

Question answering over temporal knowledge graphs: a self-improving hierarchical retrieval-augmented approach

Ruishen LIU^{1†}, Shaorong XIE^{1,2†}, Xiangfeng LUO¹, Xinzhi WANG^{1*} & Hang YU^{1*}¹*School of Computer Engineering and Science, Shanghai University, Shanghai 200444, China*²*Engineering Research Center of Unmanned Intelligent Marine Equipment, Ministry of Education, Shanghai 200444, China*

Received 18 August 2025/Revised 24 December 2025/Accepted 4 February 2026/Published online 10 June 2026

Citation Liu R S, Xie S R, Luo X F, et al. Question answering over temporal knowledge graphs: a self-improving hierarchical retrieval-augmented approach. *Sci China Inf Sci*, 2026, 69(9): 199101, <https://doi.org/10.1007/s11432-025-4858-2>

Reasoning over dynamic, time-evolving information is a fundamental challenge for artificial intelligence. Temporal knowledge graph question answering (TKGQA) addresses this challenge by answering temporal questions with dynamic intent over the temporal knowledge graph (TKG). Different from standard knowledge graph question answering, answers to the temporal questions are based on temporal constraints. Compared with common questions, temporal questions have more complex semantics requiring multiple facts. Early embedding-based methods, such as CronKGQA [1], extend static techniques to the temporal domain but often struggle with complex semantic structures. Recently, the integration of large language models (LLMs) with TKGs has shown promise. Methods like TimeR⁴ [2] employ a “retrieve-rewrite” strategy to handle temporal queries. However, these approaches face two critical challenges. First, they struggle to distinguish the implicit temporal constraints embedded within complex queries. Temporal questions often possess a hierarchical structure where temporal conditions act as implicit constraints for the target answer. Mapping the entire question into a flat embedding mixes these constraints with entity information, making it difficult for the model to distinguish between the reference event and the target inquiry. Second, they suffer from a loss of intermediate information. The key to temporal knowledge graph question answering is to fetch the relevant temporal facts based on the temporal questions. Therefore, it is essential to employ an accurate retriever. However, the current dataset only contains questions and answers without the intermediate temporal facts, resulting in ineffective training of fact retrievers.

To solve the above issues, we propose SHRA, a self-improving hierarchical retrieval-augmented framework. Firstly, the complex questions are decomposed into subquestions reflecting the hierarchical temporal semantics. Subsequently, for each subquestion, a temporal-fact retriever is utilized to fetch the relevant facts from the TKG. Since generating the reasoning process can improve the LLM reasoning capability [3], the subquestions and the facts are processed by a closed-source LLM to generate the reasoning process, forming an augmented dataset. Afterward, the augmented dataset is used to fine-tune an open-source LLM through supervised fine-tuning (SFT). Due to the lack of intermediate tem-

poral facts used to train the retriever, we further propose self-improvement strategies to improve the retriever accuracy and reasoning performance. (1) The temporal-fact retriever is refined via a contrastive learning objective, using automatically generated temporally sensitive negative samples. (2) The LLM’s reasoning capabilities are optimized using group relative policy optimization (GRPO), a reinforcement learning strategy that rewards higher accuracy and adjusts the reasoning process. The overall framework is shown in Figure 1, which consists of five components: question understanding, temporal fact retrieval, data augmentation, LLM fine-tuning, and self-improvement.

Question understanding. Complex temporal questions often contain nested entities and constraints. We employ a classification module to determine whether a question is complex. If so, a closed-source LLM decomposes the question q into a set of subquestions q' to reflect hierarchical temporal semantics. This decomposition allows the model to focus on distinct information needs, such as separating “When did Antonis Samaras praise the Prime Minister of Germany?” and “Who was the first person to be praised by Antonis Samaras after subquestion1?” from “Who was the first person to be praised by Antonis Samaras after the Prime Minister of Germany?”.

Temporal fact retrieval. After understanding the questions, we need to retrieve the temporal facts related to them. A significant challenge in TKGs is data redundancy. The same core fact can appear multiple times with different timestamps, leading to the retrieval of more similar information. To handle this efficiently, we preprocess the TKG into a structured temporal knowledge base (TKB). In the TKB, all temporal instances of a core fact are consolidated into a single group. Subsequently, we employ a fine-tuned temporal-fact retriever LM_t to encode questions and TKB facts. The retriever identifies the top- k relevant temporal fact groups based on cosine similarity between the processed question embedding $E_{q'}$ and fact embedding E_f .

Data augmentation and fine-tuning. We utilize the closed-source LLM to perform data augmentation, generating explicit reasoning paths p for the training set. We filter these based on ground truth (y) alignment. The open-source LLM is then trained via SFT on the augmented dataset $d = (q, q', f, p, y^*, y)$, learning

* Corresponding author (email: wxz2017@shu.edu.cn, yuhang@shu.edu.cn)

† These authors contributed equally to this work.

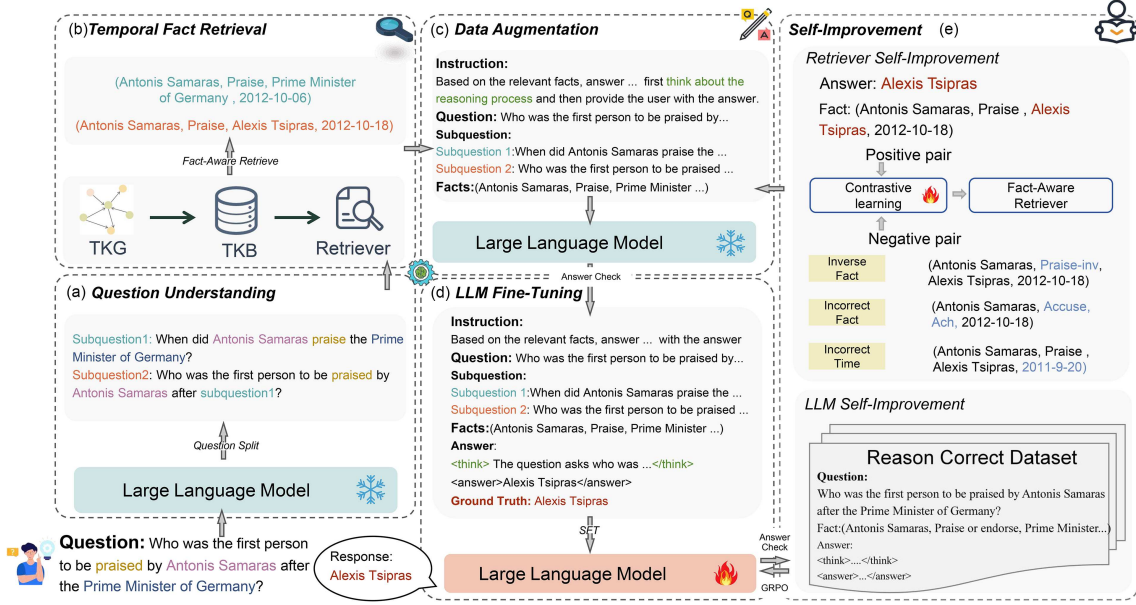


Figure 1 (Color online) The overall framework of the proposed model SHRA. The model includes five components: question understanding, temporal fact retrieval, data augmentation, LLM fine-tuning and self-improvement. The right part illustrates two self-improvement strategies: retriever self-improvement and LLM self-improvement.

to generate the reasoning process p and answer y^* given the question and facts.

Self-improvement strategies. To further improve the reasoning performance of the proposed model, SHRA incorporates a dual self-improvement loop. This mechanism progressively enhances the two most critical components of the framework: the temporal-fact retriever precision and the LLM reasoning capabilities.

Retriever self-improvement. We generate weak supervision signals from the final answer. The closed-source LLM generates a reasoning process and an answer based on retrieved facts. If the generated answer matches the ground truth, the retrieved temporal facts used are designated as positive pairs. As shown in Figure 1, we then randomly corrupt relations, entities, or timestamps of these correct facts to generate negative pairs. The retriever LM_t is refined via a contrastive learning objective [4]:

$$\mathcal{L} = \sum_i [w_p Y \cdot \exp(\phi_{TKB}) + w_n (1 - Y) \cdot \exp(1 - \phi_{TKB})], \quad (1)$$

where ϕ_{TKB} is the cosine similarity, and $Y = 1$ for positive pairs and 0 for negative pairs. This forces the retriever to distinguish subtle temporal and semantic discrepancies.

LLM self-improvement. Compared with SFT, RL can dynamically adjust the LLM reasoning process, thereby achieving superior performance in complex reasoning tasks. Therefore, we further enhance the model performance through reinforcement learning. In the initial stage, the SFT model generates the reasoning path and a final answer for each question in the training set. Each generated answer is then validated against the ground truth. If an answer is correct, the corresponding data is added to the correct dataset. Then the model is further trained on the correct dataset by using reinforcement learning. We employ GRPO [5] as the reinforcement learning algorithm. This approach enables the model to refine its own reasoning policy by learning from its successful problem-solving attempts.

Experiment. We conduct experiments on two benchmark datasets. The proposed SHRA method is evaluated against various TKGQA baselines, including traditional embedding-based models

and recent LLM-based approaches. The results show that SHRA achieves state-of-the-art performance in TKGQA, particularly excelling in questions requiring multiple temporal facts. We verify the effectiveness of SHRA in handling hierarchical temporal semantics through question decomposition and demonstrate the superiority of the dual self-improvement strategies. Moreover, we conduct ablation studies and fine-grained analyses to validate the contribution of each component. Case studies provide deep insights into the reasoning process of SHRA. Furthermore, error analysis points out the limitations of the proposed model. More details are shown in Appendix D.

Conclusion. We propose SHRA, a novel framework for TKGQA that integrates hierarchical question decomposition with retrieval-augmented generation. By introducing self-improvement strategies for the retriever and the LLM, SHRA effectively overcomes the scarcity of intermediate supervision signals. Extensive experiments demonstrate that SHRA achieves state-of-the-art performance, particularly in handling complex temporal semantics.

Acknowledgements This work was supported by National Natural Science Foundation of China (Grant No. 91746203) and Outstanding Academic Leader Project of Shanghai (Grant No. 20XD1401700).

Supporting information Appendixes A–D. The supporting information is available online at info.scichina.com and link.springer.com. The supporting materials are published as submitted, without typesetting or editing. The responsibility for scientific accuracy and content remains entirely with the authors.

References

- Saxena A, Chakrabarti S, Talukdar P P. Question answering over temporal knowledge graphs. In: Proceedings of Annual Meeting of the Association for Computational Linguistics, 2021
- Qian X, Zhang Y, Zhao Y, et al. TimeR⁴: time-aware retrieval-augmented large language models for temporal knowledge graph question answering. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2024. 6942–6952
- DeepSeek-AI, Guo D, Yang D, et al. Deepseek-r1: incentivizing reasoning capability in LLMs via reinforcement learning. 2025. ArXiv:2501.12948
- Son J, Oh A. Time-aware representation learning for time-sensitive question answering. 2023. ArXiv:2310.12585
- Shao Z, Wang P, et al. Deepseekmath: pushing the limits of mathematical reasoning in open language models. 2024. ArXiv:2402.03300