

# Knowledge Editing with Dynamic Knowledge Graphs for Multi-Hop Question Answering

Yifan Lu<sup>1</sup>, Yigeng Zhou<sup>1</sup>, Jing Li<sup>1\*</sup>, Yequan Wang<sup>2</sup>,  
Xuebo Liu<sup>1</sup>, Daojing He<sup>1</sup>, Fangming Liu<sup>3</sup>, Min Zhang<sup>1</sup>

<sup>1</sup>Harbin Institute of Technology, Shenzhen, China

<sup>2</sup>Beijing Academy of Artificial Intelligence, Beijing, China

<sup>3</sup>Pengcheng Laboratory, Shenzhen, China

lu.yifan@foxmail.com, jingli.phd@hotmail.com

## Abstract

Multi-hop question answering (MHQA) poses a significant challenge for large language models (LLMs) due to the extensive knowledge demands involved. Knowledge editing, which aims to precisely modify the LLMs to incorporate specific knowledge without negatively impacting other unrelated knowledge, offers a potential solution for addressing MHQA challenges with LLMs. However, current solutions struggle to effectively resolve issues of knowledge conflicts. Most parameter-preserving editing methods are hindered by inaccurate retrieval and overlook secondary editing issues, which can introduce noise into the reasoning process of LLMs. In this paper, we introduce KEDKG, a novel knowledge editing method that leverages a dynamic knowledge graph for MHQA, designed to ensure the reliability of answers. KEDKG involves two primary steps: dynamic knowledge graph construction and knowledge graph augmented generation. Initially, KEDKG autonomously constructs a dynamic knowledge graph to store revised information while resolving potential knowledge conflicts. Subsequently, it employs a fine-grained retrieval strategy coupled with an entity and relation detector to enhance the accuracy of graph retrieval for LLM generation. Experimental results on benchmarks show that KEDKG surpasses previous state-of-the-art models, delivering more accurate and reliable answers in environments with dynamic information.

## Introduction

Large language models (LLMs) have gained widespread adoption due to their advanced language understanding and reasoning capabilities (Zhao et al. 2023; Huang and Chang 2023). However, as the world changes and information becomes outdated, the datasets or corpora used to pre-train LLMs may no longer be relevant, potentially leading to unreliable outcomes in certain question-and-answer applications. To circumvent the substantial time and costs associated with retraining LLMs, researchers have shifted focus to knowledge editing (Zhang et al. 2024; Wang et al. 2023). The goal of knowledge editing is to precisely modify and update the knowledge within LLMs in a cost-effective manner, thereby enabling them to deliver dependable answers.

\*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

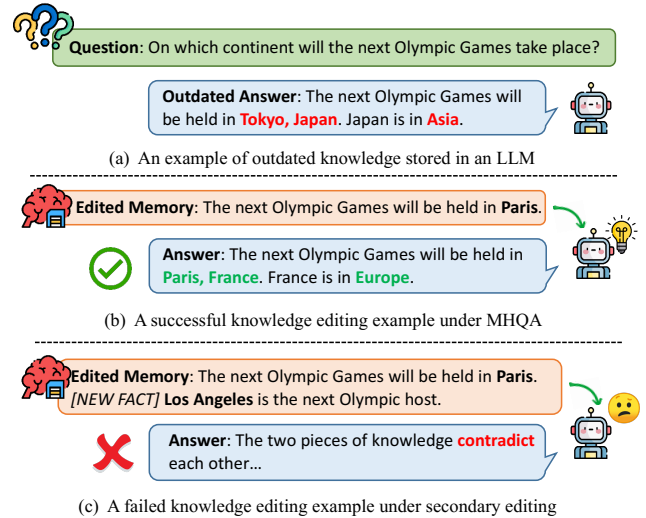


Figure 1: Examples of knowledge editing for MHQA. (a) an example of outdated information stored in LLMs. (b) a successful update with a parameter-preserving editing method. (c) a failure occurring during secondary editing.

For LLMs, multi-hop question answering (MHQA) is a challenging task that requires high-order comprehension and reasoning abilities (Mavi, Jangra, and Jatowt 2022). In particular, knowledge editing in MHQA is often vulnerable due to the **cascade effect**. For example, as shown in Figure 1.(a) with the multi-hop question: “On which continent will the next Olympic Games take place?”, if the information about the next Olympic host is updated from *Tokyo* to *Paris*, then the answer to the multi-hop question should also be updated to reflect *Paris*’ location. Currently, approaches to knowledge editing for MHQA can be categorized into two main types: (1) *Parameter-based editing methods*: These methods directly adjust the parameters of the LLM to change its output (Zhu et al. 2020; Meng et al. 2022; De Cao, Aziz, and Titov 2021; Meng et al. 2023). While effective for single-hop questions, recent studies have shown that they face challenges with MHQA, including issues like catastrophic forgetting and degradation of the original model’s performance (Gu et al. 2024b). (2) *Parameter-preserving editing*

*methods*: These methods retain the original model’s parameters and alter its behavior by adding extra modules (Mitchell et al. 2022; Huang et al. 2023; Zhong et al. 2023; Gu et al. 2024a; Cheng et al. 2024). A typical strategy involves storing the edited fact directly in memory and using retrieval-enhanced techniques to prompt the LLM to modify its responses. Figure 1.(b) illustrates a successful application of this approach in MHQA.

However, existing approaches exhibit two primary limitations. First, the retrieval process often relies on semantic similarity matching, which can be inadequate and overlook information pertinent to the original question. In multi-hop questions, only the initial entity is known, requiring the LLM to infer intermediate entities. Second, as illustrated in Figure 1.(c), edited facts may require updates over time, a process we refer to as **secondary editing**. Directly storing secondarily edited knowledge in memory can conflict with previously edited facts, leading to interference in the retrieval process and errors during the LLM’s inference phase.

To address the aforementioned limitations, we propose a novel method called KEDKG: **K**nowledge **E**ditin**G** with **D**ynamic **K**nowledge **G**raphs for MHQA. Inspired by the evolving knowledge storage and convenient modification features of knowledge graphs, KEDKG initially converts edited knowledge into structured triples to establish a knowledge graph. This knowledge graph can **dynamically** expand, modify, or delete information as the world’s knowledge evolves. Utilizing this dynamic knowledge graph, KEDKG tackles the issue of secondary editing that may lead to potential conflicts, accurately providing the necessary edited knowledge for the LLM. When handling multi-hop questions, inspired by the chain-of-thought (CoT) approach (Wei et al. 2022), KEDKG use a fine-tuned LLM to decompose a multi-hop question into multiple sub-questions. For each sub-question, KEDKG perform fine-grained retrieval within the newly created knowledge graph and use detectors to filter the results, ensuring both the speed and accuracy of retrieval. Extensive experiments demonstrate that KEDKG surpasses existing state-of-the-art knowledge editing methods. We summarize the key contributions of our work as follows:

- To the best of our knowledge, this is the first approach to automatically construct dynamic knowledge graphs specifically for knowledge editing, and to investigate the knowledge conflict issue of secondary editing.
- We propose a novel knowledge editing method, named KEDKG, for MHQA tasks that leverages knowledge graphs to resolve knowledge conflicts in secondary editing. KEDKG employs a fine-grained retrieval and filtering strategy to improve the accuracy of the retrieval process. Additionally, the method is lightweight and applicable with all open-source and black-box LLMs.
- We conduct extensive experiments across various LLMs and datasets to validate the effectiveness and usability of KEDKG. Our empirical results and analysis demonstrate that KEDKG significantly outperforms the advanced existing baselines, achieving superior performance.

## Related Work

### Parameter-Preserving Editing Methods

Parameter-preserving editing methods do not modify the original model’s parameters. Instead, they introduce additional modules (such as neurons, retrievers, etc.) to influence the model’s behavior. SERAC (Mitchell et al. 2022) introduces a scope classifier to determine if an input request requires editing and employs a counterfactual model to handle requests. T-patcher (Huang et al. 2023) and CaliNet (Dong et al. 2022) incorporate additional trainable parameters in feed-forward layers of PLMs to provide edited knowledge. IKE (Zheng et al. 2023) utilizes in-context learning based on demonstration storage to prompt LLMs to edit knowledge with predefined templates. MeLLO (Zhong et al. 2023) is a simple editing method for multi-hop QA tasks, storing all edited facts externally while iteratively prompting LLMs to generate answers consistent with the edited facts. PokeMQA (Gu et al. 2024a) improves retrieval and model answer accuracy by dividing the task into 2 steps: question decomposition and decoupled self-checking. TEMPLEMQA (Cheng et al. 2024) directly generates multi-hop reasoning paths and constructs temporal graphs to enhance support for handling temporal multi-hop QA tasks. Unlike existing solutions, KEDKG constructs a structured knowledge graph to store the edited knowledge and supports dynamic updates to the edited knowledge.

### Parameter-Based Editing Methods

Parameter-based editing methods directly modify some or all parameters to incorporate edited knowledge into the model. Current methods can be categorized into three types: fine-tuning, meta-learning, and locate-then-edit methods. Traditional fine-tuning methods (Chen et al. 2020; Zhu et al. 2020) can enable the model to learn new knowledge, but it suffers from catastrophic forgetting. Meta-learning methods use a hyperparameter network to learn to modify model weights. For example, MEND (De Cao, Aziz, and Titov 2021) learns to transform the gradient of fine-tuned LLMs through low-rank decomposition of the gradient, but this approach incurs additional training overhead. Locate-then-edit methods directly identify specific parameters within the model and modify them. For instance, ROME (Meng et al. 2022) uses causal mediation analysis to locate the editing region and modifies the feed-forward layer parameters in that region to update knowledge. However, this method does not perform well for complex multi-hop QA problems. Recent research (Gu et al. 2024b) indicates that parameter-based editing methods are likely to significantly impact the model’s original performance.

## Preliminaries

**Notations.** We define each piece of factual knowledge embedded in the LLM as a triple  $f = (s, r, o) \in F$ , where  $F$  represents the collection of all facts. Each triple  $f$  consists of a subject  $s$ , a relation  $r$  and an object  $o$ . Edited knowledge shares the same subject and relation as the original knowledge, and the object is updated from  $o$  to  $o^*$ , denoted as  $e = (s, r, o \rightarrow o^*)$ . For example, a possible edit could be (Brazil,

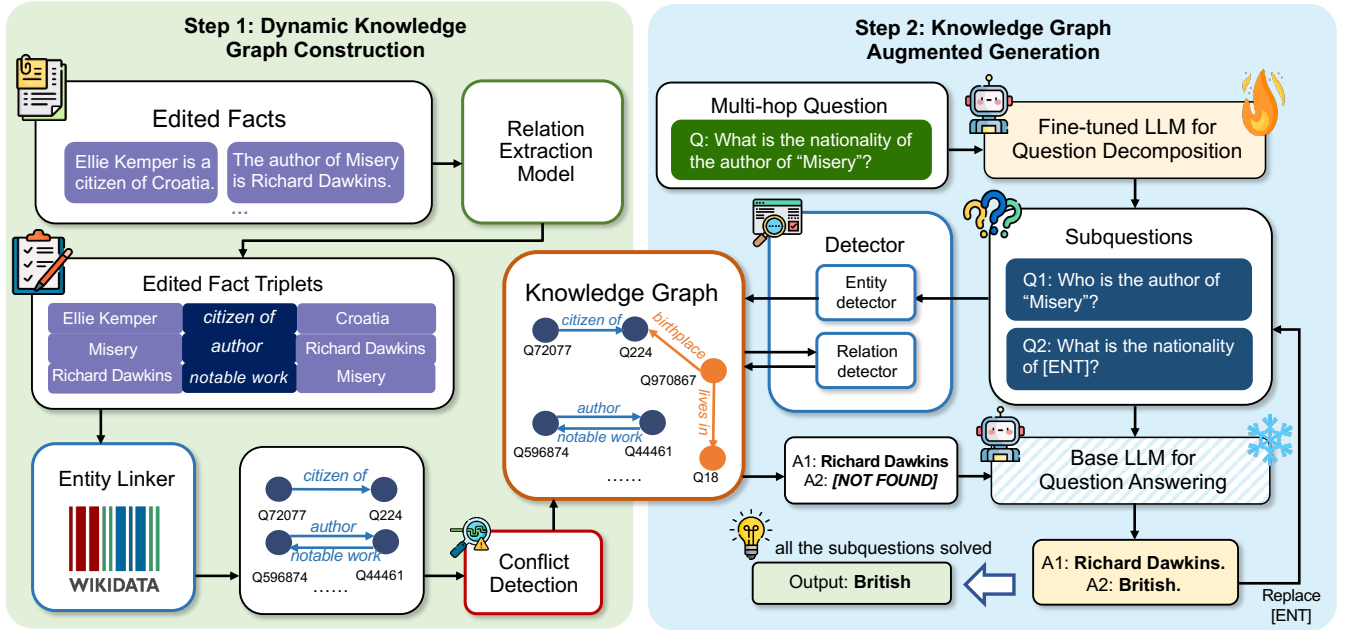


Figure 2: An overview of our proposed method KEDKG, which consists of two main steps: Dynamic Knowledge Graph Construction and Knowledge Graph Augmented Generation. KEDKG begins by dynamically building a new knowledge graph to store edited knowledge and address potential conflicts. Subsequently, KEDKG employs a fine-grained retrieval strategy, utilizing entity and relation detectors to enhance the accuracy of graph retrieval for LLM generation.

continent, South America  $\rightarrow$  Asia). A knowledge editing operation can be represented as  $E = \{e_1, e_2, \dots, e_m\}$ , potentially encompassing multiple edited triples of knowledge.

**Multi-hop QA under knowledge editing.** Given a multi-hop question  $Q$ , answering  $Q$  requires sequentially querying and retrieving multiple facts. Following the order in which these facts are retrieved, these facts can be represented a chain of facts  $C = [(s_1, r_1, o_1), \dots, (s_n, r_n, o_n)]$ , where  $s_{i+1} = o_i$  and  $o_n$  is the final answer to the question  $Q$ . If we replace old knowledge in the chain of facts with new knowledge  $(s_i, r_i, o_i^*)$ , the cascading effect inherent in multi-hop question tasks affects the entire chain. Consequently, the final reasoning chain becomes  $C^* = [(s_1, r_1, o_1), \dots, (s_i, r_i, o_i^*), \dots, (s_n, r_n, o_n^*)]$ , where  $o_n^*$  is the updated answer. It should be noted that any number of old knowledge triples can be replaced in the chain of facts. The goal of the knowledge editing task for MHQA can be outlined as follows, given an editing set  $E$  and a language model  $M$ , generate a conditionally edited language model  $M^*$ . For each multi-hop question influenced by the editing operations in  $E$ ,  $M^*$  should be able to infer the correct answer  $o_n^*$ . Meanwhile, the reasoning chain needs to align with  $C^*$ , which we call **the golden path** for question  $Q$ . KEDKG freezes the parameters of  $M$  and utilizes a dynamic knowledge graph  $\mathcal{G}$  generated by the editing set  $E$  as an external knowledge base to guide  $M$  in generating the answer.

## Methodology

KEDKG is a lightweight memory-based knowledge editing approach that stores edited knowledge in the form of a knowledge graph. This facilitates subsequent retrieval and enables dynamic updates to the graph in response to real-world knowledge changes. Additionally, we decouple multi-hop question decomposition from the original model to preserve its performance, handling this decomposition by fine-tuning an additional language model. As illustrated in the Figure 2, KEDKG consists of two main steps: (1) *dynamic knowledge graph construction*, and (2) *knowledge graph augmented generation*.

### Dynamic Knowledge Graph Construction

**Edited knowledge triples extraction.** Given a set of plain texts  $T = \{t_1, t_2, \dots, t_n\}$  that represent edited knowledge in natural language form, KEDKG employs a relation extraction model to transform these natural texts into structured triples. This process can be described as follows:

$$E_i = \{e_1^i, e_2^i, \dots, e_m^i\} = \text{Retriever}(t_i), i \in [1, n] \quad (1)$$

where *Retriever* is the relation extraction model.

Previous approaches typically stored natural texts directly in memory, retrieving the most relevant knowledge based on semantic similarity to the question. However, this method is often influenced by syntactic structures and a variety of comparable words, which significantly affects accuracy. By pre-extracting relations from natural texts, we can more effectively capture the relationships between entities. As

shown in Figure 2, from the sentence “*The author of Misery is Richard Dawkins*”, we can extract two pieces of edited knowledge: (Misery, author, Richard Dawkins) and (Richard Dawkins, notable work, Misery), the latter of which is often overlooked by conventional editing methods.

**Entity linking.** To store the edited knowledge in triple form into the knowledge graph, we note that different entity names might refer to the same entity (e.g., United States and U.S. denote the same entity). We link the extracted entities to Wikidata and store their corresponding node IDs and aliases in the knowledge graph.

**Conflict detection and modification.** Traditional knowledge editing methods do not account for the necessity of updating edited knowledge as the world evolves, a process we refer as **secondary editing**. Secondary editing can introduce conflicting knowledge which, if stored directly in memory, negatively impacts subsequent retrieval and is difficult to detect. The utilization of knowledge graphs can simplify this process. For a new piece of edited knowledge  $e_{new} = (s, r, o_{new}^*)$ , we locate the old edited knowledge  $e_{old} = (s, r, o_{old}^*)$  based on the subject  $s$  and relation  $r$ , then remove it and add  $e_{new}$  to the knowledge graph. If  $e_{old}$  is not found, this indicates that the knowledge has not undergone secondary editing, and  $e_{new}$  can be added directly.

### Knowledge Graph Augmented Generation

Unlike the MeLLO (Zhong et al. 2023) approach, which iteratively decomposes multi-hop questions, we propose an approach that utilizes a template  $P_{divide}$  to guide a trained LLM to decompose a multi-hop question  $Q$  into multiple sub-questions in a single step:

$$\{q_1, q_2, \dots, q_n\} = LLM(P_{divide}(Q)) \quad (2)$$

where  $P_{divide}(Q)$  is the prompt obtained by filling the multi-hop question  $Q$  into the template  $P_{divide}$ . Except for the first sub-question  $q_1$ , which contains the subject  $s$  from  $Q$ , the subjects in other sub-questions are replaced by a special marker  $[ENT]$ . We believe this can effectively reduce the burden on LLMs in understanding contexts during iteration, thereby improving the accuracy of question decomposition.

To search for relevant edited knowledge in the knowledge graph, we employ a fine-grained retrieval scheme and filter the retrieval results using **entity and relation detectors**. When processing each generated sub-question  $q_i$ , we first use a relation extraction model to generate a set of possible entities  $K = \{k_1, \dots, k_n\}$ . Then, we use an entity detector  $g_\phi(q_i, k_j) \rightarrow [0, 1]$  to predict the probability that entity  $k_j$  is the subject of  $q_i$ . We select  $s = \arg \max_{k_i \in K} g_\phi(q_i, k_j)$  as the subject and link it to the knowledge graph. If the subject does not exist in the knowledge graph, it indicates that no information related to the subject needs to be modified, and we directly call the original LLM for output. Otherwise, we can easily obtain the set of relations  $R = \{r_1, \dots, r_m\}$  related to the entity and use a relation detector  $g_\psi(q_i, r_j) \rightarrow [0, 1]$  to detect the probability of the relation  $r_j$  being relevant to the subject in  $q$ . If the highest probability  $p$  exceeds a threshold  $\alpha$ , which is set to 0.5 in our experiments, we can retrieve the

corresponding fact triple  $(s, r, o^*)$  and use  $o^*$  as the retrieval answer. Otherwise, we directly call the original LLM for the answer. The final answer  $a_i$  to the sub-question  $q_i$  can be expressed using the formula:

$$a_i = \begin{cases} LLM(P_{retrieve}(q_i, o^*)), & \text{if } p > \alpha \\ LLM(P_{answer}(q_i)), & \text{if } s \notin G \text{ or } p \leq \alpha \end{cases} \quad (3)$$

$$o^* = \mathcal{G}(s, \arg \max_{r_j \in R} g_\psi(q_i, r_j)) \quad (4)$$

where  $P_{answer}$  is the template guiding the original LLM in generating answers,  $P_{retrieve}$  is a template that prompts the LLM to refine the response based on the retrieved entities, and  $\mathcal{G}$  is the knowledge graph constructed in Step 1.

Entity detector and relation detector can be implemented using any text classification model. After obtaining the answer  $a_i$  to  $q_i$ , we use  $a_i$  as the subject to fill in the special marker in the next sub-question  $q_{i+1}$ . This process is iteratively repeated until all sub-questions are resolved, culminating in the output of the final answer.

### Training Objectives

**Question decomposition module.** This module is designed to enable the LLM to learn both explicit and implicit atomic sentence meanings and the relationships between sentences. It enhances the LLM’s ability to decompose a complex question into multiple sub-questions. This objective focuses on training the LLM to predict subsequent tokens based on previous tokens. Specifically, given the predefined prompt template  $\mathbf{p}$  and question  $\mathbf{q}$ , the objective function for generating the sub-questions sequence  $o = [o_1, \dots, o_T]$  is outlined as follows:

$$\mathcal{L}_{dec}(\theta) = - \sum_{t=1}^T \log p_\theta(o_t | o_{<t}, \mathbf{p}, \mathbf{q}) \quad (5)$$

where  $\theta$  represents the parameters of the model, and the output sequence  $o = [q_1, \dots, q_H]$  is composed of multiple sub-clauses derived from decomposition.  $q_i$  represents each individual sub-question, and  $H$  represents the number of sub-questions.

**Entity detector and relation detector.** Specifically, given the input pair  $(q_i, k_j)$  or  $(q_i, r_j)$ , the training objective of the entity detector  $g_\phi$  and relation detector  $g_\psi$  is to minimize the average binary cross-entropy loss over the training dataset:

$$\mathcal{L}_{det} = - \frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \quad (6)$$

where  $N$  denotes the number of training samples.  $y_i$  represents the true label of the  $i$ -th sample, and  $p_i$  represents the probability that the  $i$ -th sample is predicted to be a subject or relation in sub-question  $q_i$ . For more details about training, please refer to Appendix C and D.

## Experiments

### Experimental Settings

**Datasets.** We evaluate KEDKG using the MQuAKE dataset. MQuAKE is a knowledge editing benchmark for

Method	MQUAKE-CF-3K						MQUAKE-T			
	1 edited		100 edited		All edited		1 edited		All edited	
	M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.
<b>LLaMa 2-7B</b>										
FT <sub>COT</sub> *	22.30	-	2.13	-	OOM	-	47.32	-	3.75	-
FT*	28.20	7.30	2.37	0.03	OOM	OOM	56.48	33.89	1.02	0.37
ROME <sub>COT</sub> *	11.17	-	2.87	-	2.77	-	28.96	-	14.40	-
ROME*	13.13	5.37	3.50	0.03	3.63	0.10	24.89	17.99	1.71	0.32
MEMIT <sub>COT</sub> *	11.83	-	9.23	-	5.57	-	36.88	-	31.58	-
MEMIT*	14.97	6.43	9.40	2.47	2.30	0.37	30.89	23.98	25.21	20.13
MeLLo†	33.57	9.90	20.00	10.07	17.33	9.90	65.78	55.27	57.69	44.55
PokeMQA*	<u>44.13</u>	<u>30.60</u>	<u>37.33</u>	<u>27.83</u>	<u>32.83</u>	<u>23.87</u>	<b>75.43</b>	<u>60.44</u>	<b>74.36</b>	<u>60.22</u>
KEDKG (Ours)	<b>66.80</b>	<b>63.67</b>	<b>58.50</b>	<b>55.37</b>	<b>48.30</b>	<b>43.90</b>	<u>73.13</u>	<b>69.06</b>	<u>71.15</u>	<b>66.76</b>
<b>Vicuna-7B</b>										
MeLLo*	30.70	20.84	24.75	12.25	22.35	10.18	60.72	48.55	51.55	42.97
PokeMQA*	<u>45.83</u>	<u>34.80</u>	<u>38.77</u>	<u>31.23</u>	<u>31.63</u>	<u>25.30</u>	<b>74.57</b>	<u>55.19</u>	<u>73.07</u>	<u>55.09</u>
KEDKG (Ours)	<b>68.60</b>	<b>65.13</b>	<b>62.43</b>	<b>58.20</b>	<b>51.10</b>	<b>44.67</b>	<u>71.90</u>	<b>67.23</b>	<b>74.68</b>	<b>66.54</b>
<b>GPT-3.5-turbo-instruct</b>										
MeLLo*	57.43	28.80	40.87	28.13	35.27	25.30	<b>88.12</b>	52.84	74.57	53.53
PokeMQA*	<u>67.27</u>	<u>56.37</u>	<u>56.00</u>	<u>49.63</u>	<u>45.87</u>	<u>39.77</u>	76.98	<u>68.09</u>	<b>78.16</b>	<u>67.88</u>
KEDKG (Ours)	<b>68.00</b>	<b>65.33</b>	<b>59.50</b>	<b>56.80</b>	<b>49.10</b>	<b>43.17</b>	<u>78.75</u>	<b>76.18</b>	<u>77.19</u>	<b>73.77</b>

Table 1: The experimental results on MQUAKE-CF-3K and MQUAKE-T benchmarks. The best results are indicated in **bold**, and the second-best results are underlined. Results marked with \* indicate they are sourced from (Gu et al. 2024a), and results marked with † indicate they are sourced from (Cheng et al. 2024). The notation “ $k$  edited” refers to the size of the edit batch is  $k$ . “COT” implies that the current method employs chain-of-thought prompting, otherwise, a decomposition prompt is used. The evaluation metrics are multi-hop accuracy (M-Acc.) and hop-wise answering accuracy (H-Acc.). “OOM” denotes that the method does not work due to high GPU memory usage.

multi-hop QA, comprising MQUAKE-CF based on counterfactual editing and MQUAKE-T based on temporal knowledge updates. We use MQUAKE-CF as the training set, which contains 9,218 data points, and MQUAKE-CF-3k as the test set, which includes 3,000 data points. It is important to note that there are no overlapping data points between these two sets. These datasets include numerous  $k$ -hop questions ( $k \in \{2, 3, 4\}$ ), with each question corresponding to  $n$  edits ( $n \in [1, k]$ ). See Appendix A for more details.

**Baselines.** We compare KEDKG with common knowledge editing methods, including parameter-preserving and parameter-based editing methods. The parameter-preserving editing methods include MeLLo (Zhong et al. 2023) and PokeMQA (Gu et al. 2024a), while the parameter-based editing methods include FT (Zhu et al. 2020), ROME (Meng et al. 2022), and MEMIT (Meng et al. 2023). We also report the performance of these methods under chain-of-thought (CoT) and question decomposition (QD) prompts. See Appendix B for more details.

**Evaluation metrics.** Following previous work (Zhong et al. 2023; Gu et al. 2024a), we use Multi-hop Accuracy (M-Acc) and Hop-wise Answering Accuracy (H-Acc) as evaluation metrics. Multi-hop accuracy measures the accuracy of the LLM in answering multi-hop questions. However, in some cases, the LLM might produce an incorrect reasoning process but still arrive at the correct answer. Hop-wise answering accuracy measures the accuracy where every an-

swer in the reasoning path is correct, avoiding such interference. This is the primary metric we focus on to assess the model’s ability to use edited knowledge. For all metrics, higher values indicate better performance.

**Experimental setup.** We train an entity detector and a relation detector based on the DistilBERT (Sanh 2019) model and fine-tune the Llama 2-7B model for the question decomposition task. In addition, we use REBEL (Cabot and Navigli 2021) as our relation extraction model and *spacy entity linker* as entity linking model. It is important to note that there is no overlap between the training and test datasets we use. We conducted training on multiple base models, including mainstream open-source models Llama 2-7B (Touvron et al. 2023) and Vicuna-7B (Zheng et al. 2024), as well as the black-box model GPT-3.5-turbo-instruct. Parameter-preserving editing methods are applicable to all models, while parameter-based editing methods are only applicable to open-source models. Therefore, we only reported the experimental results of parameter-based editing methods on open-source models. All our experiments are carried out on a NVIDIA 8×A800-SXM4-80G machine.

## Main Results

Table 1 shows the experimental results on MQUAKE datasets. We can draw the following observations:

- (1) KEDKG demonstrates outstanding performance across all metrics, notably surpassing all baselines in the H-Acc metric. Utilizing Llama 2-7B as the base model,

$g_\phi$	$g_\psi$	CDM	MQUAKE-CF-3k						MQUAKE-T			
			1 edited		100 edited		All edited		1 edited		All edited	
			M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.	M-Acc.	H-Acc.
-	✓	✓	35.23	31.63	31.17	26.93	23.10	18.87	54.39	50.05	53.80	48.50
✓	-	✓	20.03	16.03	19.13	15.10	16.80	12.57	54.98	52.19	56.48	52.78
✓	✓	-	66.30	63.43	58.43	54.60	41.10	36.83	73.13	68.74	71.15	63.81
✓	✓	✓	<b>66.80</b>	<b>63.67</b>	<b>58.50</b>	<b>55.37</b>	<b>48.30</b>	<b>43.90</b>	<b>73.13</b>	<b>69.06</b>	<b>71.15</b>	<b>66.76</b>

Table 2: Ablation results of KEDKG and its variants in terms of Multi-hop Accuracy (M-Acc.) and Hop-wise Answering Accuracy (H-Acc.) on MQUAKE datasets. ‘‘CDM’’ refers to the Conflict Detection and Modification module.

KEDKG achieves an improvement of 108.1% in 1 edit batch, 96.2% in 100 edit batches, and 83.9% in all edit batches compared to the best previous baseline. This robust performance underscores KEDKG’s capability to effectively answer multi-hop questions under conditions of knowledge editing and adapt seamlessly to varying numbers of edits.

(2) Notably, other methods exhibit a significant discrepancy between M-Acc and H-Acc, indicating that their reasoning processes do not align with the task logic presented in the prompt, which can be seen as a mismatch in question decomposition. In contrast, KEDKG demonstrates a smaller difference between M-Acc and H-Acc, primarily because we fine-tune the model specifically on the question decomposition task, decoupling it from the original model and reducing the burden of understanding complex prompts. This results in more reliable reasoning outcomes.

(3) Interestingly, on the MQUAKE-CF-3k benchmark, KEDKG’s performance with a 7B parameter open-source base model matches or even surpasses that of GPT-3.5-turbo-instruct. This comparison reveals that the knowledge capability of the 7B parameter model is on par with that of the black-box model, though the latter demonstrates a stronger ability to understand and follow human instructions. KEDKG effectively bridges this gap, enhancing the 7B parameter model’s ability to adeptly handle multi-hop QA problems under conditions of knowledge editing.

(4) We also observe that all baselines perform best under the 1 edit batch setting, with performance declining as the edit batch size increases, a trend that is particularly pronounced in the MQUAKE-CF-3k dataset. This phenomenon can be attributed to two main factors: first, conducting retrieval within large batches of edits is a challenging task, especially when the entities and relationships involved are very similar. Second, in large batch edits, the edited knowledge in different cases within the MQUAKE-CF-3k dataset may conflict with each other. Thanks to the integration of a conflict detector and the application of dynamic knowledge graphs, we have effectively improved KEDKG’s performance in scenarios involving large batch edits. Even under the all edit batch setting, KEDKG can still achieve an accuracy rate close to 50%.

### Ablation Study

We conduct ablation experiments with the base model Llama 2-7B on the MQUAKE-CF-3k and MQUAKE-T datasets to analyze the impact of  $g_\phi$ ,  $g_\psi$ , and conflict detection modification (CDM) modules.

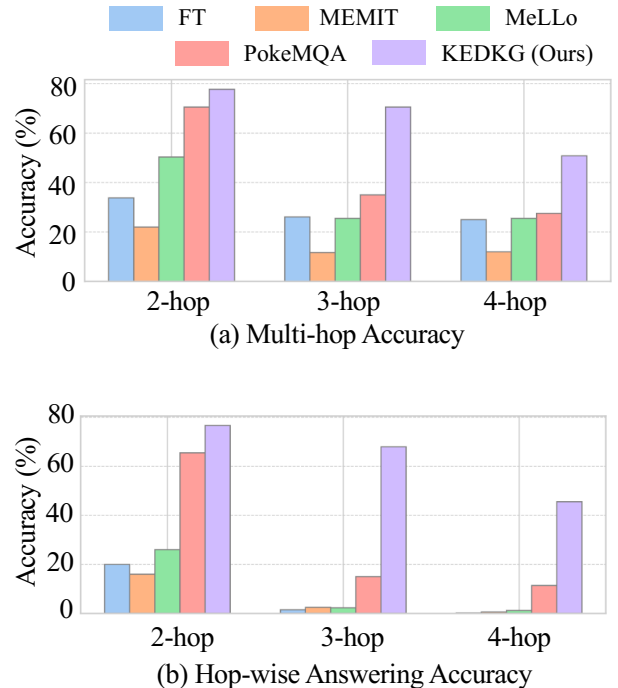


Figure 3: Multi-hop Accuracy and Hop-wise Answering Accuracy results on MQUAKE-CF-3K, utilizing different knowledge editing methods. The experiments are conducted on Llama 2-7B and the edit batch size is 1.

cation (CDM) modules.

**The trained entity detector  $g_\phi$  and relation detector  $g_\psi$  have proven to be effective.** As shown in Table 2, replacing  $g_\phi$  or  $g_\psi$  with the semantic similarity model, all-MiniLM-L6-v2 from Huggingface (Wolf et al. 2020), resulted in a decrease in answer accuracy across all experimental settings. The most significant performance decline was observed when the relation detector was replaced, with a 31.5% drop in Multi-hop Accuracy under the ‘‘all edited’’ setting. This decline highlights the limitations of relying solely on semantic similarity to select relations from the knowledge graph, as a specific relation may have various expressions, making such an approach unreliable.

```

Edited Knowledge: Association football was created in the country of Brazil. Brazil is located in the continent of Africa.
Knowledge Graph: (Association football, country of origin, Brazil); (Brazil, sport, Association football); (Brazil, continent, Africa)

Question: What is the continent of origin for the sport associated with Watford F.C.?
Subquestions: 1. Which sport is Watford F.C. associated with? 2. Which country was [ENT] created in? 3. Which continent is [ENT] located in?
Thought 1: Which sport is Watford F.C. associated with?
Retrieved answer: [None]
Generated answer: Watford F.C. is associated with Association Football (Soccer).
Thought 2: Which country was Association Football (Soccer) created in?
Retrieved answer: Brazil.
Generated answer: Association Football (Soccer) was created in Brazil.
Thought 3: Which continent is Brazil located in?
Retrieved answer: Africa.
Generated answer: Brazil is located in Africa.
Final answer: Africa

```

Table 3: A case study of KEDKG solving one 3-hop question in MQUAKE-CF-3K. Yellow parts are generated by the fine-tuned LLM for question decomposition. Green parts are generated by the base LLM for question answering, and blue parts are answers retrieved in the knowledge graph.

**The CDM module effectively enhances performance in batch editing.** The MQUAKE-CF dataset includes some conflicting edits, providing an opportunity to evaluate KEDKG’s conflict resolution capabilities. As shown in Table 2, the CDM module has no significant effect in small batch editing, which is intuitive since there are rarely conflicting edits in small batches. However, as the number of edits increases, the probability of conflicts also rises. Removing the CDM module led to a 7.2% drop in Multi-hop Accuracy. This decrease is attributed to the presence of conflicting knowledge triples within the knowledge graph, which disrupt the retrieval process. Introducing the CDM module effectively minimizes this interference and facilitates the execution of secondary edits. On the MQUAKE-T dataset, we observe that applying the CDM module does not significantly impact the results. This is intuitive because MQUAKE-T itself does not contain conflicting edits.

## Analysis

**KEDKG excels in 3-hop and 4-hop question answering tasks.** As shown in Figure 3, both KEDKG and PokeMQA perform well in two-hop questions. However, in 3-hop and 4-hop questions, KEDKG significantly outperforms PokeMQA, almost doubling the performance metrics. In the most complex 4-hop questions, KEDKG achieves over 50% Multi-hop Accuracy, while the accuracy of other methods remains below 25%. This superior performance is attributed to our training on the decomposition task, which effectively minimizes the likelihood of errors or hallucinations in the LLM during the decomposition process. Additionally, the fine-grained retrieval based on the knowledge graph significantly improves the accuracy of answers at each hop.

Meanwhile, we have also observed that **parameter-based editing methods generally underperform compared to parameter-preserving editing methods in multi-hop QA tasks**, exhibiting notably low Hop-wise Answering Accu-

racy. This suggests that current parameter-based editing methods still struggle to enable LLMs to flexibly utilize edited knowledge during the reasoning process.

## Case Study

We conduct a case study as presented in Table 3. KEDKG first extracts edited triples from the given edited knowledge. After passing through the entity linking and conflict detection modules, these triples are stored in the knowledge graph  $\mathcal{G}$ . Note that Table 3 only shows the edited knowledge relevant to this case. In large-scale editing, there are many instances of edited knowledge and conflicting edits.

In the knowledge graph-based QA phase, the fine-tuned LLM used for question decomposition breaks down the multi-hop question into three sub-questions, where the subjects of the latter two sub-questions are replaced with special markers. Subsequently, the entity detector and relation detector are applied sequentially to retrieve answers for the three sub-questions. For the first sub-question, since no relevant facts about *Watford F.C.* are found in  $\mathcal{G}$ , the question is directly handed to the base LLM to produce an answer, and the entity in the answer is then filled into the next sub-question. For the latter two questions, KEDKG successfully retrieves the edited facts from  $\mathcal{G}$  and prompts the LLM to provide answers, ultimately yielding the correct answer.

## Conclusion

In this paper, we propose a novel knowledge editing method, KEDKG for MHQA. KEDKG constructs a dynamic knowledge graph to store edited knowledge and to handle conflicting edits. Additionally, KEDKG employs a fine-grained retrieval scheme to fetch edited knowledge from the knowledge graph, guiding the model to modify its answers. Extensive experiments on the MQUAKE dataset demonstrate that KEDKG outperforms previous knowledge editing methods.

## Acknowledgements

This work was supported in part by National Science Foundation of China (62476070), Shenzhen Science and Technology Program (JCYJ20241202123503005, GXWD20231128103232001) and Department of Science and Technology of Guangdong (2024A1515011540). This work was also supported in part by the Major Key Project of PCL under Grant PCL2024A06 and PCL2022A05, and in part by the Shenzhen Science and Technology Program under Grant RCJC20231211085918010. This research is supported by the NSFC (Grant: 62376074), the Shenzhen Science and Technology Program (Grants: SGDX20230116091244004, JSGGKQTD20221101115655027, KJZD20231023095959002).

## References

- Cabot, P.-L. H.; and Navigli, R. 2021. REBEL: Relation extraction by end-to-end language generation. In *Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2021*, 2370–2381.
- Chen, S.; Hou, Y.; Cui, Y.; Che, W.; Liu, T.; and Yu, X. 2020. Recall and Learn: Fine-tuning Deep Pretrained Language Models with Less Forgetting. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7870–7881.
- Cheng, K.; Lin, G.; Fei, H.; Yu, L.; Ali, M. A.; Hu, L.; Wang, D.; et al. 2024. Multi-hop question answering under temporal knowledge editing. *arXiv preprint arXiv:2404.00492*.
- De Cao, N.; Aziz, W.; and Titov, I. 2021. Editing Factual Knowledge in Language Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 6491–6506.
- Dong, Q.; Dai, D.; Song, Y.; Xu, J.; Sui, Z.; and Li, L. 2022. Calibrating Factual Knowledge in Pretrained Language Models. In *Proceedings of the Findings of the Association for Computational Linguistics (EMNLP)*, 5937–5947.
- Gu, H.; Zhou, K.; Han, X.; Liu, N.; Wang, R.; and Wang, X. 2024a. PokeMQA: Programmable knowledge editing for Multi-hop Question Answering. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, 8069–8083.
- Gu, J.-C.; Xu, H.-X.; Ma, J.-Y.; Lu, P.; Ling, Z.-H.; Chang, K.-W.; and Peng, N. 2024b. Model editing can hurt general abilities of large language models. *arXiv preprint arXiv:2401.04700*.
- Huang, J.; and Chang, K. C.-C. 2023. Towards Reasoning in Large Language Models: A Survey. In *Proceedings of the Findings of the Association for Computational Linguistics: ACL 2023*, 1049–1065.
- Huang, Z.; Shen, Y.; Zhang, X.; Zhou, J.; Rong, W.; and Xiong, Z. 2023. Transformer-Patcher: One Mistake Worth One Neuron. In *Proceedings of the The Eleventh International Conference on Learning Representations (ICLR)*.
- Mavi, V.; Jangra, A.; and Jatowt, A. 2022. A survey on multi-hop question answering and generation. *arXiv preprint arXiv:2204.09140*.
- Meng, K.; Bau, D.; Andonian, A.; and Belinkov, Y. 2022. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 17359–17372.
- Meng, K.; Sharma, A. S.; Andonian, A. J.; Belinkov, Y.; and Bau, D. 2023. Mass-Editing Memory in a Transformer. In *Proceedings of The Eleventh International Conference on Learning Representations (ICLR)*.
- Mitchell, E.; Lin, C.; Bosselut, A.; Manning, C. D.; and Finn, C. 2022. Memory-based model editing at scale. In *Proceedings of the International Conference on Machine Learning (ICML)*, 15817–15831.
- Sanh, V. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. In *Proceedings of Thirty-third Conference on Neural Information Processing Systems (NeurIPS)*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, S.; Zhu, Y.; Liu, H.; Zheng, Z.; Chen, C.; et al. 2023. Knowledge editing for large language models: A survey. *arXiv preprint arXiv:2310.16218*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems (NeurIPS)*, 35: 24824–24837.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations (EMNLP)*, 38–45.
- Zhang, N.; Yao, Y.; Tian, B.; Wang, P.; Deng, S.; Wang, M.; Xi, Z.; Mao, S.; Zhang, J.; Ni, Y.; et al. 2024. A comprehensive study of knowledge editing for large language models. *arXiv preprint arXiv:2401.01286*.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Zheng, C.; Li, L.; Dong, Q.; Fan, Y.; Wu, Z.; Xu, J.; and Chang, B. 2023. Can We Edit Factual Knowledge by In-Context Learning? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4862–4876.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems (NeurIPS)*, 36.
- Zhong, Z.; Wu, Z.; Manning, C. D.; Potts, C.; and Chen, D. 2023. MQuAKE: Assessing Knowledge Editing in Language Models via Multi-Hop Questions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 15686–15702.

Zhu, C.; Rawat, A. S.; Zaheer, M.; Bhojanapalli, S.; Li, D.; Yu, F.; and Kumar, S. 2020. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*.