortcomings and challenges of the existing LLM-based model about how to effectively provide comprehensive high-order historical information for LLM. Then we point out that relying solely on the reasoning capability of LLMs is still limited for TKG prediction. To resolve these issues, we propose CoH reasoning which achieves effective utilization of high-order histories for LLM. And we design CoH as plug-and-play, serving to complement and enhance the performance of graph-based models. Extensive experimental results demonstrate the superiority of CoH, and its effectiveness in enhancing the performance of graph-based models for TKG prediction.

7 Limitations

Since CoH reasoning is conducted in multi-step, LLMs need to be invoked multiple times, resulting in an increased complexity of the inference process. Moreover, we design CoH as plug-and-play, fusing its answers with predicted results of graph-based models. As this entire process does not involve any training, the fusion weight can only be controlled by the hyper-parameter w , making it impossible to achieve adaptive fusion which can automatically learn the weight allocation of the scores obtained from CoH for different queries. Moving forward, how to design an adaptive fusion strategy that optimally enhances the performance of graph-based models with results of CoH without compromising efficiency is worth exploring.

Acknowledgements

This work is sponsored by the National Key Research and Development Program (2023YFC3305203), and the National Natural Science Foundation of China (NSFC) (Grant 62376265, 62141608, 62206291).

References

- Mohamed Aghzal, Erion Plaku, and Ziyu Yao. 2023. Can large language models be good path planners? a benchmark and investigation on spatial-temporal reasoning. *arXiv preprint arXiv:2310.03249*.
- Elizabeth Boschee, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz, and Michael Ward. 2015. Icews coded event data. *Harvard Data-verse*, 12.
- Ziyang Chen, Jinzhi Liao, and Xiang Zhao. 2023. Multi-granularity temporal question answering over knowledge graphs. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11378–11392.
- Zifeng Ding, Heling Cai, Jingpei Wu, Yunpu Ma, Ruotong Liao, Bo Xiong, and Volker Tresp. 2023. Zero-shot relational learning on temporal knowledge graphs with large language models. *arXiv preprint arXiv:2311.10112*.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Yifu Gao, Yongquan He, Zhigang Kan, Yi Han, Linbo Qiao, and Dongsheng Li. 2023. Learning joint structural and temporal contextualized knowledge embeddings for temporal knowledge graph completion. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 417–430.
- Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. 2018. Learning sequence encoders for temporal knowledge graph completion. *arXiv preprint arXiv:1809.03202*.
- Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. 2021a. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *ICLR*.
- Zhen Han, Zifeng Ding, Yunpu Ma, Yujia Gu, and Volker Tresp. 2021b. Learning neural ordinary equations for forecasting future links on temporal knowledge graphs. In *EMNLP*, pages 8352–8364.
- Zhen Han, Ruotong Liao, Beiyan Liu, Yao Zhang, Zifeng Ding, Jindong Gu, Heinz Koepl, Hinrich Schuetze, and Volker Tresp. 2022. Enhanced temporal knowledge embeddings with contextualized language representations.
- Zhen Han, Yunpu Ma, Yuyi Wang, Stephan Günnemann, and Volker Tresp. 2020. Graph hawkes neural network for forecasting on temporal knowledge graphs. In *AKBC*.
- Raghav Jain, Daivik Sojitra, Arkadeep Acharya, Sriparna Saha, Adam Jatowt, and Sandipan Dandapat. 2023. Do language models have a common sense regarding time? revisiting temporal commonsense reasoning in the era of large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6750–6774.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Structgpt: A general framework for large language model to reason over structured data. *arXiv e-prints*, pages arXiv–2305.
- W. Jin, M. Qu, X. Jin, and X. Ren. 2020. Recurrent event network: Autoregressive structure inference over temporal knowledge graphs. In *EMNLP*, pages 6669–6683.
- Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. 2019. Recurrent event network: Autoregressive structure inference over temporal knowledge graphs. *arXiv preprint arXiv:1904.05530*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Dong-Ho Lee, Kian Ahrabian, Woojeong Jin, Fred Morstatter, and Jay Pujara. 2023. Temporal knowledge graph forecasting without knowledge using in-context learning. *arXiv preprint arXiv:2305.10613*.
- Yujia Li, Shiliang Sun, and Jing Zhao. 2022. Tirgn: Time-guided recurrent graph network with local-global historical patterns for temporal knowledge graph reasoning. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2152–2158.
- 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 and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4732–4743.

- Zixuan Li, Xiaolong Jin, Wei Li, Saiping Guan, Jiafeng Guo, Huawei Shen, Yuanzhuo Wang, and Xueqi Cheng. 2021b. Temporal knowledge graph reasoning based on evolutionary representation learning. In *SIGIR*, pages 408–417.
- Ke Liang, Lingyuan Meng, Meng Liu, Yue Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, and Xinwang Liu. 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*, pages 1559–1568.
- Qiang Liu, Shu Wu, and Liang Wang. 2017. Multi-behavioral sequential prediction with recurrent log-bilinear model. *IEEE Transactions on Knowledge and Data Engineering*, 29(6):1254–1267.
- Wanlong Liu, Shaohuan Cheng, Dingyi Zeng, and Hong Qu. 2023. Enhancing document-level event argument extraction with contextual clues and role relevance. *arXiv preprint arXiv:2310.05991*.
- Wanlong Liu, Li Zhou, Dingyi Zeng, Yichen Xiao, Shaohuan Cheng, Chen Zhang, Grandee Lee, Malu Zhang, and Wenyu Chen. 2024. Beyond single-event extraction: Towards efficient document-level multi-event argument extraction. *arXiv preprint arXiv:2405.01884*.
- 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 AAAI conference on artificial intelligence*, volume 36, pages 4120–4127.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Xiaoming Shi, Siqiao Xue, Kangrui Wang, Fan Zhou, James Y Zhang, Jun Zhou, Chenhao Tan, and Hongyuan Mei. 2023. Language models can improve event prediction by few-shot abductive reasoning. *arXiv preprint arXiv:2305.16646*.
- Haohai Sun, Jialun Zhong, Yunpu Ma, Zhen Han, and Kun He. 2021. TimeTraveler: Reinforcement learning for temporal knowledge graph forecasting. In *EMNLP*, pages 8306–8319.
- Xin Sun, Qiang Liu, Shu Wu, Zilei Wang, and Liang Wang. 2023. Noise-robust semi-supervised learning for distantly supervised relation extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13145–13157.
- Qingyu Tan, Hwee Tou Ng, and Lidong Bing. 2023. Towards benchmarking and improving the temporal reasoning capability of large language models. *arXiv preprint arXiv:2306.08952*.
- Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *ICML*, pages 3462–3471.
- Liang Wang, Hao Fu, Shu Wu, Qiang Liu, Xuelei Tan, Fangsheng Huang, Mengdi Zhang, and Wei Wu. 2024. Camlo: Cross-attentive multi-view network for long-term origin-destination flow prediction. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*, pages 454–462. SIAM.
- Yuqing Wang and Yun Zhao. 2023. Tram: Benchmarking temporal reasoning for large language models. *arXiv preprint arXiv:2310.00835*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yuwei Xia, Mengqi Zhang, Qiang Liu, Liang Wang, Shu Wu, Xiaoyu Zhang, and Liang Wang. 2024. Metatkg++: Learning evolving factor enhanced meta-knowledge for temporal knowledge graph reasoning. *Pattern Recognition*, page 110629.
- Yuwei Xia, Mengqi Zhang, Qiang Liu, Shu Wu, and Xiao-Yu Zhang. 2022. Metatkg: Learning evolutionary meta-knowledge for temporal knowledge graph reasoning. In *EMNLP*, pages 7230–7240.
- Sheng Xiang, Dawei Cheng, Chencheng Shang, Ying Zhang, and Yuqi Liang. 2022. Temporal and heterogeneous graph neural network for financial time series prediction. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 3584–3593.
- Siheng Xiong, Ali Payani, Ramana Kompella, and Faramarz Fekri. 2024. Large language models can learn temporal reasoning. *arXiv preprint arXiv:2401.06853*.
- Wenjie Xu, Ben Liu, Miao Peng, Xu Jia, and Min Peng. 2023a. Pre-trained language model with prompts for temporal knowledge graph completion. *arXiv preprint arXiv:2305.07912*.
- Yi Xu, Junjie Ou, Hui Xu, and Luoyi Fu. 2023b. Temporal knowledge graph reasoning with historical contrastive learning. In *AAAI*.
- Chenhan Yuan, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2023. Back to the future: Towards explainable temporal reasoning with large language models. *arXiv preprint arXiv:2310.01074*.
- Mengqi Zhang, Yuwei Xia, Qiang Liu, Shu Wu, and Liang Wang. 2023a. Learning latent relations for temporal knowledge graph reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12617–12631.

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

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 *AAAI*, pages 4732–4740.

A Prompt Examples for CoH Reasoning

Taking the query (*Government_(Nigeria), Make_an_appeal_or_request, ?, 340*) as an example, whose ground-truth is “Member_of_the_Judiciary_(Nigeria)”. The prompts of CoH reasoning are shown as follows.

Prompt for Step 1

There is a given text consisting of multiple historical events in the form of “{id}:{subject} {relation} {object} {time}”;. And there is a query in the form of: “{subject} {relation} {whom} time?” If you must infer several {object} that you think may be the answer to the given query based on the given historical events, what important historical events do you base your predictions on? Please list the top 30 most important histories and output their {id}.

Here are the given historical events:

- 0:[Government_(Nigeria) Engage_in_diplomatic_cooperation with Independent_Corrupt_Practices_Commission on the 339th day];
- 1:[Government_(Nigeria) Threaten Education_(Nigeria) on the 338th day];
- 2:[Government_(Nigeria) Criticize_or_denounce Boko_Haram on the 337th day];
- 3:[Government_(Nigeria) Threaten Education_(Nigeria) on the 337th day];
- 4:[Government_(Nigeria) Provide_aid for Citizen_(Nigeria) on the 337th day];
- 5:[Government_(Nigeria) Make_optimistic_comment on Citizen_(Nigeria) on the 336th day];
- 6:[Government_(Nigeria) Use_conventional_military_force to Boko_Haram on the 335th day];
- 7:[Government_(Nigeria) Make_statement to Barack_Obama on the 333rd day];
- 8:[Government_(Nigeria) Accuse Boko_Haram on the 332nd day];
- 9:[Government_(Nigeria) Express_intent_to_engage_in_diplomatic_cooperation to Citizen_(Nigeria) on the 332nd day];
- 10:[Government_(Nigeria) Engage_in_diplomatic_cooperation with Citizen_(Nigeria) on the 331st day];
-

Here is the query:

Government_(Nigeria), Make_an_appeal_or_request to, whom, on the 340th day?

Please only output {id} of the historical events that your inferred answers are based on. Note that you must only output no more than 30 {id} without any explanation. Please strictly follow the above demands for output.

Answer:

0, 1, 2, 5, 7, 8, 9, 10, 11, 14, 17, 21, 24, 25, 26, 27, 28, 30, 31, 32, 33, 34, 36, 37, 38, 39, 45, 49, 50, 56

Prompt for Step 2

You must be able to correctly predict the {whom} of the given query from a given text consisting of multiple historical events in the form of “{subject} {relation} {object} {time}” and the query in the form of “{subject} {relation} {whom} {time}?” You must output several {object} that you think may be the answer to the given query based on the given historical events. Please list all possible {object} which may be answers to the query. Please assign each answer a serial number to represent its probability of being the correct answer. Note that answers with a high probability of being correct should be listed first.

Here are the given historical events:

Government_(Nigeria), Engage_in_diplomatic_cooperation with, Independent_Corrupt_Practices_Commission, on the 339th day; Independent_Corrupt_Practices_Commission, Arrest_or_detain_or_charge_with_legal_action to, Citizen_(Nigeria), on the 308th day;

Government_(Nigeria), Criticize_or_denounce, Boko_Haram, on the 337th day; Boko_Haram, Use_conventional_military_force to, Citizen_(Nigeria), on the 336th day;

Government_(Nigeria), Threaten, Education_(Nigeria), on the 337th day; Education_(Nigeria), Make_statement to, Muslim_(Nigeria), on the 332nd day;

Government_(Nigeria), Make_optimistic_comment on, Citizen_(Nigeria), on the 336th day; Citizen_(Nigeria), Make_an_appeal_or_request to, Member_of_the_Judiciary_(Nigeria), on the 331st day;

... ..

Here is the query:

Government_(Nigeria), Make_an_appeal_or_request to, whom, on the 340th day?

Please list all possible {object} which may be answers (one per line) without explanations. Note that answers with high probability should be listed first.

For example:

""

Possible answers:

1. XXX

2. XXX

3. XXX

... ..

""

Please strictly follow the above demands for output.

Answer:

1. Citizen_(Nigeria)

2. Education_(Nigeria)

3. Member_of_the_Judiciary_(Nigeria) 4. Barack_Obama

5. Xi_Jinping

6. Boko_Haram

7. Head_of_Government_(Nigeria)

8. Court_Judge_(Nigeria)

Datasets	ICEWS14	ICEWS05-15	ICEWS18
# \mathcal{E}	6,869	10,094	23,033
# \mathcal{R}	230	251	256
# Train	74,845	368,868	373,018
# Valid	8,514	46,302	45,995
# Test	7,371	46,159	49,545
Time gap	24 hours	24 hours	24 hours

Table 8: The statistics of the datasets.

B Datasets

In this paper, we utilize three representative TKG datasets for experimental analysis: ICEWS14 (García-Durán et al., 2018), ICEWS18 (Jin et al., 2019), and ICEWS05-15 (García-Durán et al., 2018), all sourced from the Integrated Crisis Early Warning System (Boschee et al., 2015), documenting events in 2014, 2018, and from 2005 to 2015, respectively. Detailed statistics of the three datasets are shown in Table 8.

C Backbones

Since CoH is plug-and-play, we plug it into several following state-of-the-art TKG reasoning models to evaluate the effectiveness of our proposed model.

- *RE-NET* (Jin et al., 2020) deals with TKGs as KG sequences. RE-NET utilizes the RGCN to capture the structural dependencies of entities and relations within each KG. Then RNN is adopted to associate KGs with different time stamps for capturing the temporal dependencies of entities and relations.
- *RE-GCN* (Li et al., 2021b) proposes a recurrent evolution module based on relational GNNs to obtain embeddings that contain dynamic information for entities and relations. In particular, RE-GCN designs a static module that utilizes the static properties of entities to enrich the embeddings for prediction.
- *TiRGN* (Li et al., 2022) utilizes a recurrent graph encoder to capture local temporal dependencies, and designs a history encoder network to capture global temporal dependencies by collecting repeated facts in history. TiRGN fuses the obtained local and global temporal dependencies for final prediction.

D Implementation Details

In this study, we developed CoH using the PyTorch framework (Paszke et al., 2019), integrat-

ing functionalities from the HuggingFace’s Transformers library (Wolf et al., 2019) and the efficient LLM inference framework vLLM (Kwon et al., 2023). All experiments were executed on an NVIDIA A100 GPU with AMD EPYC 7763 CPU processor. Our experimental setup primarily utilized the *TheBloke/Mixtral-8x7B-Instruct-v0.1-GPTQ* model with a parameter size of 6.07B. This model represents a quantized variant of the *Mixtral-8x7B MoE* model, employing the GPTQ technique (Frantar et al., 2022) to achieve a balanced compromise between computational speed and model performance. Regarding the generation hyper-parameters, we set the Maximum tokens as 8000, Top-p sampling as 1, and Temperature as 0.

Moreover, the three graph-based TKG prediction models are all implemented with the most optimal hyper-parameters reported in their corresponding papers. Especially, for evaluating CoH on reversed testing sets like graph-based models, we manually process relations in datasets into reversed ones. For example, we reverse “Express intent to cooperate” into “Receive intent to cooperate”, and “threaten” into “be threatened”. For the score fusion module, the hyper-parameter α is set to 0.3 for all datasets. And the weights w of ICEWS14s, ICWES18, and ICEWS05-15 are set to 0.35, 0.45, and 0.35, respectively.

E Prompt Examples for Data Leakage Analysis

To avoid the possible data leakage issue, we directly ask the LLM Mixtral-8x7B regarding the facts contained in our used datasets, to check whether it knows the knowledge. And we remove all queries whose answer is “Yes” from the testing set for all three datasets. A few prompt examples are shown as follows.

Prompt for Example 1
Do you know the fact that United Arab Emirates reduced or broke diplomatic relations with Qatar on 2014-12-04?

Answer No.

Prompt for Example 2

Do you know the fact that Police (Egypt) used tactics of violent repression against Protesters (Egypt) on 2014-12-02?

Answer Yes.

Prompt for Example 3

Do you know the fact that Abdullah Abdullah met at a 'third' location with Jens Stoltenberg on 2014-12-02?

Answer: No.