

# RFKG-CoT: Relation-Driven Adaptive Hop-count Selection and Few-Shot Path Guidance for Knowledge-Aware QA

Chao Zhang, Minghan Li\*, Tianrui Lv, Guodong Zhou

School of Computer Science and Technology, Soochow University, China  
 czhang1@stu.suda.edu.cn, mhli@suda.edu.cn, trlvtrl@stu.suda.edu.cn, gdzhou@suda.edu.cn

## Abstract

Large language models (LLMs) often generate hallucinations in knowledge-intensive QA due to parametric knowledge limitations. While existing methods like KG-CoT improve reliability by integrating knowledge graph (KG) paths, they suffer from rigid hop-count selection (solely question-driven) and underutilization of reasoning paths (lack of guidance). To address this, we propose RFKG-CoT: First, it replaces the rigid hop-count selector with a relation-driven adaptive hop-count selector that dynamically adjusts reasoning steps by activating KG relations (e.g., 1-hop for direct “brother” relations, 2-hop for indirect “father-son” chains), formalized via a relation mask. Second, it introduces a few-shot in-context learning path guidance mechanism with CoT (think) that constructs examples in a “question-paths-answer” format to enhance LLMs’ ability to understand reasoning paths. Experiments on four KGQA benchmarks show RFKG-CoT improves accuracy by up to 14.7 pp (Llama2-7B on WebQSP) over KG-CoT. Ablations confirm the hop-count selector and the path prompt are complementary, jointly transforming KG evidence into more faithful answers.

## Introduction

In recent years, large language models (LLMs) such as the GPT series (Achiam et al. 2023) and Llama series have made breakthrough progress in natural language processing, excelling in text generation, question answering (QA), and logical reasoning (Wei et al. 2022). However, in knowledge-intensive tasks, LLMs still face two core challenges: first, hallucination—due to limited reliable knowledge, LLMs sometimes struggle to generate credible answers based on black-box parametric knowledge (Ji et al. 2023); second, factual staleness—static parameters make it hard to obtain real-time domain knowledge (Wang, Ma, and Chen 2023). These issues severely restrict applications in high-accuracy scenarios like medical diagnosis and legal consulting.

To address the above problems, researchers have proposed knowledge enhancement methods, which provide factual support for LLMs through external knowledge sources (e.g., knowledge graphs, text corpora). Among them, knowledge graphs (KGs) have become an important knowledge

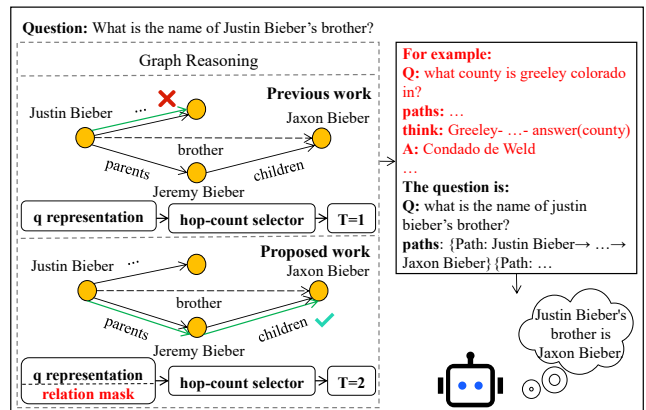


Figure 1: Comparison of RFKG-CoT and KG-CoT, highlighting two key improvements: relation-driven hop-count selector (left red) and few-shot ICL path guidance (right red). The green arrows indicate the reasoning paths of different models.

enhancement tool due to their structured and interpretable characteristics. Their triple structure, composed of entities (nodes) and relations (edges), can clearly express semantic associations between entities (Bollacker, Cook, and Tufts 2007). KG-based enhancement methods fall into two main categories: pre-training fusion, which embeds KG knowledge into LLM parameters (e.g., fusing KG information during pre-training (Zhao et al. 2024a)), and reasoning-phase fusion, which dynamically introduces KG knowledge during inference (e.g., Retrieval-Augmented Generation (RAG) (Lewis et al. 2020)). As a technique to improve LLMs’ reasoning ability, Chain-of-Thought (CoT) prompting significantly enhances model performance in logical reasoning tasks by guiding the model to generate intermediate reasoning steps and decomposing complex problems into simple subproblems (Wei et al. 2022). However, traditional CoT relies solely on LLMs’ parametric knowledge, still suffering from factual errors. To this end, KG-CoT has been proposed. It generates reasoning paths using structured KG knowledge and inputs these paths into LLMs as prompts to enable knowledge-aware reasoning (Zhao et al. 2024b).

Zhao et al. (2024b) attempts to dynamically adjust rea-

\*Corresponding author

soning steps using a graph reasoning model and hop-count selector. However, two critical limitations remain: (1) The hop-count selector underutilizes KG relational features, relying solely on question features to determine hops. For example, answering “Who is Justin Bieber’s brother?” requires 1 hop with a direct “brother” relation but multiple hops with indirect “father-son” chains (his father’s other sons), yet the selector fails to adapt. (2) LLMs underutilize reasoning paths, as paths are input without guidance on interpretation, limiting the utilization of logical connections.

To address these issues, this paper proposes Relation-Driven Adaptive Hop-count Selection and Few-Shot Path Guidance for Knowledge-Aware QA (RFKG-CoT), an enhanced framework built on KG-CoT, as shown in Figure 1. This paper makes two key contributions: First, we propose a relation-based dynamic reasoning step mechanism that enables reasoning steps to adaptively adjust path length based on entity relations in the knowledge base. Second, we introduce a few-shot in-context learning (ICL) strategy (Dong et al. 2024) with CoT (think) to guide LLMs in better utilizing reasoning paths through examples, thereby enhancing knowledge-aware reasoning capabilities.

## Related Work

The fusion of KGs and LLMs has emerged as a pivotal research direction in recent years, aiming to leverage KG-structured knowledge to compensate for LLMs’ factual inaccuracies in knowledge-intensive tasks (Hu et al. 2023). Based on the integration stage, such methods are categorized into two paradigms: pre-training fusion and reasoning-phase fusion. Pre-training fusion injects KG information during model pre-training to encode structured knowledge into LLM parameters. However, existing approaches (e.g., K-BERT (Liu et al. 2020), KG-BART (Liu et al. 2021)) suffer from two critical limitations: inapplicability to closed-source LLMs (e.g., GPT-4) due to parameter inaccessibility, and difficulty in updating dynamic KG knowledge (e.g., newly emerging entities) without costly re-pre-training. In contrast, reasoning-phase fusion dynamically introduces KG knowledge during LLM inference, offering superior flexibility and interpretability. Retrieval-Augmented Generation (RAG) (Lewis et al. 2020) is a representative method, which retrieves relevant triples from KGs, converts them into text, and feeds them to LLMs (Zhu et al. 2025). Recent advancements include integrating graph structure for precise subgraph retrieval (Mavromatis and Karypis 2024) and iterative retrieval with LLM feedback (Yang et al. 2025).

The challenge of dynamically determining the optimal reasoning depth in knowledge-intensive tasks represents a critical limitation in contemporary Chain-of-Thought (CoT) methodologies. Prior approaches largely fall into two categories: fixed-step reasoning and task-driven step selection. Traditional CoT frameworks (Wei et al. 2022) typically employ a predetermined number of reasoning steps, regardless of query complexity or underlying knowledge structure. This rigidity proves particularly suboptimal for KGs. KG-CoT (Zhao et al. 2024b) introduced question-driven hop-count selection, dynamically setting hop counts based on

question intent classification. However, its hop-count selection relies solely on the initial question embedding and lacks an explicit guidance mechanism for the reasoning path. This makes it difficult for LLMs to effectively utilize intermediate reasoning results, limiting their ability to reason with complex logic. We introduce a relation mask mechanism to fuse KG relational features with the question embedding to jointly generate the hop-count decisions.

Few-Shot Learning leverages models’ pre-trained knowledge and strong contextual understanding, enabling rapid adaptation to new tasks (e.g., style-specific translation, text classification, code generation) with only a small number of examples. In KG-related tasks, few-shot is mainly used for relationship prediction and question answering. For instance, FewRel (Han et al. 2018) trains models to predict new relations through a small number of triple examples; KG-FewShotQA (Li et al. 2023) constructs examples containing knowledge graph paths to guide LLMs to generate answers. However, such designs primarily focus on path-answer mapping and lack explicit components to guide how LLMs should interpret path elements or derive answers step-by-step from paths. To address this gap, we extend the path-based framework by fusing questions with reasoning paths to generate structured thinking chains, including a symbolic *Think* template that maps path elements to step-by-step logic. This strengthens path-reasoning connections, enhancing LLMs’ utilization of path semantics for answer derivation.

## Methodology

RFKG-CoT first performs reasoning on the knowledge graph through a small-scale graph reasoning model, then generates reasoning paths. By collaborating with large language models, it combines the model’s own knowledge and generated reasoning paths to derive the final result. Compared with the previous method (Zhao et al. 2024b), we have improved the graph reasoning model by optimizing the hop-count selection. We further provide few-shot exemplars that teach the LLM how to ground each path, ensuring it attends to the KG evidence. The framework of RFKG-CoT is illustrated in Figure 2.

### Graph Reasoning Model

We employ a graph reasoning model that decomposes problems into multiple steps via a hop-count selector, calculating relationship scores at each step to form executable multi-step reasoning paths on the knowledge graph.

**Initialization.** We let  $\mathcal{G}$  denote the knowledge graph, with  $n$  representing the number of entities in the entity set and  $m$  representing the number of relations in the relation set. The state of the topic entity in the problem is initialized using one-hot encoding as  $\mathbf{e}^0 \in [0, 1]^n$ , where the position corresponding to the topic entity in the problem is set to 1 (and all others to 0). In addition, we initialize a triple matrix  $\mathbf{M} \in [0, 1]^{n \times n}$ , where  $M_{ij} = k$  indicates that the relation index between entity  $i$  and entity  $j$  is  $k$  (e.g.,  $k=3$  maps to “law.judge.cases” in WebQSP). Finally, we de-

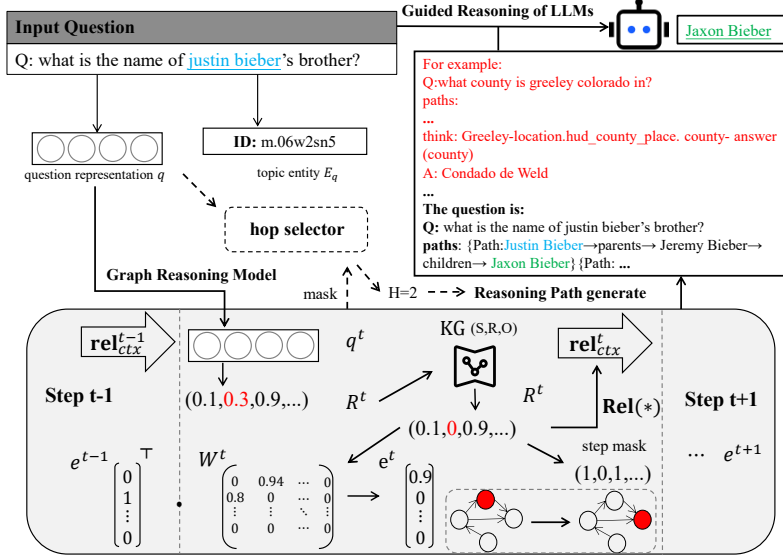


Figure 2: Overview of RFKG-CoT. (1) A graph reasoning model performs reasoning over the knowledge graph, and a hop-count selector that incorporates both the question and knowledge graph relationships determines the required reasoning hops for answering the question. (2) Reasoning paths are generated via a dedicated path generation module. (3) The question is concatenated with generated paths, and selected few-shot examples are appended as prompts to input into the large language model (LLM) for final reasoning.

fine the golden answer to the question as a one-hot vector  $\mathbf{a} \in [0, 1]^n$ , that indicates which entity is the answer.

**Relation Score Calculation.** We divide the graph reasoning process into  $T$  steps (determined by the dataset). Starting from the topic entity, at step  $t$  ( $t < T$ ), we calculate the scores for all relations in  $\mathcal{G}$  based on the problem information focused on at the current stage, as  $\mathbf{R}^t \in [0, 1]^m$ . The score of each relation  $r_i^t \in \mathbf{R}^t$  represents its activation probability, and the model transmits entity scores through these activated relations. The calculation of  $\mathbf{R}^t$  is as follows:

$$\mathbf{R}^t = \text{Sigmoid}(\text{MLP}_{KG}(\mathbf{q}^t)), \quad (1)$$

where  $\text{MLP}_{KG}$  converts the 768 dimensional question embedding  $\mathbf{q}^t$  to an  $m$  dimensional vector of relation scores, and  $\mathbf{q}^t$  denotes the query representation at step  $t$ . Since the model’s focus on key parts of the question context may shift across different reasoning stages, at step  $t$ , we focus on a specific part of the question, and the calculation of the question representation  $\mathbf{q}^t$  is as follows:

$$\mathbf{q}_t(h_1, \dots, h_{|q|}) = \text{Encoder}(q), \quad (2)$$

$$\mathbf{Q}_t = f^t([\mathbf{q}, \text{rel}_{ctx}]), \quad (3)$$

$$\mathbf{b}^t = \text{Softmax}(\mathbf{Q}_t[h_1; \dots; h_{|q|}]), \quad (4)$$

$$\mathbf{q}^t = \sum_{i=1}^{|q|} b_i^t h_i, \quad (5)$$

where  $\mathbf{q}$  denotes the question embedding, and  $(h_1, \dots, h_{|q|})$  is the hidden state sequence of the question. The function  $f^t$  is used to project the extracted question information and

the relation information ( $\text{rel}_{ctx}$  initialized as a zero vector with the same shape as  $\mathbf{q}^t$  at early steps) from the previous step onto the attention query  $\mathbf{Q}$  at step  $t$ . Then, we calculate the attention weight for each word, and obtain the question representation at step  $t$  by performing a weighted sum of these weights.

**Reasoning.** We define a state transition matrix  $\mathbf{W}^t \in [0, 1]^{n \times n}$  based on relation scores  $\mathbf{R}^t$ :

$$W_{ij}^t = \begin{cases} R_k^t & k = M_{ij}, R_k^t \in \mathbf{R}^t, M_{ij} \in \mathbf{M}, \\ 0 & \text{Otherwise,} \end{cases} \quad (6)$$

where  $k$  is the relation index between entities  $i$  and  $j$  in the triple matrix  $\mathbf{M}$ , and  $R_k^t$  is the score of relation  $k$  at step  $t$ . The reasoning process propagates entity states via:

$$e_j^t = \sum_{i=1}^n e_i^{t-1} \times W_{ij}^t. \quad (7)$$

$e^{t-1}$  denotes the entity state vector at step  $t - 1$ . This operation diffuses entity probabilities along 1-hop neighbors weighted by relation scores.

**Adapted Hop-Count Selection.** Unlike the static hop-count selector proposed in Zhao et al. (2024b) that determines hops solely based on question features, we propose a relation-driven mechanism that adapts reasoning steps to the KG structure. Specifically, we introduce a relation activation mask  $\text{mask} \in \{0, 1\}^m$  (initialized to all zeros) to track which relations are activated during reasoning, as detailed in Algorithm 1:

---

**Algorithm 1: Relation Activation Mask Generation**

---

**Input:** Triple matrix  $\mathbf{M}$ , relation score  $\mathbf{R}^t$ .  
**Initialize:**  $\mathbf{mask} \in \{0, 1\}^m \leftarrow \mathbf{0}$ .  
**Output:** Relation activation mask  $\mathbf{mask}$ .

```
1: for  $t = 1, \dots, T$  do
2:   Initialize step-wise mask  $\mathbf{mask}_i \in \{0, 1\}^m \leftarrow \mathbf{0}$ .
3:   Get triple indices  $\mathbf{I}_t = \{j \mid \text{obj\_p}[j] > 0\}$ .
4:   Extract relation index  $\mathbf{I}_r = \{M_{i,j} \mid j \in \mathbf{I}_t, i \in \text{sub\_p}[j]\}$  (relations from subject entities  $i$  to target entities  $j$ ).
5:   Activate step-wise mask  $\mathbf{mask}_i[k] = 1$  for all  $k \in \mathbf{I}_r$ .
6:   Filtered relation scores  $\mathbf{R}^t \leftarrow \mathbf{R}^t \odot \mathbf{mask}_i$ .
7:   Update global mask  $\mathbf{mask} = \mathbf{mask} \vee (\mathbf{R}^t > 10^{-6})$ .
8: end for
9: return  $\mathbf{mask}$ 
```

---

In each of the  $T$  reasoning steps, we first identify activated relations via non-zero target entity probabilities (stored in  $\text{obj\_p}$ , where  $\text{obj\_p}[i]$  denotes the probability of the target entity in the  $i$ -th triple, computed as the product of subject probability  $\text{sub\_p}$  and relation probability  $\text{rel\_p}$  for the triple), update a step-wise mask to filter the current relation scores  $\mathbf{R}^t$ , and accumulate valid activations into the  $\mathbf{mask}$  (retaining relations with scores exceeding  $10^{-6}$ ). The filtered relation scores are then used to extract relation information (via  $\text{Rel}(\ast)$ ) for the next step of question encoding (see Eq. (3)), enabling the model to incorporate KG structure into reasoning adaptively.

$$\text{rel}_{ctx} = \text{Rel}(\mathbf{R}^t). \quad (8)$$

After  $T$  steps, the  $\mathbf{mask}$  captures all relations that contributed to the reasoning process. We combine this mask with the original question  $\mathbf{q}$  to determine the optimal number of hops  $H$ :

$$\mathbf{c} = \text{Softmax}(MLP_T([\mathbf{q}, \mathbf{mask}])), \quad (9)$$

where  $MLP_T$  converts the  $768+m$  dimensional concatenated vector of  $\mathbf{q}$  and  $\mathbf{mask}$  to  $T$  dimensional vector,

$$H = \arg \max_{t \in \{1, 2, \dots, T\}} c_t, \quad (10)$$

and subsequently compute the final entity score  $\bar{e}$ :

$$\bar{e} = \sum_{t=1}^T c_t e^t. \quad (11)$$

**Training.** We optimize the graph reasoning model by computing the L2 distance between  $\bar{e}$  and  $\mathbf{a}$ :

$$\mathcal{L} = \|\bar{e} - \mathbf{a}\|^2. \quad (12)$$

### Generation of Reasoning Paths

In the path generation phase, two lists  $L_c$  and  $L_m$  are maintained to store candidate reasoning paths and intermediate

---

**Algorithm 2: Generate reasoning paths**

---

**Input:** Input question  $q$ , retrieved knowledge subgraph  $\mathcal{G}$ .  
**Initialize:** Entity score  $e^0 \leftarrow$  extract topic entity entity  $\mathbf{E}_q$ , triplet matrix  $\mathbf{M} \leftarrow$  extract triples from  $\mathcal{G}$ .  
**Output:** Reasoning paths

```
1: for  $t = 1, \dots, T$  do
2:   Compute the question representation  $\mathbf{q}^t$  using (2)-(5).
3:   Compute the relation score  $\mathbf{R}^t$  using (1).
4:   Update relation activation mask  $\mathbf{mask}$  using algorithm 1.
5:   compute transition matrix  $\mathbf{W}^t$  using (6).
6:   Entity score  $e^t \leftarrow$  reasoning using (7).
7:   Compute the relation information  $\text{rel}_{ctx}$  using (8).
8: end for
9: Compute final scores  $\bar{e}$  using (9), (11) and select top- $K$  entities  $\mathbf{E}^K$ .
10: Initialize  $L_m \leftarrow$  topic entity  $E^0$ .
11: for  $t = 1, \dots, H$  do
12:   Extract t-hop paths  $\mathbf{P}^t$  using  $\mathbf{W}^t$  and paths in  $L_m$ .
13:   Update intermediate path list  $L_m$  with  $\mathbf{P}_{qj}^t$ .
14:   if Object entity  $\mathbf{E}_j \in \mathbf{E}^K$  then
15:     Update reasoning path list  $L_c$  with  $\mathbf{P}_{qj}^t$ .
16:   end if
17: end for
18: Select  $N$  paths for each top- $K$  entity  $\mathbf{E}^K$  from  $L_c$ .
19: return list of reasoning paths
```

---

reasoning paths, respectively. Starting from the topic entity, we identify the predicted entities for each sample via the hop-count selector, select the top- $K$  answer entities, and then generate reasoning paths from the topic entity to these top- $K$  predicted entities using the transition matrix  $W^1, W^2, \dots, W^T$ , as shown in Algorithm 2.  $\mathbf{P}_{qj}^t$  denotes the path starting from the topic entity to the target entity. Each answer entity may correspond to multiple candidate paths in  $L_c$ . We calculate the average of the relationship scores in each path and take it as the final path score.

### Guided Reasoning of LLMs

For the top- $K$  candidate entities with high confidence scores, we extract the path with the highest path score for each entity. Consequently, we generate  $K$  reasoning paths per question via the reasoning path generation method.

We use ‘ $\rightarrow$ ’ to connect entities and relations, and a 1-hop reasoning path is represented as:

$$\mathbf{E}_i \rightarrow \text{Rel}_{ij} \rightarrow \mathbf{E}_j. \quad (13)$$

We serialize the  $K$  reasoning paths and concatenate them with the question context to form the final input sequence. To enhance LLMs’ ability to utilize reasoning paths, we introduce a manually curated few-shot prompt containing three fixed examples. Each example follows a structured ‘‘Question-Paths-Answer’’ format with four essential components: the original natural language query, serialized KG paths in the form of  $\mathbf{E}_i \rightarrow \text{Rel}_{ij} \rightarrow \mathbf{E}_j$ , a symbolic reasoning template (Think) mapping path elements to answer con-

```

LLM Input
for example:
Q: {what did st augustine do?}
paths: {Path: Augustine of Hippo → people.person.profession → Escritor}
{Path: Augustine of Hippo → people.person.profession → Physician}
{Path: Augustine of Hippo → book.written_work.author_reverse → City of God}
{Path: Augustine of Hippo → book.written_work.author_reverse → Herramme
vuorisarna}
{Path: Augustine of Hippo → people.person.profession → Ahli Falsafah}
{Path: Augustine of Hippo → book.written_work.author_reverse → Enchiridion}
think: St. Augustine- corresponds to - Augustine of Hippo - people. person.profession-
answer(profession)
A: Augustine of Hippo held professions as an 'Escritor' (writer), 'Physician', and 'Ahli
Falsafah' (philosopher)
Q: {what county is greeley colorado in?}
paths: {Path: Greeley → location.hud_county_place.county → Condado de Weld}
A: Condado de Weld
...
The Question is
Q: ...
paths: ...
answer:

```

Figure 3: Few-shot prompt template used in the LLM stage.

straints, and explicit answer format constraints. The prompt template is shown in Figure 3. As shown, we prepend few-shot ICL exemplars with CoT (think). They calibrate the model’s attention to the critical fields of each path (entity, relation, direction), thereby reducing free-form hallucination. The later ablation study confirms that removing the exemplars drops accuracy.

## Experiments

### Datasets

We evaluate RFKG-CoT on four widely used KGQA benchmarks, covering single-hop, multi-hop, and open-domain scenarios to test its knowledge-aware reasoning capabilities comprehensively:

- **WebQSP** (Yih et al. 2016): Multi-hop QA benchmark (1-2 hops) based on Freebase with compositional questions.
- **CompWebQ** (Talmor and Berant 2018): Extension of WebQSP with increased complexity and constraints.
- **SimpleQuestions** (Petrochuk and Zettlemoyer 2018): Single-hop QA requiring direct KG lookups.
- **WebQuestions** (Berant et al. 2013): Open-domain QA from real user queries.

Following the method of Zhao et al. (2024b), we construct knowledge subgraphs via neighborhood retrieval:

- **SimpleQuestions**: 1-hop neighborhood of topic entities.
- **CompWebQ**: We adopt the method of Shi et al. (2021) to retrieve subgraphs for each question.
- **WebQSP & WebQuestions**: Bidirectional 2-hop expansion with reverse relations added to enhance connectivity.

The detailed information of the datasets is shown in Table 1.

### Baselines

**Prompting Baselines.** We compare with:

- **IO** (Patel et al. 2023): Direct input-output prompting.
- **CoT** (Wei et al. 2022): Chain-of-thought prompting.

Dataset	Total	Train	Test
WebQSP	4637	2998	1639
CompWebQ	31154	27623	3531
SimpleQuestions	97597	75910	21687
WebQuestions	5810	3778	2032

Table 1: Dataset statistics.

- **SC** (Wang et al. 2022): Self-consistency with multiple reasoning paths.

**Retrieval-Augmented Baselines.** We selected retrieval-augmented methods for each benchmark, including GNN-RAG (Mavromatis and Karypis 2024) that leverages graph networks for context-aware retrieval, DiFaR (Baek et al. 2023) with a dynamic factual alignment mechanism, CBR (Das et al. 2021) utilizing analogical reasoning over knowledge snippets, and FiE (Kedia, Zaidi, and Lee 2022) that integrates knowledge through encoder modifications.

**Knowledge Base Question Answering Baselines.** We compare with representative traditional knowledge base question answering models (non-LLM-based) including UniKGQA (Jiang et al. 2023) and RNG (Ye et al. 2022).

**LLM+KG Baseline.** We take KG-CoT (Zhao et al. 2024b) as the key LLM+KG baseline, as it integrates knowledge graphs with chain-of-thought reasoning. Our RFKG-CoT builds upon KG-CoT with two critical improvements: During graph reasoning, we introduce relation activation information to optimize hop-count selection, enabling the model to flexibly handle semantically equivalent information by prioritizing relevant relations. We incorporate few-shot examples during LLM reasoning to enhance the model’s ability to utilize extracted reasoning paths.

### Implementation Setting

We incorporated reversed relations into the relation graph. For experimental settings, the number of reasoning steps T was set to 2 for WebQSP and WebQuestions, while for SimpleQuestions and CompWebQ, it was set to 1 and 3, respectively. In all experiments, we used the pre-trained ‘bert-base-uncased’ model (Devlin et al. 2019) as the question encoder and fine-tuned its parameters. We trained the graph reasoning model using the RAdam optimizer for 60 epochs with a learning rate of 1e-3 (for graph layers) and 1e-5 (for BERT parameters). For large language models, we conducted evaluations by calling ChatGPT and GPT-4 via the OpenAI API, selecting ‘gpt-3.5-turbo’ and ‘gpt-4’ as the backbones of the large language models and using the default settings of the OpenAI API. Specifically, for each question, we generated one reasoning path for each of the top-10 candidate answer entities, and used these paths together with the question as input to enable the large language models to generate answers for evaluation directly.

### Main Results

As shown in Table 2, RFKG-CoT achieves strong performance across knowledge-intensive QA benchmarks. On

Model	AccessKB	Multi-hop QA		Single-hop QA	Open-domain QA
		WebQSP	CompWebQ	SimpleQuestions	WebQuestions
ChatGPT+IO (Patel et al. 2023)	×	63.3	37.6	20.0	48.7
ChatGPT+CoT (Wei et al. 2022)	×	62.2	38.8	20.3	48.5
ChatGPT+SC (Wang et al. 2022)	×	61.1	45.4	18.9	50.3
Previous RA	✓	90.7 <sup>α</sup>	<b>70.4<sup>β</sup></b>	85.8 <sup>γ</sup>	56.3 <sup>δ</sup>
Previous KBQA	✓	76.6 <sup>ε</sup>	52.2 <sup>ε</sup>	71.1 <sup>η</sup>	–
ChatGPT+KG-CoT (Zhao et al. 2024b)	✓	82.1	51.6	77.8	66.5
GPT-4+KG-CoT (Zhao et al. 2024b)	✓	84.9	62.3	86.1	68.0
<b>ChatGPT+RFKG-CoT (ours)</b>	✓	89.9	61.4	79.2	<b>77.0</b>
<b>GPT-4+RFKG-CoT (ours)</b>	✓	<b>91.5</b>	65.1	<b>87.0</b>	<b>78.2</b>

Table 2: Accuracy Comparison between RFKG-CoT and Other Baseline Methods(e.g.,  $\alpha$ : GNN-RAG (Mavromatis and Karypis 2024),  $\beta$ : CBR (Das et al. 2021),  $\gamma$ : DiFaR (Baek et al. 2023),  $\delta$ : FiE (Kedia, Zaidi, and Lee 2022),  $\epsilon$ : UniKGQA (Jiang et al. 2023), and  $\eta$ : RNG (Ye et al. 2022) ) on Multi-hop QA, Single-hop QA, Open-domain QA.

Method	WebQSP	CompWebQ
Llama2-7B		
KG-CoT	72.4	46.7
RFKG-CoT	<b>87.1</b>	<b>51.9</b>
$\Delta$	(+14.7)	(+5.2)
Llama2-13B		
KG-CoT	74.6	50.0
RFKG-CoT	<b>88.7</b>	<b>55.5</b>
$\Delta$	(+14.1)	(+5.5)
ChatGPT		
KG-CoT	82.1	51.6
RFKG-CoT	<b>89.9</b>	<b>61.4</b>
$\Delta$	(+7.8)	(+9.8)
GPT-4		
KG-CoT	84.9	62.3
RFKG-CoT	<b>91.5</b>	<b>65.1</b>
$\Delta$	(+6.6)	(+2.8)

Table 3: Performance comparison (%) on WebQSP and CompWebQ across LLM architectures.

multi-hop QA tasks, it performs well on WebQSP (91.5%) and WebQuestions (78.2%), surpassing GPT4+KG-CoT by 6.6 and 10.2 points, respectively. For CompWebQ, the GPT-4 variant reaches 65.1%, significantly outperforming KG-CoT (62.3%) despite subgraph coverage limitations. In single-hop scenarios, it performs strongly on SimpleQuestions (87%), with the modest gain partly attributed to two factors: the hop-count selector (designed for multi-hop tasks) has limited utility here, and entity IDs in generated paths (lacking semantic alignment) may have constrained further improvements. Notably, when evaluating only answerable CompWebQ queries (2,611 samples), RFKG-CoT reaches 88.1% accuracy, which is a 3.8 pp improvement over KG-CoT (84.3%), demonstrating its robustness to subgraph incompleteness. This confirms that performance gaps

primarily stem from knowledge coverage constraints rather than limitations in reasoning capability.

Method	WebQSP	CompWebQ
KG-CoT (Baseline)	82.1	51.6
+ Relation Activation Mask Only	85.5	59.8
+ Few-Shot Guidance Only	87.7	57.8
<b>RFKG-CoT (Full)</b>	89.9	61.4

Table 4: Ablation study on WebQSP and CompWebQ with ChatGPT for verifying improved components (Accuracy %).

### Robustness across LLM Architectures

To assess the generalizability of RFKG-CoT, we evaluate its performance with diverse LLMs spanning open-source (Llama2-7B, Llama2-13B) and commercial APIs (ChatGPT, GPT-4). As shown in Table 3, RFKG-CoT consistently enhances KG-CoT across all configurations with two key findings: First, we observe an inverse scaling relationship where smaller LLMs exhibit substantially larger relative gains (e.g., Llama2-13B achieves +14.1 pp on WebQSP and +5.5 pp on CompWebQ) compared to advanced models like GPT-4 (+6.6 pp/+2.8 pp). This pattern emerges because RFKG-CoT’s path guidance compensates more effectively for parametric knowledge gaps in capacity-constrained models. Second, performance plateaus approach the intrinsic limits of our graph reasoning module: analysis shows the correct answer appears in generated reasoning paths for 92.3% of WebQSP. Similar constraints apply to CompWebQ. Thus, while LLM scaling improves absolute performance, the remaining gap is primarily constrained by the knowledge grounding component rather than LLM capabilities.

### Ablation study

To verify the effectiveness of the improved components, we conducted ablation studies on the WebQSP and CompWebQ datasets, using ChatGPT for evaluation. Table 4 compares the full model against KG-CoT and single-component vari-

ants. The ‘‘Relation Activation Mask Only’’ variant outperforms KG-CoT by 3.4 pp on WebQSP but 8.2 pp on CompWebQ. This larger gain on CompWebQ indicates the mask’s strength in improving path quality by prioritizing relevant relations, particularly valuable for more complex reasoning scenarios. In contrast, ‘‘Few-Shot Guidance Only’’ yields balanced gains (+5.6 pp on WebQSP, +6.2 pp on CompWebQ), enhancing path utilization by teaching LLMs to interpret structured paths. Unlike the mask (optimizing the selection of reasoning hops), few-shot guidance optimizes how paths are used. The full RFKG-CoT achieves the highest accuracy (89.9%/61.4%). This synergy arises because high-quality paths (from the mask) are better utilized via few-shot guidance, forming a closed loop. Both components are critical for robust reasoning.

### Hyperparameter Sensitivity Analysis

To investigate how key input parameters affect performance, we analyze the sensitivity of RFKG-CoT to two critical hyperparameters consistent with the settings in Zhao et al. (2024b) for direct comparability:  $K$  (number of candidate answer entities, tested values: 5, 10, 15) and  $N$  (reasoning paths per entity, tested values: 1, 5, 10).

Experiments are conducted on WebQSP using ChatGPT. As shown in the left part of Figure 4, both models exhibit improved accuracy with increasing  $K$ , following a trend of ‘‘early rapid gain, late saturation’’, which is consistent with KG-CoT’s observations. Expanding  $K$  from 5 to 10 boosts performance for both models, driven by increased coverage of valid answer entities in reasoning paths. However, further increasing  $K$  to 15 yields marginal gains, as lower-ranked candidates contribute minimally to reasoning due to reduced confidence and semantic relevance.

The right part of Figure 4 characterizes the effect of the number of reasoning paths per candidate answer entity. Consistent with KG-CoT’s findings, both models exhibit minimal performance gains as  $N$  increases. This trend stems from a shared mechanism: paths are selected by confidence scores, where the top-ranked paths already capture the most critical reasoning signals. As  $N$  increases, incremental paths inherently have lower confidence, contributing marginally to LLM inference.

### Impact of Few-Shot Example Quantity

To investigate how the number of few-shot examples affects reasoning performance, we conducted a controlled experiment on WebQSP using ChatGPT+RFKG-CoT, with the baseline configuration  $E = 0$ . We then systematically increased the number of examples from  $E = 1$  to  $E = 5$  to evaluate performance sensitivity to guidance quantity.

As shown in Figure 5, RFKG-CoT’s accuracy exhibits a distinct non-monotonic relationship with  $E$ . The baseline  $E = 0$  demonstrates competent performance using only relation-driven path selection. Adding minimal guidance ( $E = 1$ ) yields substantial gains, with progressive improvement up to the  $E = 3$  optimum. Beyond this optimum, additional examples progressively degrade results, revealing a critical trade-off between guidance completeness and cognitive load. This pattern aligns with transformer architecture

constraints, validating our default three-example design and confirming that calibrated few-shot guidance complements relation-driven path selection.




Figure 4: Effect of the number of candidate answer entities (left) and paths per entity (right) on performance in WebQSP.




Figure 5: Accuracy sensitivity to few-shot example quantity on WebQSP (ChatGPT). The dashed line indicates the baseline (RFKG-CoT without few-shot examples).

## Conclusion

We presented RFKG-CoT, a knowledge-aware QA framework that fuses symbolic graph reasoning with LLMs. It improves from previous work like KG-CoT via (1) a relation-driven hop selector that records activated relations with a mask and dynamically chooses the right reasoning depth, and (2) a few-shot path prompt that guides the LLM to ground each retrieved path.

On WebQSP, CompWebQ, SimpleQuestions and WebQuestions, RFKG-CoT consistently outperforms prior work, boosting accuracy by up to 14.7 pp. Gains are largest on smaller open-source LLMs, underlining the value of explicit external knowledge. Ablations show the hop mask and path prompt each help independently and, when combined, reinforce one another by selecting better paths and exploiting them more effectively.

Beyond performance, RFKG-CoT produces interpretable reasoning traces and naturally supports incremental KG updates without retraining the LLM, making it attractive for real-world, rapidly evolving domains such as legal or biomedical search.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62376178), and the Priority Academic Program Development of Jiangsu Higher Education Institutions.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Leoni Aleman, F.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. GPT-4 Technical Report. *arXiv e-prints*, arXiv-2303.
- Baek, J.; Aji, A. F.; Lehmann, J.; and Hwang, S. J. 2023. Direct Fact Retrieval from Knowledge Graphs without Entity Linking. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 10038–10055.
- Berant, J.; Chou, A.; Frostig, R.; and Liang, P. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 1533–1544.
- Bollacker, K.; Cook, R.; and Tufts, P. 2007. Freebase: A shared database of structured general human knowledge. In *AAAI*, volume 7, 1962–1963.
- Das, R.; Zaheer, M.; Thai, D.; Godbole, A.; Perez, E.; Lee, J.-Y.; Tan, L.; Polymenakos, L.; and McCallum, A. 2021. Case-based Reasoning for Natural Language Queries over Knowledge Bases. *arXiv e-prints*, arXiv-2104.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Dong, Q.; Li, L.; Dai, D.; Zheng, C.; Ma, J.; Li, R.; Xia, H.; Xu, J.; Wu, Z.; Chang, B.; et al. 2024. A Survey on In-context Learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 1107–1128.
- Han, X.; Zhu, H.; Yu, P.; Wang, Z.; Yao, Y.; Liu, Z.; and Sun, M. 2018. FewRel: A Large-Scale Supervised Few-Shot Relation Classification Dataset with State-of-the-Art Evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4803–4809.
- Hu, L.; Liu, Z.; Zhao, Z.; Hou, L.; Nie, L.; and Li, J. 2023. A survey of knowledge enhanced pre-trained language models. *IEEE Transactions on Knowledge and Data Engineering*, 36(4): 1413–1430.
- Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y. J.; Madotto, A.; and Fung, P. 2023. Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12): 1–38.
- Jiang, J.; Zhou, K.; Zhao, X.; and Wen, J.-R. 2023. UniKGQA: Unified Retrieval and Reasoning for Solving Multi-hop Question Answering Over Knowledge Graph. In *The Eleventh International Conference on Learning Representations*.
- Kedia, A.; Zaidi, M. A.; and Lee, H. 2022. FiE: Building a Global Probability Space by Leveraging Early Fusion in Encoder for Open-Domain Question Answering. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4246–4260.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474.
- Li, T.; Ma, X.; Zhuang, A.; Gu, Y.; Su, Y.; and Chen, W. 2023. Few-shot in-context learning for knowledge base question answering. *arXiv preprint arXiv:2305.01750*.
- Liu, W.; Zhou, P.; Zhao, Z.; Wang, Z.; Ju, Q.; Deng, H.; and Wang, P. 2020. K-bert: Enabling language representation with knowledge graph. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 2901–2908.
- Liu, Y.; Wan, Y.; He, L.; Peng, H.; and Yu, P. S. 2021. Kg-bart: Knowledge graph-augmented bart for generative commonsense reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 6418–6425.
- Mavromatis, C.; and Karypis, G. 2024. GNN-RAG: Graph Neural Retrieval for Large Language Model Reasoning. *ICLR*.
- Patel, A.; Li, B.; Rasooli, M. S.; Constant, N.; Raffel, C.; and Callison-Burch, C. 2023. Bidirectional Language Models Are Also Few-shot Learners. In *The Eleventh International Conference on Learning Representations (ICLR 2023)*.
- Petrochuk, M.; and Zettlemoyer, L. 2018. SimpleQuestions Nearly Solved: A New Upperbound and Baseline Approach. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 554–558.
- Shi, J.; Cao, S.; Hou, L.; Li, J.; and Zhang, H. 2021. TransferNet: An Effective and Transparent Framework for Multi-hop Question Answering over Relation Graph. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 4149–4158.
- Talmor, A.; and Berant, J. 2018. The Web as a Knowledge-Base for Answering Complex Questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 641–651.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Wang, Y.; Ma, X.; and Chen, W. 2023. Augmenting black-box llms with medical textbooks for clinical question answering. *CoRR*.
- 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*, 35: 24824–24837.

- Yang, R.; Xue, H.; Razzak, I.; Hacid, H.; and Salim, F. D. 2025. Beyond Single Pass, Looping Through Time: KG-IRAG with Iterative Knowledge Retrieval. *arXiv preprint arXiv:2503.14234*.
- Ye, X.; Yavuz, S.; Hashimoto, K.; Zhou, Y.; and Xiong, C. 2022. RNG-KBQA: Generation Augmented Iterative Ranking for Knowledge Base Question Answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6032–6043.
- Yih, W.-t.; Richardson, M.; Meek, C.; Chang, M.-W.; and Suh, J. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 201–206.
- Zhao, R.; Zhao, F.; Hu, L.; and Xu, G. 2024a. Graph reasoning transformers for knowledge-aware question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19652–19660.
- Zhao, R.; Zhao, F.; Wang, L.; Wang, X.; and Xu, G. 2024b. Kg-cot: Chain-of-thought prompting of large language models over knowledge graphs for knowledge-aware question answering. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, 6642–6650. International Joint Conferences on Artificial Intelligence.
- Zhu, X.; Xie, Y.; Liu, Y.; Li, Y.; and Hu, W. 2025. Knowledge Graph-Guided Retrieval Augmented Generation. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8912–8924. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-189-6.