

# Question-guided Knowledge Graph Re-scoring and Injection for Knowledge Graph Question Answering

Yu Zhang<sup>1</sup>, Kehai Chen<sup>1\*</sup>, Xuefeng Bai<sup>1</sup>, Zhao Kang<sup>2</sup>, Quangjiang Guo<sup>2</sup>, Min Zhang<sup>1</sup>

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

<sup>2</sup>University of Electronic Science and Technology of China, Chengdu, China

yuzhang2717@gmail.com, {chenkehai, baixuefeng, zhangmin2021}@hit.edu.cn,

zkang@uestc.edu.cn, guochance1999@163.com

## Abstract

Knowledge graph question answering (KGQA) involves answering natural language questions by leveraging structured information stored in a knowledge graph. Typically, KGQA initially retrieve a targeted subgraph from a large-scale knowledge graph, which serves as the basis for reasoning models to address queries. However, the retrieved subgraph inevitably brings distraction information for knowledge utilization, impeding the model’s ability to perform accurate reasoning. To address this issue, we propose a Question-guided Knowledge Graph Re-scoring method (Q-KGR) to eliminate noisy pathways for the input question, thereby focusing specifically on pertinent factual knowledge. Moreover, we introduce Knowformer, a parameter-efficient method for injecting the re-scored knowledge graph into large language models to enhance their ability to perform factual reasoning. Extensive experiments on multiple KGQA benchmarks demonstrate the superiority of our method over existing systems. <sup>1</sup>

## 1 Introduction

Knowledge graph question answering (KGQA) aims to provide factual answers to natural language questions by leveraging information from knowledge graphs (Cheng et al., 2024; Jiang et al., 2023; Zhu et al., 2024b). Sitting at the intersection between graphs and texts, this task can further facilitate the applicability of knowledge graphs in more downstream tasks, such as recommendation systems (Damianou et al., 2024; Yao et al., 2024), information extraction (Zhang and Kang, 2024; Zhu et al., 2024a), time series forecasting (Yu et al., 2024; Gao et al., 2024; Gui et al., 2024).

Recent work on KGQA attempts to synergize large language models (LLMs) with knowledge

\*Corresponding Author

<sup>1</sup>The code and data are released on <https://github.com/EchoDreamer/Q-KGR>.

## QA context

This is the most famous bank where children keep their assets

A. repository B. bank vault C. piggybank\* D. pig

## Retrieved subgraph




Figure 1: Given the QA context (question and answers), we retrieve a subgraph including question entities (red), answer entities (blue) and other neighbor entities (grey). The correct answer is piggybank.

graphs (KGs) by utilizing retrieved knowledge from KGs to assist LLMs in conducting factual reasoning (Luo et al., 2023; Zhang et al., 2023; Sun et al., 2023). A prevailing approach is to heuristically retrieve a subgraph from a large-scale KG using topic entities (KG entities mentioned in the given question and candidate answers) and their few-hop neighbors, and then use this subgraph as the foundation for LLMs to perform reasoning. Despite the promising performance, this approach heavily rely on the retrieved graph which inevitably includes distracting information. This might leads the LLM to focus on irrelevant knowledge, thereby hindering its ability to conduct accurate reasoning. Figure 1 depicts an example from the RiddleSense dataset (Lin et al., 2021). In the retrieved subgraph, the question entity *bank* connects a distracting answer entity *repository* with two noisy pathways: (I) *bank*→*river bank*→*human*→*repository* and (II) *bank*→*robber*→*repository*. Directly incorporating knowledge from these paths into the LLM would significantly distort its reasoning and judgment. However, these noisy paths can be easily identified and excluded by considering and analysing the

input question. For example, we can exclude the noisy path (I) by considering that the question is irrelevant to "river". In addition, the noisy path (II) can be excluded by taking into account the semantic meaning more deeply such as a specific restriction phrase, "where children keep their assets".

Building on the insights from above analyses, this paper introduces a question-guided knowledge graph re-scoring method (**Q-KGR**) to eliminate the negative effect of noisy pathways for specific questions, thereby focusing attention on precise factual knowledge. Specifically, we develop an edge re-scoring method to assign a relevance score to each edge by computing the semantic similarity between the specific question and each edge triplet to obtain a re-scored knowledge graph. And then we extend the graph attention network to incorporate edge scores during feature aggregation. In addition, we design **Knowformer**, a customized transformer where a knowledge projector is utilized to align and inject the information of knowledge graphs into the standard transformer architecture, thereby enhancing its reasoning capabilities.

Furthermore, Knowformer can naturally adapt to LoRA (Hu et al., 2021), and the combined use can effectively achieve task adaptation and knowledge injection. We conduct extensive experiments on four KGQA benchmarks, and the experimental results demonstrate that our method successfully eliminates noisy pathways and injects knowledge into LLM to perform better factual reasoning. Our main contributions are summarized as follows:

- We propose Q-KGR to re-score the edges of knowledge graph and introduce Knowformer to align and inject structured knowledge into LLM, thereby assisting it in conducting accurate factual reasoning.
- Extensive experiments on multiple KGQA benchmarks demonstrate the superiority of our method across different LLM architectures, sizes and setups.
- Further experimental analysis indicates that our approach effectively guides the LLM to concentrate on accurate factual knowledge and achieves parameterized knowledge alignment and injection.

## 2 Related Work

**Knowledge Graph Question Answering** Knowledge Graph Question Answering (KGQA) focuses on enabling machines to retrieve and utilize relevant knowledge from a KG (Speer et al., 2017; Bollacker et al., 2008). This task requires models to excel in capturing accurate factual knowledge and complex reasoning abilities.

Capturing accurate factual knowledge requires maximizing recall during retrieval, followed by meticulous screening to filter out irrelevant information. Many methods have embarked on this journey (Oguz et al., 2022; Jiang et al., 2022b; Li et al., 2023). Lin et al. 2019 and Zhang et al. 2022b propose to heuristically retrieve a subgraph from the KG for reasoning over KG. Similar to our method, Yasunaga et al. 2021 introduces node relevance scores to filter irrelevant nodes. However, edges including subject, relation and object in KG capture richer information about interactions between entities. By comparison, using edge scores can provide more comprehensive views which in turn enhances the model's ability to accurately capture and utilize relevant pathways for more precise and meaningful answers. Besides, Jiang et al. 2022a and Wang et al. 2021 indicate that node relevance scores brings weak or even redundant effects in model optimization and our subsequent experiments also validate prior conclusions. On the other hand, because LLMs have demonstrated remarkable reasoning abilities on various tasks (Lu et al., 2022; Chen et al., 2024a; Huang and Chang, 2023; Chen et al., 2024b; Zhang et al., 2024), many methods apply them to conduct reasoning tasks on KG (Luo et al., 2023; Zhang et al., 2023; Tian et al., 2024). Our method simultaneously achieves capturing accurate factual knowledge and evoking complex reasoning skills.

**Knowledge Editing and Injecting for Large Language Model** In the field of knowledge editing, many works have explored how factual knowledge is stored in LLM (Meng et al., 2022a,b; Yao et al., 2023). Through causal analysis and quantitative experiments, Geva et al., 2021 and Dai et al., 2022 demonstrates that parameterized knowledge in LLM is stored as key-value memory pairs in the feed forward Neural Network (FFN) of the transformer. Building on this, Yao et al., 2022 is the first to attempt injecting text knowledge into parameter space of language models (LM). Dong et al., 2022 constructed an external knowledge

vector database to calibrate factual knowledge in LLM. By comparison, we utilize GNNs to model and inject structured knowledge into LLMs in a layer-wised manner.

### 3 Preliminary

**Knowledge Graph** We define knowledge graph (KG) as a multi-relation graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  where  $\mathcal{V}$  is the set of entity nodes in the KG,  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$  is the set of edges that connect nodes in  $\mathcal{V}$  and  $\mathcal{R}$  represents the set of relation types in the KG.

**Knowledge Retrieval** Given a question  $q$  and all candidate answers  $\mathcal{A}_q$ , we follow the previous work (Lin et al., 2019) to associate the entities mentioned in the question and all candidate answers from  $\mathcal{G}$ . We use  $\mathcal{V}_{qa} := \mathcal{V}_q \cup \mathcal{V}_a$  to represent all the entities that appear in the question and answers, named topic entities. Then, we retrieve inner node entities  $\mathcal{V}_{in}$  and corresponding edges  $\mathcal{E}_{in}$  from KG on the  $k$ -hop paths among all topic entities.

**Knowledge Graph Question Answering (KGQA)** is a typical reasoning task based KG. The task aims to design a function  $f$  to predict correct answer  $a \in \mathcal{A}_q$  based knowledge from  $\mathcal{G}$  i.e.,  $a = f(q, \mathcal{G})$ .

## 4 Methodology

The overall framework of the proposed method is illustrated in Figure 2. Firstly, we employ an edge re-scoring to assign a relevance score to each edge based on its relevance to the question (§4.1). Additionally, we modify GAT to encode the re-scored subgraph, obtaining structured knowledge representations at each layer (§4.2). Furthermore, we develop a layer-wised knowledge injection method to inject knowledge from last  $M$  layers in GNNs into corresponding Knowformer layers (§4.3). Finally, we process the question text sequentially through  $N$  layers of the transformer and  $M$  layers of Knowformer to predict the probability distribution of the candidate answers.

### 4.1 Question-guided Knowledge Graph Re-scoring

For eliminating the noisy pathways, we introduce a question-guided edge re-scoring to assign each triplet a relevance score by computing the semantic similarity between the specific question and each edge. Specifically, we sample each node pair (subject and object) from the subgraph and feed

them into a pre-trained language model (PLM) to gain the embeddings as  $e_s$  and  $e_o$ :

$$e_s, e_o = PLM(\text{subject}, \text{object}). \quad (1)$$

To leverage the core semantic information in question  $q$ , we extract entity representations from encoded question by the same PLM and aggregate them through an average pooling layer as  $e_q$ . Given the computational constraints posed by the large number of edges in the original graph, we keep the parameters of the PLM frozen to ensure efficient processing. Finally, we utilize a bilinear layer to estimate the relevance score for each edge as:

$$\eta = \text{Normalize}(\text{Bilinear}([e_s, e_o], e_q)), \quad (2)$$

Regarding normalization, we investigate two strategies: (1) Gumbel-max (Jang et al., 2016): modeling predicted score for an edge as a hard (one-hot) label,  $\eta \in \mathbb{R}^{2 \times 1}$  indicating whether the edge exists. (2) Gumbel-softmax (Jang et al., 2016): modeling the score as a soft label, representing the score of relevance and irrelevance. Since there is no gold graph structure as supervised signal to train the proposed edge re-scoring individually, we incorporate the predicted score as a latent variable to train an end-to-end model. Moreover, to address the issue of gradient propagation when using hard labels, we further refine the one-hot label as follows:

$$\eta = \eta + \text{one\_hot}(\eta) - sg(\eta) \quad (3)$$

where  $sg$  stands for the "stop gradient" operator (also called `.detach()` in PyTorch) which does not affect the forward pass, but in the backward pass it does not back-propagate the gradient to its input, meaning it is treated as a constant. We find that both strategies are effective, but in the main experiments, we use the soft label.

### 4.2 Re-scored Knowledge Graph Modeling

In order to better aggregate and extract extensive knowledge from knowledge graph, we first expand the subgraph by adding a *global node* as  $\mathcal{V}_g$  connecting to all topic nodes. Based on this, we extend the graph attention network (GAT) (Veličković et al., 2017) to encode the re-scored knowledge graph, which induces node representations via iterative message passing between neighbors on the graph.




Figure 2: The overall framework of our method. Given a retrieved subgraph from origin KG, we revise it through question-guided knowledge graph re-scoring (Q-KGR, I). Then we utilize question text and revised subgraph to conduct knowledge modeling, injection and reasoning (II). Knowformer (III) consists of self-attention layers and a customized feed-forward network (FFN) layer. In the FFN layer, *key* and *value* represent original weight matrices,  $LoRA_k$  and  $LoRA_v$  correspond to LoRA weight parameters, and  $\phi_k$  and  $\phi_v$  are the knowledge vectors mapped from graph latent space to the parameter space of FFN. Gray denotes frozen model parameters, while green indicates updated model parameters.

In the  $l + 1$  layer, we update the representation  $h_t^{l+1} \in \mathbb{R}^D$  of each node  $\mathcal{V}_t$  by

$$h_t^{(l+1)} = f_m \left( \sum_{s \in N_t \cup \{t\}} \alpha_{st} m_{st} \right) + h_t^{(l)} \quad (4)$$

where  $f_m : \mathbb{R}^D \rightarrow \mathbb{R}^D$  is a two-layer MLP,  $N_t$  denotes the neighbors of node  $\mathcal{V}_t$ ,  $m_{st} \in \mathbb{R}^D$  and  $\alpha_{st}$  represent the information and weight from each neighbor node  $\mathcal{V}_s$  to node  $\mathcal{V}_t$ . And  $\alpha_{st}$  is calculated as:

$$\alpha_{st} = \frac{\eta \cdot \exp(\gamma_{st})}{\sum_{s' \in N_t \cup \{t\}} \exp(\gamma_{s't})}, \gamma_{st} = \frac{q_s^T k_t}{\sqrt{D}} \quad (5)$$

where  $q_s$  and  $k_t$  correspond to the query and key in the GAT, obtained from node and relation type embeddings through multiple MLPs<sup>2</sup> and we use edge relevance score  $\eta$  from Section 4.1 to scale the original attention weights. For each node, we use the same PLM as section 4.1 to initialize the representation, and we initialize the global node as 0.

<sup>2</sup>The calculation method for  $m_{st}, q_s$  and  $k_t$  is based on Yasunaga et al. 2021.

Finally we obtain the embeddings of all nodes in each layer and for last  $M$  layers, we use the representation of the global node,  $h_l$  as external graph knowledge. We will align and inject them into LLMs in section 4.3.

### 4.3 Knowformer

In a LLM, we use  $M$  Knowformer layers to replace the original Transformer layers that are closest to the LLM’s final output for structured knowledge injection into LLM to assist in conducting factual reasoning. For a single Knowformer layer, it inherits the parameters of the original Transformer layer at the corresponding position. A standard transformer consists of a multi-head self-attention layer and a feed-forward network (FFN) layer, typically comprising two or three linear layers with an activation function. Supposing that the input of FFN is  $x$  and the output of FFN is computed as:

$$FFN(x) = (x \cdot key) \cdot value \quad (6)$$

where  $key, value \in \mathbb{R}^{d_m \times d}$  are parameter matrices of the first and the second layer of



LLM	Method	OBQA	Riddle	ARC	PIQA	Average
FLAN-T5-XL	LLM-only	69.20	53.73	68.24	58.43	62.24
	REL	61.80	43.33	64.12	57.56	57.71
	KAPING TH	58.80	40.78	63.52	52.34	53.86
	KAPING OH	60.00	41.37	63.09	51.69	54.04
	Prompt Tuning	72.20	53.33	70.64	60.83	64.25
	Full Fine-tuning	82.8	74.12	73.30	63.55	73.44
	LoRA	80.40	72.94	71.33	63.76	72.11
	LoRA+GNP	83.40	75.49	72.45	64.31	73.91
	LoRA+Ours	<b>85.00</b>	<b>78.43</b>	<b>73.38</b>	<b>87.03</b>	<b>80.96</b>
FLAN-T5-XXL	LLM-only	76.8	61.37	68.93	56.58	65.92
	REL	72.80	53.53	66.78	56.80	62.48
	KAPING TH	60.60	48.43	57.25	53.21	54.87
	KAPING OH	60.00	47.65	56.65	51.69	54.00
	Prompt Tuning	78.80	61.37	74.85	61.26	69.07
	Full Fine-tuning	89.40	80.78	76.82	65.61	78.15
	LoRA	88.60	74.90	78.54	65.61	76.91
	LoRA+GNP	89.60	76.67	78.71	65.94	77.73
	LoRA+Ours	<b>90.00</b>	<b>80.98</b>	<b>79.78</b>	<b>90.08</b>	<b>85.21</b>

Table 1: Overall experimental results on commonsense reasoning tasks. The best results across different LLM sizes and setups are highlighted in bold. Accuracy is used as the evaluation metric.

FFN, respectively.<sup>3</sup> As shown in Figure 2, the Knowformer modifies the FFN by coupling *key, value* with aligned structured knowledge,  $\phi_k, \phi_v$ . The computation within the modified FFN can be described as follows:

$$FFN(x) = (x \cdot [key : \phi_k]) \cdot [value : \phi_v]. \quad (7)$$

To inject the structured knowledge, we first obtain the structured knowledge representation  $h \in \mathbb{R}^{1 \times D}$  for a certain layer from section 4.2. Subsequently, we align it with parameterized knowledge in FFN using two linear projectors:

$$\phi_k = Pr_k h, \phi_v = Pr_v h, \quad (8)$$

where  $Pr_k, Pr_v : \mathbb{R}^D \rightarrow \mathbb{R}^d$  are two parameter matrices and  $d$  is the input dimension of FFN. Inspired by Hu et al. 2021, we use random Gaussian distribution and zero to initialize  $Pr_k$  and  $Pr_v$  respectively. It should be noted that our method is highly flexible, seamlessly integrating with LoRA, and its combined use effectively facilitates task adaptation, knowledge alignment, and injection. Finally, our method is architecture-agnostic. As shown in the Figure 2, we illustrate general Transformer layers and do not

<sup>3</sup>We simplify the FFN layer by omitting the activation layer and bias and take the FFN composed of two linear layers as an example.

distinguish architecture of LLM e.g., encoder-decoder or decoder-only. Our knowledge injection is performed in the several layers that are nearest to the LLM’s final output.

## 5 Experiments

In this section, we conduct extensive experiments on four KGQA benchmarks including OpenBookQA (OBQA) (Mihaylov et al., 2018), AI2 Reasoning Challenge (ARC) (Clark et al., 2018), RiddleSense (Riddle) (Lin et al., 2021) and Physical Interaction Question Answering (PIQA) (Bisk et al., 2020) to compare the performances of different methods. We use ConceptNet (Speer et al., 2017), a general-domain knowledge graph as our structured knowledge source  $\mathcal{G}$ . A detailed explanation of the datasets and the knowledge graph is provided in the appendix A.

### 5.1 Baselines

We use encoder-decoder architecture LLM, FLAN-T5-XL (3B) and FLAN-T5-XXL (11B) (Chung et al., 2024) as the backbone to conduct experiments. We compare eight baselines including **LLM-only**, KG Flattening that flattens the nodes in the graph into a sequence via relevance score (**REL**) ranking (Yasunaga et al., 2022), **KAPING** (Baek et al., 2023) that injects the important KG

LM	Method	OBQA
RoBERTa-Large	Full fine-tuning	64.80
	QAGNN	67.80
	GreaseLM	66.90
	DRLK	70.20
	GSC	70.33
	Rumination	70.30
	Ours	<b>71.00</b>
LLaMA2-7B	LLM-only	31.80
	Full Fine-tuning	74.00
	LoRA	71.80
	LoRA-Ours	75.40
	Full Fine-tuning-Ours	<b>77.00</b>

Table 2: Results on OBQA dataset using Encoder-only (RoBERTa-Large) and Decoder-only (LLaMA2-7B) architectures.

triples within one-hop (OH) and two-hop (TH) neighborhoods, **Prompt Tuning** (Lester et al., 2021) that regards embedding knowledge (soft prompts) as prefix of the text question after embedding layer in LLM, **LoRA** (Hu et al., 2021) that updates partial LLM parameters, **Full Fine-Tuning** and **GNP** (Tian et al., 2024) which uses GNN to model retrieved subgraph as soft prompts to argument LLM. Besides, we also compare other baselines using encoder-only model (RoBERTa-Large) (Zhuang et al., 2021) and decoder-only model (LLaMA2-7B) Touvron et al. 2023. Appendix A provides more details on these baselines and experimental setups.

## 5.2 Main Results

Table 1 compares the results of the proposed method with previous systems on four widely-used benchmark datasets. It can be firstly observed that our method with LoRA achieves the best results compared to other systems. In particular, our approach achieves a substantial breakthrough on the PIQA dataset, surpassing all baselines by a significant margin of absolute 20 percentage points. Additionally, compared to the state-of-the-art system GNP, which uses a GNN to directly encode the retrieved subgraph as soft prompts and inject them into an LLM, our method obtains consistently better results. This achievement underscores the effectiveness of our approach in reconstructing the knowledge graph and injecting knowledge.

We further verify the generalization ability of the proposed method across both encoder-only and decoder-only architectures. Table 2 presents the performance of various systems using

Model	Variant	Riddle		ARC	
		dev	test	dev	test
FLAN-T5-XL	w/o Q-KGR	72.99	73.73	77.26	72.61
	NRS	72.73	74.31	77.26	72.55
	HL	73.91	71.57	76.59	<b>73.81</b>
	w/o PKI	73.02	72.94	74.21	71.33
	SP	72.45	76.30	73.12	72.56
FLAN-T5-XXL	w/o global	<b>74.95</b>	74.31	75.25	72.53
	Ours	74.67	<b>78.43</b>	<b>77.59</b>	73.38
	w/o Q-KGR	78.82	79.61	79.60	78.50
	NRS	78.72	79.31	80.27	78.10
FLAN-T5-XXL	HL	<b>80.82</b>	<b>82.16</b>	79.60	78.75
	w/o PKI	74.02	74.90	77.45	78.54
	SP	74.65	76.42	76.53	78.69
	w/o global	80.63	79.61	76.60	76.35
	Ours	79.84	80.98	<b>80.27</b>	<b>79.78</b>

Table 3: Results of ablation study.

Method	OBQA	CSQA	MedQA
QAGNN	67.80	73.41	38.0
QAGNN-ours	68.56	75.02	38.49

Table 4: Performances using node relevance scoring from Yasunaga et al., 2021 and edge relevance re-scoring on the OBQA (Mihaylov et al., 2018), CSQA (Talmor et al., 2019) and MedQA (Jin et al., 2021) dataset.

RoBERTa-Large and LLaMA2-7B as the underlying architectures. It is evident that our approach surpasses a series of strong baselines within the RoBERTa-Large model architecture. With the LLaMA2-7B model, our method, particularly when integrated with LoRA, exceeds full fine-tuning by 1.4 percentage points. When combined with full fine-tuning, our approach significantly outperforms other methods. Overall, our method has demonstrated exceptional performance across various model architectures, sizes, and training setups, confirming its effectiveness.

## 5.3 Ablation Study

We conduct extensive ablation experiments on Riddle and ARC datasets using FLAN-T5-XL and FLAN-T5-XXL to investigate the effectiveness of three key components in our methods: (1) question-guided knowledge graph re-scoring (**Q-KGR**); (2) Re-scored knowledge graph modeling (**KM**); (3) parameterized knowledge injection (**PKI**).

Targeted at Q-KGR, we design the following variants: removing Q-KGR (**w/o Q-KGR**); using node relevance scores from QAGNN (Yasunaga et al., 2021) to replace Q-KGR (**NRS**); utilizing

LLM	Method	0	10%	20%	30%	40%	50%
<b>FLAN-T5-XL</b>	w/o Q-KGR	73.73	73.30	72.86	71.53	69.25	65.45
	Ours	78.43 (+4.7)	78.44 (+5.14)	78.30 (+5.44)	78.14 (+6.61)	77.54 (+8.29)	76.28 (+10.83)
<b>FLAN-T5-XXL</b>	w/o Q-KGR	79.61	79.53	78.70	76.01	74.31	72.27
	Ours	80.98(+1.37)	80.93(+1.40)	80.56(+1.86)	79.63 (+3.62)	78.52 (+4.21)	77.89 (+5.62)

Table 5: Performances by adjusting the number of distractor nodes on the OBQA datasets.

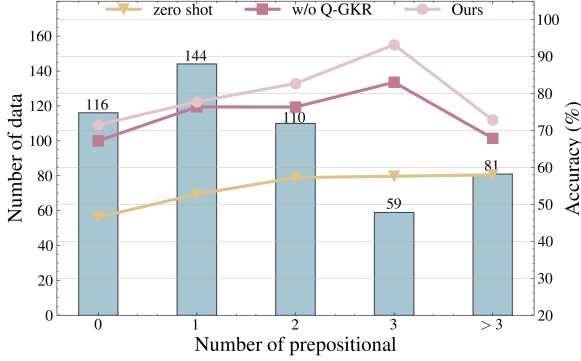


Figure 3: Performance of different methods across divisions varying number of prepositional phrases.

hard labels from Section 4.1 for edge re-scoring (**HL**). Targeted at KM, we compare our method with removing the global node and using the mean over all nodes representations as the graph representation (**w/o global**). Targeted at PKI, we design the following variants: removing PKI (**w/o PKI**); representing the graph knowledge as soft prompt (Tian et al., 2024) and inject it into embedding layer of LLM (**SP**).

As illustrated in Table 3, removing Q-KGR or replacing it with node scoring causes notable declines in the overall performance of the model, particularly on the Riddle dataset, highlighting the indispensable role played by Q-KGR. In addition, we compare the effects of using hard labels versus soft labels (notably, our method ultimately employs soft labels) as discussed in § 4.1 across various setups. It can be observed that soft labels generally leads to superior performance in the majority of cases. Furthermore, when we remove the global node, there is a significant performance drop in the majority of experiments which indicates the crucial role of the global node in knowledge graph modeling at a global level. Lastly, when we remove PKI or instead use SP to inject graph knowledge as soft prompts into the embedding layers of LLM (Tian et al., 2024), the overall results significantly declined compared to our method in all setups.

To further compare Q-KGR with the node relevance scores (NRS) in QAGNN (Yasunaga et al., 2021), we reproduce the method of QAGNN

and replace NRS in QAGNN with our proposed edge relevance scoring to conduct experiments on the OBQA, CSQA (Lin et al., 2019), and MedQA (Jin et al., 2021) dataset in Table 4. The results demonstrate that proposed method has obvious improvements compared to node relevance scoring in QAGNN and can better filter out retrieved irrelevant information.

## 6 Analyses and Discussions

We conduct comprehensive quantitative and qualitative analysis from various perspectives to validate the effectiveness of knowledge graph re-scoring and parameterized knowledge injection.

### 6.1 Knowledge Graph Re-scoring

**Quantitative Analysis** To show the denoising effectiveness of proposed method, we build a controllable setting where we gradually add the irrelevant nodes to the extracted subgraph and observe whether the performance is degrading. Specifically, we limit the number of retrieved nodes to 200 and then collecte distractor nodes that are not in the 1-hop path among topic entities. We replace  $k\%$  of the nodes in the subgraph with distractor nodes, and then manually connect these distractor nodes to the topic nodes as a way to increase noise pathways. We test the performance of the proposed method and w/o Q-KGR on the Riddle dataset, varying  $k$  at 0, 10, 20, 30, 40, and 50, respectively. In Figure 5, we observe that as  $k$  increases, the proposed method shows a greater improvements relative to w/o Q-KGR. From the experimental results, we conclude that our method effectively reduces the noise in retrieving external knowledge, enabling more truthful factual reasoning.

To further validate whether the observed improvements indeed originates from filtering out irrelevant paths using edge re-scoring, we conduct another quantitative analysis. Since there is no gold standard graph clue to measure the number of noisy pathways in the retrieved subgraph, we use the number of prepositional phrases as a proxy. For example, the question "What is that which








Figure 5: With the increase in training steps, it becomes apparent that the injected knowledge gradually integrates into the distribution of the FFN parameters.




Figure 6: As the depth of the layers increases, the injected knowledge gradually integrates into the distribution of the FFN parameters.




Figure 7: Performance by adjusting the layers of knowledge injection, as well as the number of layers in the GAT encoder on the Riddle and OBQA datasets.

Furthermore, we design Knowformer, a customized transformer, through which we align and inject structured knowledge into the parameter space of LLM. Extensive experiments on multiple KGQA datasets across different LLM architectures, sizes and setups demonstrate the superiority of our method. And in-depth analysis experiments indicates that our model is particularly effective in handling complex questions characterized with significant noisy pathways in external knowledge. We also illustrate that our approach achieves effective knowledge alignment and injection for LLM.

## 8 Limitation

Our method improves the model’s knowledge-based question answering by injecting commonsense knowledge but symbolic reasoning in Question Answering remains an open challenge.

From the technical perspective, we extract a representative node from a single layer of GNN and inject it into a Knowformer layer. However, incorporating more nodes intuitively might capture richer knowledge and potentially enhance model performance. Yet, this necessitates further validation regarding graph smoothing issues. Therefore, additional analysis is required to effectively address and balance this concern.

Additionally, aligning with human preferences (Jiang et al., 2024; Liu et al., 2023) during the implementation of Retrieval-Augmented Generation (RAG) is another area that requires further exploration.

## Acknowledgment

We sincerely thank all the reviewers for their valuable and insightful comments, which have greatly enhanced the quality of our work. This work was supported by the National Natural Science Foundation of China (U23B2055 and 62276077), Guangdong Basic and Applied Basic Research Foundation (2024A1515011205), Shenzhen College Stability Support Plan under Grants (GXWD20220811170358002, GXWD20220817123150002 and GXWD2023-1130104007001). The work of Zhao Kang was supported by Qiyuan Lab.

## References

- Jinheon Baek, Alham Fikri Aji, and Amir Saffari. 2023. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In *Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations*, pages 78–106, Toronto, Canada.
- Yonatan Bisk, Rowan Zellers, Ronan Le bras, Jianfeng Gao, and Yejin Choi. 2020. Piqa: Reasoning about physical commonsense in natural language. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7432–7439.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, page 1247–1250, New York, NY, USA.
- Andong Chen, Lianzhang Lou, Kehai Chen, Xuefeng Bai, Yang Xiang, Muyun Yang, Tiejun Zhao, and Min Zhang. 2024a. Benchmarking llms for translating classical chinese poetry: Evaluating adequacy, fluency, and elegance. *arXiv preprint arXiv:2408.09945*.
- Andong Chen, Lianzhang Lou, Kehai Chen, Xuefeng Bai, Yang Xiang, Muyun Yang, Tiejun Zhao, and Min Zhang. 2024b. Dual-reflect: Enhancing large language models for reflective translation through dual learning feedback mechanisms. *arXiv preprint arXiv:2406.07232*.
- Yuhang Cheng, Kaiwen Li, and Zhao Kang. 2024. Emkg: Efficient matchings for knowledge graph integration in stance detection. In *2024 International Joint Conference on Neural Networks*, pages 1–8.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland.
- Andreas Damianou, Francesco Fabbri, Paul Giglioli, Marco De Nadai, Alice Wang, Enrico Palumbo, and Mounia Lalmas. 2024. Towards graph foundation models for personalization. In *Companion Proceedings of the ACM on Web Conference 2024*, page 1798–1802, New York, NY, USA.
- Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui, and Lei Li. 2022. Calibrating factual knowledge in pretrained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5937–5947, Abu Dhabi, United Arab Emirates.
- Xianqiang Gao, Pingrui Zhang, Delin Qu, Dong Wang, Zhigang Wang, Yan Ding, Bin Zhao, and Xuelong Li. 2024. Learning 2d invariant affordance knowledge for 3d affordance grounding. *arXiv preprint arXiv:2408.13024*.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, Online and Punta Cana, Dominican Republic.
- Haokun Gui, Xiucheng Li, and Xinyang Chen. 2024. Vector quantization pretraining for EEG time series with random projection and phase alignment. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pages 16731–16750.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065, Toronto, Canada.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Xin Zhao, and Ji-Rong Wen. 2023. StructGPT: A general framework for large language model to reason over structured data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9237–9251, Singapore.
- Jinhao Jiang, Kun Zhou, Ji-Rong Wen, and Xin Zhao. 2022a. *great truths are always simple* : a rather simple knowledge encoder for enhancing the commonsense reasoning capacity of pre-trained models. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1730–1741, Seattle, United States.
- Jinhao Jiang, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. 2022b. Unikgqa: Unified retrieval and reasoning for solving multi-hop question answering over knowledge graph. *arXiv preprint arXiv:2212.00959*.
- Ruili Jiang, Kehai Chen, Xuefeng Bai, Zhixuan He, Juntao Li, Muyun Yang, Tiejun Zhao, Liqiang Nie, and Min Zhang. 2024. A survey on human preference

- learning for large language models. *arXiv preprint arXiv:2406.11191*.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Nicolas Keriven. 2024. Not too little, not too much: a theoretical analysis of graph over smoothing. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic.
- Shiyang Li, Yifan Gao, Haoming Jiang, Qingyu Yin, Zheng Li, Xifeng Yan, Chao Zhang, and Bing Yin. 2023. Graph reasoning for question answering with triplet retrieval. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3366–3375, Toronto, Canada.
- Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. KagNet: Knowledge-aware graph networks for commonsense reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2829–2839, Hong Kong, China.
- Bill Yuchen Lin, Ziyi Wu, Yichi Yang, Dong-Ho Lee, and Xiang Ren. 2021. RiddleSense: Reasoning about riddle questions featuring linguistic creativity and commonsense knowledge. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1504–1515, Online.
- George C. Linderman, Manas Rachh, Jeremy G. Hoskins, Stefan Steinerberger, and Yuval Kluger. 2019. Fast interpolation-based t-sne for improved visualization of single-cell rna-seq data. *Nature Methods*, 16(3):243–245.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. On the variance of the adaptive learning rate and beyond. *arXiv 2019. arXiv preprint arXiv:1908.03265*.
- Xiao Liu, Hanyu Lai, Hao Yu, Yifan Xu, Aohan Zeng, Zhengxiao Du, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. Webglm: Towards an efficient web-enhanced question answering system with human preferences. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 4549–4560, New York, NY, USA.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Advances in Neural Information Processing Systems*, volume 35, pages 2507–2521.
- Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023. Reasoning on graphs: Faithful and interpretable large language model reasoning. *arXiv preprint arXiv:2310.01061*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. In *Advances in Neural Information Processing Systems*, volume 35, pages 17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium.
- Todor Mihaylov and Anette Frank. 2018. Knowledgeable reader: Enhancing cloze-style reading comprehension with external commonsense knowledge. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 821–832, Melbourne, Australia.
- Barlas Oguz, Xilun Chen, Vladimir Karpukhin, Stan Peshterliev, Dmytro Okhonko, Michael Schlichtkrull, Sonal Gupta, Yashar Mehdad, and Scott Yih. 2022. UniK-QA: Unified representations of structured and unstructured knowledge for open-domain question answering. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1535–1546, Seattle, United States.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: an open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, page 4444–4451.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Jiashuo Sun, Chengjin Xu, Luminyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-graph: Deep and responsible reasoning of large language model with knowledge graph. *arXiv preprint arXiv:2307.07697*.

- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A question answering challenge targeting commonsense knowledge. 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)*, pages 4149–4158, Minneapolis, Minnesota.
- Yijun Tian, Huan Song, Zichen Wang, Haozhu Wang, Ziqing Hu, Fang Wang, Nitesh V. Chawla, and Panpan Xu. 2024. Graph neural prompting with large language models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):19080–19088.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Kuan Wang, Yuyu Zhang, Diyi Yang, Le Song, and Tao Qin. 2021. Gnn is a counter? revisiting gnn for question answering. *arXiv preprint arXiv:2110.03192*.
- Yunzhi Yao, Shaohan Huang, Li Dong, Furu Wei, Huajun Chen, and Ningyu Zhang. 2022. Kformer: Knowledge injection in transformer feed-forward layers. In *Natural Language Processing and Chinese Computing: 11th CCF International Conference, NLPCC 2022, Guilin, China, September 24–25, 2022, Proceedings, Part I*, page 131–143, Berlin, Heidelberg.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. Editing large language models: Problems, methods, and opportunities. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10222–10240, Singapore.
- Zhiyu Yao, Xinyang Chen, Sinan Wang, Qinyan Dai, Yumeng Li, Tanchao Zhu, and Mingsheng Long. 2024. Recommender transformers with behavior pathways. In *Proceedings of the ACM Web Conference 2024*, page 3643–3654, New York, NY, USA.
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2022. Retrieval-augmented multimodal language modeling. *arXiv preprint arXiv:2211.12561*.
- Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy Liang, and Jure Leskovec. 2024. Deep bidirectional language-knowledge graph pretraining. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA.
- Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. QA-GNN: Reasoning with language models and knowledge graphs for question answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 535–546, Online.
- Guoqi Yu, Jing Zou, Xiaowei Hu, Angelica I Aviles-Rivero, Jing Qin, and Shujun Wang. 2024. Revitalizing multivariate time series forecasting: Learnable decomposition with inter-series dependencies and intra-series variations modeling. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pages 57818–57841.
- Hongbin Zhang, Kehai Chen, Xuefeng Bai, Yang Xiang, and Min Zhang. 2024. Paying more attention to source context: Mitigating unfaithful translations from large language model. *arXiv preprint arXiv:2406.07036*.
- Miao Zhang, Rufeng Dai, Ming Dong, and Tingting He. 2022a. DRLK: Dynamic hierarchical reasoning with language model and knowledge graph for question answering. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5123–5133, Abu Dhabi, United Arab Emirates.
- Qinggong Zhang, Junnan Dong, Hao Chen, Xiao Huang, Daochen Zha, and Zailiang Yu. 2023. Knowgpt: Black-box knowledge injection for large language models. *arXiv preprint arXiv:2312.06185*.
- Xikun Zhang, Antoine Bosselut, Michihiro Yasunaga, Hongyu Ren, Percy Liang, Christopher D Manning, and Jure Leskovec. 2022b. Greaselm: Graph reasoning enhanced language models for question answering. *arXiv preprint arXiv:2201.08860*.
- Yu Zhang and Zhao Kang. 2024. Tpn: Transferable proto-learning network towards few-shot document-level relation extraction. In *2024 International Joint Conference on Neural Networks*, pages 1–9.
- Xudong Zhu, Zhao Kang, and Bei Hui. 2024a. FCDS: Fusing constituency and dependency syntax into document-level relation extraction. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, pages 7141–7152, Torino, Italia.
- Yuqi Zhu, Shuofei Qiao, Yixin Ou, Shumin Deng, Ningyu Zhang, Shiwei Lyu, Yue Shen, Lei Liang, Jinjie Gu, and Huajun Chen. 2024b. Knowagent: Knowledge-augmented planning for llm-based agents. *arXiv preprint arXiv:2403.03101*.



Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. A robustly optimized BERT pre-training approach with post-training. In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 1218–1227, Huhhot, China.

## A Experiments

### A.1 Knowledge Graph and Datasets

**Knowledge Graph** We conduct experiments in the general domain (commonsense reasoning) and we consider ConceptNet (Speer et al., 2017) that contains rich commonsense knowledge regarding the daily concepts. Node embeddings are initialized using the entity embeddings prepared by Lin et al. 2019, which applies pre-trained LMs to all triples in ConceptNet and then obtains a pooled representation for each entity.

**OpenBookQA** is a kind of question-answering dataset modeled after open book exams for assessing human understanding of a subject. It consists of 5,957 multiple-choice elementary-level science questions, which probe the understanding of a small “book” of 1,326 core science facts and the application of these facts to novel situations. Answering OpenBookQA questions requires additional broad common knowledge, not contained in the book. We use the original data splits in Mihaylov and Frank 2018.

**ARC** consists of 7,787 science exam questions drawn from a variety of sources, including science questions provided under license by a research partner affiliated with AI2. These are text-only, English language exam questions that span several grade levels as indicated in the files. Each question has a multiple choice structure (typically 4 answer options). We evaluate challenge sets in it using the original data splits in Clark et al. 2018.

**Riddle** is a multiple-choice question answering task for answering riddle-style commonsense questions. It involves a challenging cognitive process, in which it requires complex commonsense reasoning abilities, an understanding of figurative language, and counterfactual reasoning skills, which are all important abilities for advanced natural language understanding (NLU). We split the dev set in half to make in-house dev/test sets in (Yasunaga et al., 2024).

**PIQA** is a dataset for commonsense reasoning, and was created to investigate the physical knowledge of existing models for naive physics

reasoning focusing on how we interact with everyday objects in everyday situations.

### A.2 Implementation & training details

For our method, we set the dimension of GNN module is 200 and number of layers of our GNN module from 5 to 7, with dropout rate (Srivastava et al., 2014) 0.2 applied to each layer. We train the model with the RAdam (Liu et al., 2019) optimizer using A100-80G GPUs. We set the learning rate for LLM module from 1e-4, 3e-4, learning rate for other modules as 1e-3, batch\_size from 16, 32. The above hyperparameters are tuned on the development set. We choose FLAN-T5-XL (3B) and XXL (11B) (Chung et al., 2024) as the LLMs used in this paper. We adjust the maximum sequence length of LLMs to best fit the question length for each dataset.

### A.3 Baseline

In this setting of RoBERTa-Large, we compare with six baselines including **Full fine-tuning**, **QAGNN** (Yasunaga et al., 2021) that proposes to use node relevance scores to filter irrelevant nodes and utilize joint training with GNN and pre-trained model (PLM), **GreaseLM** that designs a modality interaction unit to achieve information interaction between PLM and GNN, **drk** (Zhang et al., 2022a) that proposes a dynamic hierarchical reasoning method with PLM and knowledge graphs, **GSC** (Wang et al., 2021) simplifies the network structure and carries out some simple reasoning such as counting, **Rumination** that prompts PLM to generate inner knowledge and then inject it into PLM again. In the setting of LLaMA2-7B, we compare with **LLM-only**, **Full fine-tuning** and **LoRA**.

## B Case Study

In this section, we provide more results of case analysis in Figure 8. In the four method successfully predicts the correct answer *rooster* by accurately predicting the relevance scores between the question and all triplets in the knowledge graph and eliminating noisy pathways. In the second case, three methods failed to make the correct prediction. This is mainly because the question involves symbolic reasoning, and external commonsense knowledge injection could not offer substantial support which is a limitation of our method.

**Question:**


For **vigilance** and **courage** true, I've no **superior**, equals few: Which makes me by the industrious priz'd, But by the **indolent** despis'd; **Bold** and **alert** I meet the **foe**, In all **engagements** without **valour** show; And if he proves too **proud** to **yield**, One **falls** before we **quite** the **field**.

**Options:**

(a) fearless (b) vain (c) egg (d) rooster (e) challenge

**Model Prediction:**

w/o KGR: (a) fearless ❌ Ours: (d) rooster ✔️



**Question:**

My first is **ocean** but not in **sea**, my **second** in **milk** but not in me my **third** is in **three** but not in **throw**, my **fourth** in **vow** but not in **crow** my **fifth** is in **eight** but not in **night**, my **last** is in **wrong** and also **right** my **whole** is **praise** for **thoughts** or **men**; or **women**, too, or **tongue** or **pen**

**Options:**

(A) center (B) clever (C) incorrect (D) belle (E) detective

**Model Prediction:**

LLM-only: (A) incorrect ❌ w/o QG-KGR: (A) incorrect ❌  
Ours: (D) incorrect ❌




Figure 8: More case analysis results.