Enhancing Conversational Question Answering through Reinforced Question Reformulation and Prompt Refinement in Children Application

1st Fen Zhao
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
2022106@njxzc.edu.cn

2nd Xinheng Wang
Nanjing, China
xhwangcqupt@gmail.com

3rd Haifei Zhang
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
2024005@njxzc.edu.cn

4th Jie Zhu
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
zhuj@njxzc.edu.cn

5th Lingling Zhang
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
zhanglingling@njxzc.edu.cn

6th Kexin Liu
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
24130530@njxzc.edu.cn

7th Meiqi Shi
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
shi_meiqi26@outlook.com

8th Siyi He
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
15261498505@163.com

9th Yi Wang
School of Information Engineering
Nanjing Xiaozhuang University
Nanjing, China
wangyi26@njxzxy133.wecom.work

***Abstract*—While question answering (QA) technologies have progressed significantly, their deployment for primary-school children remains understudied. Childrens queries often exhibit semantic ambiguity, incomplete expressions, and unclear articulation, leading to unanswered questions or irrelevant responses. Leveraging external knowledge is key to addressing this common challenge of ambiguity. Consequently, effectively reformulating questions and refining prompts into precise SPARQL queries are crucial tasks. This paper presents a multi-task learning approach utilizing a text generation model for both question reformulation and prompt refinement. Additionally, to bridge the preference gap between the question reformulation module and the QA model, we introduce a training strategy that utilizes QA model feedback to further optimize the question reformulation processes. Evaluations across two datasets show our method outperforms baselines, achieving statistically significant gains of +5.96% to +12.70% in F1 and Acc scores. These consistent improvements confirm that unified question reformulation and prompt refinement, reinforced by downstream QA feedback, substantially advance child-oriented conversational QA.**

***Index Terms*—SPARQL queries, multi-task learning, question reformulation, prompt refinement, child-oriented conversational QA.**

This work is supported by the General Project of Basic Science Research for Universities in Jiangsu Province (Grant no. 23KJB520017), the Scientific Research Project of Nanjing Xiaozhuang University (Grant no. 2022NXY35), the National Natural Science Foundation of China (Grant nos. 12171065), Nanjing Xiaozhuang University 2024 Jiangsu Province College Students Innovation Training Program Project (Grant no. 202411460019Z), Key Project of the 14th Five-Year Plan for Education Science in Jiangsu Province (B-b/2024/01/169).

I. Introduction

With the ongoing advancement of large-scale language models (LLMs), intelligent question answering (QA) technology has become increasingly mature [1], [2]. However, the development of children's intelligent QA technology faces unique challenges when it comes to the scenario of primary education. Due to the particularity of children's language expression, children's language expression shows problems such as semantic ambiguity, incomplete expression, and unclear pronunciation [3]. Such phenomena impact question understanding and answer retrieval, leading to irrelevant answers or unanswered questions.

In the field of QA, particularly those tailored for children's applications, enhancing conversational capabilities is crucial. Recently, text generation has been used to enhance conversational QA performance through question reformulation (QR) technology [6]–[8]. QR technology is to convert a ambiguous question into a question with clear intent and complete phrasing that can be understood without additional information. For example, QR technology is used in conversational QA models to refine and clarify user queries, enabling the system to better comprehend the user's intent. The QR task not only reduces

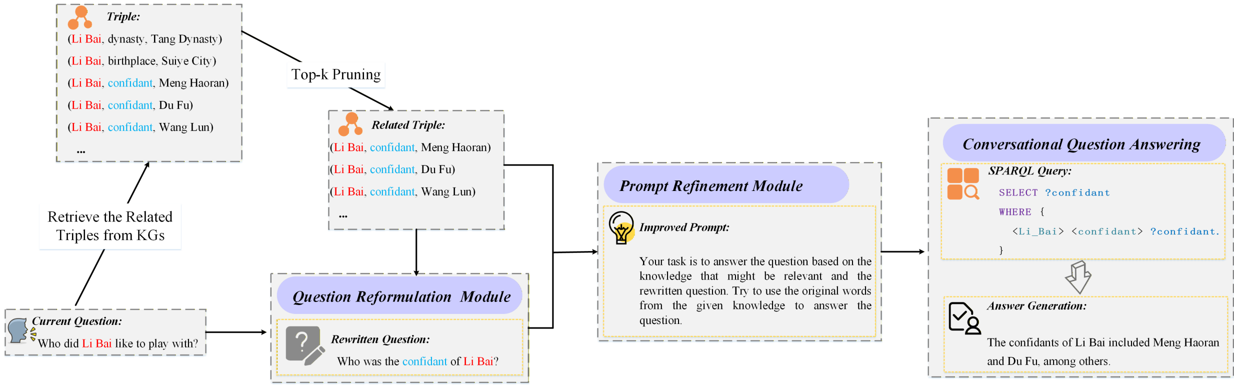

Fig. 1. An example of the conversational QA through QR and PR technologies.

the need for frequent retraining of the underlying models but also improves the system's ability to handle ambiguous or incomplete questions. This makes QR a more practical and efficient solution, which has been widely explored to improve answer accuracy and relevance [9].

Based on these observations, we propose QRPR-QAC, a novel **Q**uestion **R**eformulation and **P**rompt **R**efinement guided **QA** approach in **C**hildren application, which consists of three steps: 1) reformulating the current question into a question with clear intent and complete phrasing, and further converting it into a reformulating edit matrix; 2) refining the initial prompt into the knowledge-aware prompt by fusing the related facts retrieved from knowledge graphs (KGs) with the reformulating edit matrix; 3) injecting the knowledge-aware prompt and reformulating edit matrix into the conversational QA model. Firstly, we train and test the QR model. Secondly, inspired by [10], we propose to integrate rewritten questions into the conversational QA task in the form of a rewriting relation matrix. We found that directly concatenating the rewritten questions might mislead the model into generating incorrect answers. The reason is that noise inevitably exists in the process of rewriting. In addition, some questions are semantically complete and do not need to be rewritten. The rewritten editing matrix represents the relationship between the initial question and the candidate relations. The relationship can clearly guide the model in solving long-term dependency problems. Furthermore, based on related facts of KGs, we refine the initial prompt to become the knowledge-aware prompt, thereby explicitly aligning the rewritten question with external knowledge. Finally, we feed the jointly optimized rewritten edit matrix and knowledge-aware prompt into the conversational QA model. Fig. 1 illustrates the major process in the conversational QA through QR and prompt refinement (PR) technologies.

In brief, our main contributions are as follows:

- To enhance the clarity and completeness of childrens questions, we propose a novel QR module. The module distills the Top-3 relevant triples from background facts, and then rewrites ambiguous or incomplete queries into precise, answer-ready forms. Through the process, it grounds vague questions in structured KGs knowledge, effectively mitigating ambiguity and information gaps.
- To generate well-defined SPARQL queries, we introduce a knowledge-aware prompt generation technique. By integrating related KGs facts into prompts, it enriches the background knowledge available to the language model and guides the model to generate more accurate and contextually appropriate responses.
- Experimental results show that QRPR-QAC achieves comparable performance to recent state-of-the-art works on two datasets. Specifically, on WebQuestionsSP-C, our method achieves an F1 score of 62.5%, surpassing the second-best model (CoTKR) by 6.7%, confirming the practical value of jointly reformulating questions and refining prompts for child-oriented QA.

II. Related Work

Conversational QA: Conversational QA methods proposed belong to two major families; they either (i) model the history as additional context to answer the current question, or (ii) integrate external knowledge sources to enhance the factual accuracy and depth of the responses.

The first family has been the development of context-aware models, which can understand and leverage the conversational context to generate more coherent and accurate responses. Sarawoot et al. [8] proposed an approach that leverages conversational context to enhance conversational QA system. By utilizing dialog history for follow-up question identification and question rewriting, this work highlights the importance of leveraging conversational context to enhance the accuracy of responses. Similarly, Chai et al. [10] presented a context-dependent text-to-SQL semantic parsing approach. Another related study is conducted by Raposo et al. [11], which emphasized the significance of the question rewriting module in enhancing the overall system performance. In addition, Chen et al. [5] proposed a novel approach to conversational QA by integrating reinforcement learning for question rewriting. This method utilized QA feedback to guide the rewriting model, aiming to enhance the interpretability of conversational questions within dialogues. These studies focus on leveraging the conversational context to generate more coherent and accurate responses. However, the aforementioned methods mainly focus on leveraging the internal dialog context and may face challenges when dealing with complex or ambiguous questions that require external knowledge for accurate answering.

Further research has investigated the application of external knowledge to improve the factual precision of responses. In this regard, Kaiser et al. [9] proposed a framework that employed reinforced reformulation generation to improve conversational QA models over KGs. Similarly, Wu et al. [6] proposed the CoTKR framework, which employed chain-of-thought enhanced knowledge rewriting to generate reasoning traces and corresponding knowledge. Furthermore, Mao et al. [7] introduced a framework that leveraged database content and execution feedback to enhance Text-to-SQL parsing. Additionally, Liu et al. [12] proposed a novel model, which combined QR with reinforcement learning over KGs to achieve accurate conversational QA. Overall, these approaches demonstrate the benefits of incorporating external knowledge to enhance the performance of QA systems.

Our approach extends beyond the limitations of existing methods by not only leveraging conversational context but also deeply integrating external knowledge from KGs. Through reformulating ambiguous or incomplete questions with relevant facts and conversational content, our work aims to enhance both the clarity of questions and the accuracy of responses, providing more appropriate and contextually relevant answers for children's queries.

Conversational QR: In the realm of QR, prior studies have laid important groundwork. Approaches have evolved from early sequence-to-sequence methods, such as LSTM and GRU [13], to self-supervised learning methods and the application of transformer architectures, such as GPT-2 [14]. For instance, Elgohary et al. [15] proposed a sequence-to-sequence model to rewrite the current question by leveraging conversational history, thereby enhancing the coherence and relevance of questions in the dialogue context. Similarly, Vakulenko et al. [16] introduced a model utilizing a unidirectional transformer decoder for automatic question rewriting, which represented a significant step forward in automating the QR process. However, these methods lack the depth of integration with external knowledge, which failed to address the unique challenges posed by children's queries. Fader et al. [17] focused on mining query rewriting rules from background KGs and employing a query rewriting operator to generate new questions. More recently, Li et al. [18] explored enhancing dialogue generation with conversational concept flows by using conversation-aware KGs and a novel relation-aware graph encoder. In contrast to these studies, our work focuses on addressing these limitations by deeply integrating conversational context with external knowledge from KGs.

III. Task Formulation and Challenges

In this section, we first formalize the knowledge-aware conversational QA task. Then, we introduce the key challenges in this research direction.

A. Task Formulation

We first define conversational QR task. Formally, conversational QR task is defined as follows: given a child's ambiguous or incomplete original query q_{ori} within a conversational context, and a set of relevant facts $Facts = \{(h_1, r_1, t_1), (h_2, r_2, t_2), \ldots, (h_n, r_n, t_n)\}$ extracted from external KGs G, the task is to automatically reformulate the original query q_{ori} into a clearer, more precise, and contextually relevant form q_{rew}. The rewriting question q_{rew} incorporates relevant facts $Facts = \{(h_c, r_c, t_c)\}_{c=1}^{n}$. The reformulated question should not only be grammatically correct but also semantically richer, drawing on external knowledge introduced during rewriting.

B. Challenges

In order to accurately understand children's intent and achieve more accurate interaction between children and educational systems, some problems should be considered. That is, due to the ambiguity of children's expressions, how to use disambiguation techniques to address the ambiguity in children's expressions, and how to maintain the coherence of the dialogue and handle context dependence.

We conclude challenges as the following aspects:

- Based on existing conversational QA methods, it is difficult to disambiguate ambiguous questions due to the lack of integration of external knowledge and dialogue context.
- Due to the difficulty in generating accurate SPARQL queries from the initial prompts, there are issues such as predicate-entity misalignment and empty results. Thus, it is hard to ensure the completeness and correctness of the answers.

IV. Model

In this section, we present a method that leverages external knowledge to enhance QR. Moreover, our framework presents a knowledge-aware PR module to generate well-defined SPARQL queries. Additionally, to bridge the preference gap between modules QR and PR and the QA model, we propose a training strategy, for leveraging feedback from the QA model to further optimize the question reformulation and prompt refinement. The specific details of enhancing conversational QA through reinforced QR and PR can be given in Algorithm 1, which mainly consists of four parts: initialization, QR and PR, answer generation, and reinforced QR and PR. Input is the original question q_{ori}. Output is the final answer A. An overview of the workflow in the proposed QRPR-QAC architecture is depicted in Fig. 2. The pipeline consists of three modules, where the first two (QR and PR) are our contributions.

A. Question Reformulation

The QR module is designed to address the inherent ambiguities and incompleteness in children's queries by dynamically integrating dialogue context and external knowledge. Specifically, the module first identifies the key entity such as "moon"

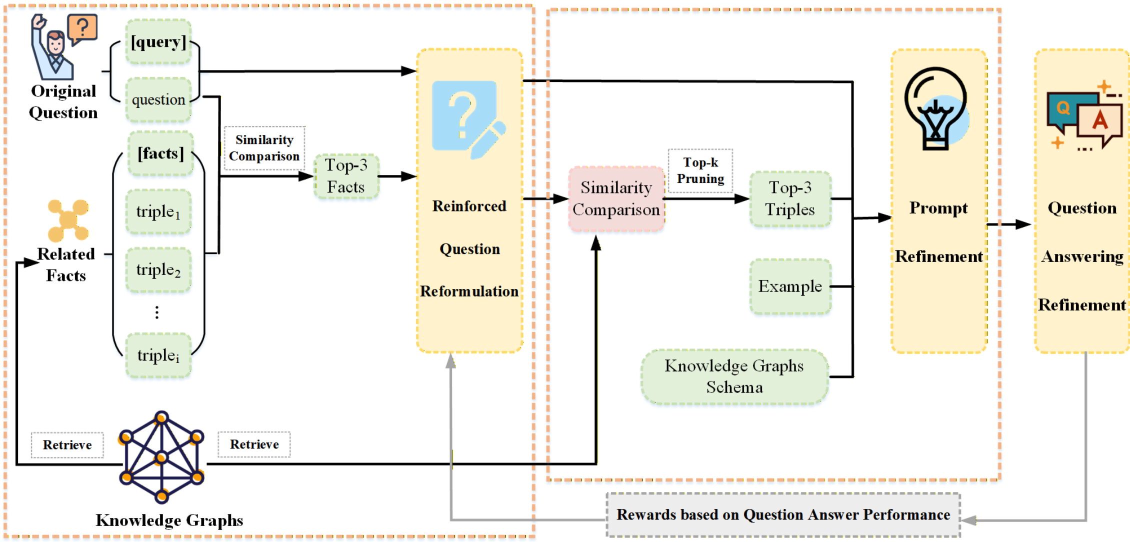

Fig. 2. Illustration QRPR-QAC framework.

Algorithm 1 Enhancing Conversational Question Answering through Reinforced Question Reformulation and Prompt Refinement

---

Input: Original quetion q_{ori}.

Output: Final answer A.

1. Retrieve related triples $\{T_1, T_2, \ldots, T_n\}$ based on the original question q_{ori};
2. Initialize the reinforced QR module and PR module;
3. **for** $i = 1$ to max_reformulation_attempts **do**
4. Use QR to rewrite the original question q_{ori} into a new question q_{rew};
5. Generate an initial prompt p_i by combining q_{rew} and $\{T_1, T_2, \ldots, T_n\}$;
6. Perform similarity-based pruning and Top-k selection to extract a subset of the relevant triples $\{T_{i1}, T_{i2}, \ldots, T_{in}\}$;
7. Optimize p_i using the PR module to generate a knowledge-aware prompt $p_{knowledge}$;
8. Obtain an answer A by invoking the QA model with the optimized prompt $p_{knowledge}$;
9. **if** A meets the termination criteria (e.g., answer quality exceeds a threshold or maximum iterations are reached) **then**
10. **break**
11. **end if**
12. **end for**
13. **return** the final answer A;
14. Calculate a reward R_{ew} based on the quality of the answer;
15. Update the parameter of the QR module to reinforce the generation of clear intention questions.

---

and extracts the relation like "look fat and look thin". Then, leveraging the KGs relation alignment technology, it encodes the relation into vector and matches it with candidate relations in the KGs. By combining entity type constraints and rule-based filtering, it narrows down to the most relevant candidate relations such as "phases of waxing and waning", "moon phase change", etc. Finally, a question generation model, powered by the T5 model [19], outputs a rewritten question, such as "Why does the moon have phases of waxing and waning?". Instead of feeding the rewritten question directly into the conversational QA model, we further convert the rewritten question into a rewriting matrix.

Consider a topic t from a set of topics M, the QR model accepts the original question q_{ori} and a set of relevant facts $Facts_c^t = \{(h_1, r_1, t_1), \ldots, (h_c, r_c, t_c), \ldots, (h_n, r_n, t_n)\}$ as inputs. We employ a pre-trained T5-base sequence generator [19] as our QR model. The QR model processes these inputs to generate a reformulated question q_{rew}. Given a original question q_{ori} with its relevant facts $Facts_c^t = \{(h_c, r_c, t_c)\}_{c=1}^n$ and a human rewritten ground truth $\tilde{q}_{hum}$, our objective is to induce a function $f(\{q_{ori}, Facts_t\}) = \tilde{q}_{hum}$. Finally, given probability P conditioned on a parameterized function $\bar{f}$, the original question q_{ori}, and the relevant facts $Facts_c^t = \{(h_c, r_c, t_c)\}_{c=1}^n$, the overall objective of the task is then defined in terms of finding the parameters φ by maximum likelihood estimation:

$$\varphi = \arg\max_{\varphi} \prod_{t=1}^{M} \prod_{c=1}^{n} P_{tc}\left(\tilde{q}_{hum} \middle| \bar{f}\left(\{q_{ori}, Facts_c^t\}, \varphi\right)\right) \tag{1}$$

As described above, we initialize the QR model using a pre-trained T5-base sequence generator. The model is subsequently fine-tuned to produce more precise rewritten questions. And then, following the inspiration from [10], we propose to incorporate the reformulated results into the conversational QA task,

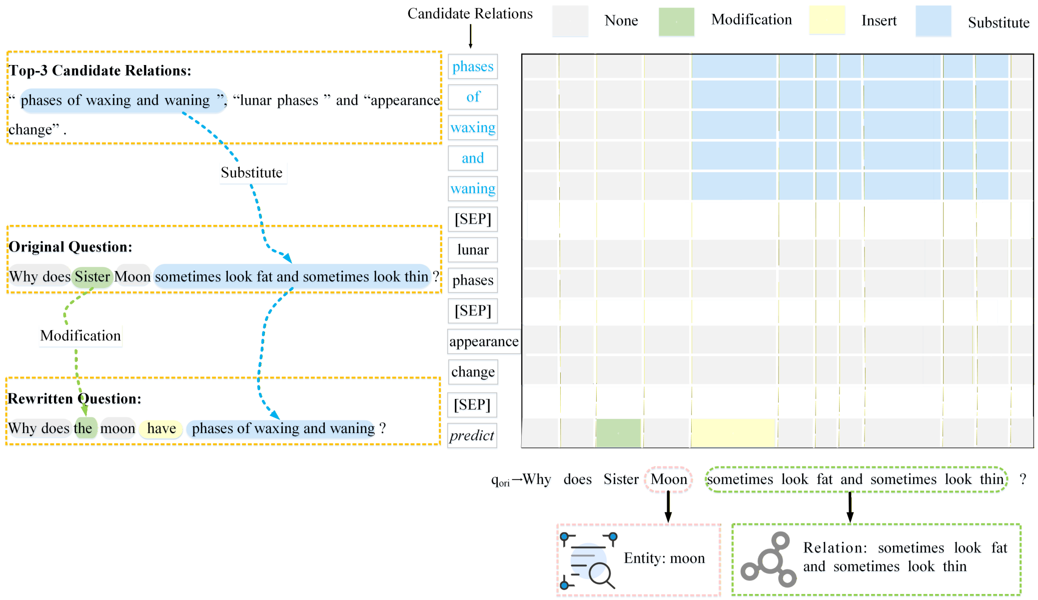

Fig. 3. An example of rewriting edit matrix.

structured as a rewriting relation matrix. The matrix reflects the connections between the original question and candidate relations of KGs. As shown in Fig. 3, the reformulated question is derived from the original one through a sequence of edit operations recorded in the rewriting edit matrix. Specifically, we substitute the phrase "sometimes look fat and sometimes look thin" with "phases of waxing and waning" extracted from KGs, and insert "have" before "phases of waxing and waning". Additionally, the modification operation increases formality via lexical modification with high-entropy words (replace "Sister" in the current question with "the").

B. Prompt Refinement

After the original question is reformulated, the initial prompt can be generated using an existing prompt framework. Subsequently, our framework refines the initial prompt into a knowledge-aware one by incorporating relevant facts from KGs.

At most k triples are included in the prompt. We set $k = 3$ for our method. Three related triples are selected based on their relevance to the rewritten question q_{rew}. We calculate the relevance of a triple to a rewritten question q_{rew}. Firstly, we convert the triple into the JSON dictionary $\tilde{D}$, using the triple's identifier as the key and its entity and relation as the value. Then, we tokenize the dictionary $\tilde{D}$, resulting in $\tilde{D} = \left\{\tilde{d}_1, \tilde{d}_2, \ldots, \tilde{d}_c\right\}$. Next, we compute the embedding of $\tilde{D}$:

$$\tilde{\mathbf{D}} = \frac{1}{|\tilde{D}|}\sum_{i=1}^{c} \text{TransE}\left(\tilde{d}_i\right) \tag{2}$$

where TransE(x) refers to the TransE embedding [20] of token x. Similarly, we tokenize the rewritten question q_{rew}, resulting in $q_{rew} = \{q_{rew-1}, q_{rew-2}, \ldots, q_{rew-m}\}$. The GloVe embedding of the rewritten question q_{rew} is represented as follows:

$$\mathbf{q}_{rew} = \frac{1}{|q_{rew}|}\sum_{i=1}^{m} \text{GloVe}\left(q_{rew-i}\right) \tag{3}$$

where GloVe(x) refers to the GloVe embedding [21] of token x. Based on the dot-product calculation, the semantic similarity between the rewritten question and the triple is computed. The dot product of $\tilde{\mathbf{D}} = \left(\tilde{D}_1, \tilde{D}_2, \ldots, \tilde{D}_n\right)$ and $\mathbf{q}_{rew} = (Q_{rew-1}, Q_{rew-2}, \ldots, Q_{rew-n})$ is calculated. Finally, we obtain three most relevant facts, namely Top-3 facts. The embeddings of Top-3 facts is represented as follows:

$$Emb\_facts = \frac{1}{3}\sum_{\kappa=1}^{3} \tilde{\mathbf{D}}_{top-\kappa} \tag{4}$$

where $\tilde{\mathbf{D}}_{top-1}$, $\tilde{\mathbf{D}}_{top-2}$, and $\tilde{\mathbf{D}}_{top-3}$ represent the embedding vectors of the Top-3 facts, respectively.

To convert the rewriting edit matrix $\mathbf{M}$ into embeddings, we first initialize an empty list called Emb_matrix. Then, we tokenize the original question q_{ori}, resulting in $q_{ori} = \left\{q_{ori}^1, q_{ori}^2, \cdots, q_{ori}^n\right\}$. The GloVe embedding of the original question $\mathbf{q}_{\text{ori}}$ is computed. Next, we append a dictionary to Emb_matrix, which contains $\mathbf{q}_{\text{ori}}$, $\mathbf{q}_{\text{rew}}$, and the operation (such as $Substitute$, $Insert$, $Modify$, or $None$). Subsequently, we concatenate the fact embedding with the matrix embedding:

$$\mathbf{V} = \text{Concatenate\_emb}\left(Emb\_facts, Emb\_matrix\right) \tag{5}$$

Finally, we generate knowledge-aware prompt by combining $\mathbf{q}_{rew}$ and $\mathbf{V}$:

$$\mathbf{p}_{knowledge} = \text{generate\_prompt}\left(\mathbf{q}_{rew}, \mathbf{V}\right) \tag{6}$$

TABLE I
EXAMPLES OF FINE-TUNING DATA.

Ambiguous Question	Type	Natural Answer	Rewritten Question
Why aren't there any dinosaurs anymore?	Colloquialism	A massive asteroid hit Earth 66 million years ago, wiping them out.	Why did the dinosaurs die out?
Sun Wukong is so powerful.	Incomplete expression	He can fly, lift 8 tons, clone himself, shapeshift, and see through lies—plus his staff grows sky-high.	What powerful skills does Sun Wukong possess?
I can not fly.	Unclear expression	Their hollow bones and wings make lift, saving time and keeping them safe.	Why do birds fly?

TABLE II
THE OVERALL RESULTS (%) OF QRPR-QAC AND THE BASELINES. THE BEST AND SECOND-BEST RESULTS ARE MARKED IN RED AND BOLD, RESPECTIVELY.

Methods	WebQuestionsSP-C				GraphQuestions-C			
	Acc	Recall	Precision	F1	Acc	Recall	Precision	F1
CoTKR	**67.2**	**51.9**	**67.3**	**58.6**	**68.4**	52.3	**68.1**	**59.1**
REIGN	61.9	51.2	58.2	54.5	62.3	**52.9**	59.8	56.1
SURF	63.0	49.3	60.7	54.8	64.2	50.4	61.5	55.4
Explaignn	59.7	50.6	59.6	54.7	60.9	51.2	60.1	55.3
QAF-SuRM	55.1	50.8	53.3	52.0	56.6	50.3	54.8	52.4
QRPR-QAC (Our work)	76.3	57.5	68.4	62.5	77.1	58.8	68.8	63.4
QRPR-QAC without QR	68.1	52.5	67.0	58.9	69.6	54.2	67.8	60.2
QRPR-QAC without PR	72.4	53.7	67.8	59.9	73.4	55.1	68.2	61.0

C. Conversational Question Answering

After receiving the refinement prompt, the conversational QA model then outputs a single SPARQL query to answer the question. We then start a self-reinforcement module, as illustrated in Fig. 2. The module operates within a feedback-driven framework: after the conversational QA module produces responses, QA performance metrics are computed. These metrics are combined into a single reward signal using the weighted-sum method, defined as:

$$\text{Combined\_Metric} = \alpha \times Acc + (1-\alpha) \times F1 \tag{7}$$

where α is a tunable parameter between 0 and 1 that adjusts the relative importance of Acc and F1. The combined metric serves as reward feedback to a reinforcement learning mechanism. These rewards iteratively refine parameters of the QR module and the PR module, ultimately optimizing the downstream QA model.

In our implementation, we utilize advanced LLMs like ChatGLM2-6B [22] to power the conversational QA system. ChatGLM2-6B is an advanced open-source language model developed by Tsinghua University and Moonshot AI.

V. EXPERIMENTAL SETUP

A. Datasets

Pre-training Dataset. We construct two unique datasets, WebQuestionsSP-C and GraphQuestions-C, using the existing WebQuestionsSP [23] and GraphQuestions [24] corpora. Leveraging the generative capacity of LLMs, we synthesize natural language questions that exhibit the linguistic features characteristic of children's speech. These characteristics include colloquialism, unclear expression, and incompleteness, among others. We utilized 20% of the dataset for validation.

Fine-Tuning Dataset. To facilitate the development of a child-centric QA system, we curate a specialized corpus for the supervised fine-tuning of LLMs. Manual collection by our team members and graduate students yielded approximately 3,500 instances. Each instance includes an ambiguous question, a rewritten question, a standard answer, and the question type. Table I presents illustrative examples of the collected data. The corpus captures colloquial, incomplete and ambiguous child-like utterances, providing a targeted benchmark for improving our child-centric QA system. We utilized 15% of the dataset for validation.

B. Baselines and Implementation

Baselines. In order to assess the effectiveness and efficiency of our work, we consider the following baselines: CoTKR [6], REIGN [9], SURF [4], Explaignn [25], and QAF-SuRM [5].

CoTKR is a chain-of-thought enhanced knowledge rewriting method designed for KGs-QA. REIGN is a robust training framework for conversational QA. It employs reinforced reformulation generation to enhance model generalization. SURF is a semantic question reformulation model designed to handle challenging user utterances in spoken QA. Explaignn is a conversational QA system that integrates heterogeneous sources via a graph neural network. QAF-SuRM proposes a reinforcement learning framework for conversational question rewriting. To ensure a fair comparison, we employ the same corpus generation method for both the baselines and our method.

VI. EXPERIMENTAL RESULTS

In this section, we evaluate QRPR-QAC on conversational QA tasks against baselines and subsequently discuss the results under each evaluation strategy.

TABLE III
ABLATION STUDY RESULTS (%) BY REMOVING THE MAIN COMPONENTS.

Methods	WebQuestionsSP-C				GraphQuestions-C			
	Acc	Recall	Precision	F1	Acc	Recall	Precision	F1
QRPR-QAC (Our work)	76.3	57.5	68.4	62.5	77.1	58.8	68.8	63.4
QRPR-QAC without QR	68.1	52.5	67.0	58.9	69.6	54.2	67.8	60.2
QRPR-QAC without PR	72.4	53.7	67.8	59.9	73.4	55.1	68.2	61.0
QRPR-QAC without reinforcement learning	74.8	56.3	68.1	61.6	76.2	56.9	68.5	62.2

A. Overall Performance

Table II presents the overall results. It is evident that our work outperforms the baselines on most metrics, where F1 score increases by 3.9 percentage points compared with the strongest baseline CoTKR on the WebQuestionsSP-C dataset, as shown in Table II. Similarly, although QAF-SuRM shows competitive performance in some metrics, the overall performance of QAF-SuRM is not good enough. Explaignn performs reasonably well but still lags behind our work. SURF also underperforms compared to our work.

Different from the baselines, our primary contribution is to combine QR and PR within a unified framework for conversational QA. By jointly leveraging the strengths of both techniques, the proposed method significantly improves answer accuracy and robustness. This is reflected in the experimental results, where QRPR-QAC consistently outperforms other methods across multiple metrics. For example, on WebQuestionsSP-C, QRPR-QAC attains a Precision of 68.4, outperforming CoTKR by 1.1 and QAF-SuRM by 15.1 percentage points, as shown in Table II. The unified QR+PR framework demonstrably resolves conversational entity ambiguity more effectively than alternative pipelines.

B. Ablation Study Results

Results of the ablation study are shown in Table III. In Table III, the last three rows correspond to (i) QRPR-QAC without QR, (ii) without PR, and (iii) without reinforcement learning, respectively. “QRPR-QAC without QR” indicates that the QR module is removed. “QRPR-QAC without PR” indicates that the PR module is removed. And, “QRPR-QAC without reinforcement learning” indicates that the reinforcement learning module is removed.

The experiment result shows that each component can effectively enhance the performance of KGs-QA. Removing any single module causes a consistent drop on all metrics. Among them, suppressing QR yields the largest degradation: F1 drops by 3.6 on WebQuestionsSP-C and by 3.2 on GraphQuestions-C. This confirms that a well-designed QR component is the most critical for boosting KGs-QA performance.

C. Impact of Semantic Fuzziness

In our experiment, the validation set is stratified into five distinct classes. Fig. 4 reports the five-class BLEU distribution on the validation set. QA pairs are stratified according to the degree of semantic ambiguity, ranging from class 0 (clear semantics) to class 4 (highly ambiguous). There is a monotonic deterioration in BLEU: class 0 attains 84.6%, whereas class 4 drops to 29.2%an absolute decrease of 55.4 percentage points. The near-linear drop in BLEU demonstrates that the rewriter’s ability to reproduce reference paraphrases deteriorates in direct proportion to the degree of semantic ambiguity. This confirms that output quality is primarily constrained by the syntacticsemantic clarity of the input question.

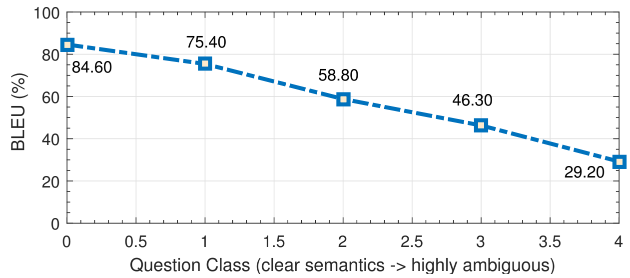

Fig. 4. 5-class BLEU scores.

D. Case Study

We explore how each QRPR-QAC module drives LLMs to produce higher-quality SPARQL queries. The representative case is listed in Table IV.

In this case, the initially ambiguous question “Who did Li Bai like to play with?” is clarified through the QR module, which reformulates it as “Who was the confidant of Li Bai?”. This process effectively resolves the ambiguity by focusing on key facts. Subsequently, the PR module generates a clear instruction, guiding LLMs to produce a concise and accurate SPARQL query. The execution of this query retrieves relevant entities, such as Meng Haoran and Du Fu, thereby illustrating QRPR-QAC’s capability to diminish semantic ambiguity and enhance query precision through the utilization of the KGs subgraph. The final execution result is correct and identical to the ground-truth output. This case study collectively highlights how the synergistic integration of QR, PR, and conversational QA modules significantly improves the model’s capacity to produce accurate and interpretable outputs.

VII. CONCLUSION

We introduce an approach for conversational QA that simultaneously trains to reformulate conversational queries and enhance prompt generation. The QR module resolves ambiguous inputs, and the PR module constructs fact-based prompts by merging rewritten questions with relevant KGs

TABLE IV
CASE STUDY OF QRPR-QAC.

Ambiguous Question	Who did Li Bai like to play with?
Subgraph Related to Question	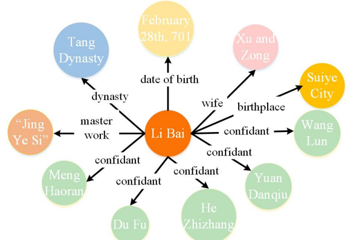
Rewritten Question	Who was the confidant of Li Bai?
Refined Prompt	Your task is to answer the question based on the knowledge that might be relevant and the rewritten question. Try to use the original words from the given knowledge to answer the question.

SPARQL Query	SELECT ?confidant
	WHERE {
	<Li_Bai> <confidant> ?confidant.
	}

Execution Result	The confidants of Li Bai included Meng Haoran and Du Fu, among others.

facts. Additionally, to mitigate the preference gap between the QR module and the QA model, we devise a training strategy that leverages QA model feedback to further refine QR. This strategy evaluates QR quality through QA model responses, then optimizes QR parameters to improve the downstream QA model. Experimental findings reveal that, relative to existing conversational QA techniques, our method consistently achieves significant gains across two datasets. On WebQuestionsSP-C, our full model boosts F1 by up to 6.7% and Acc by 13.5% over the leading baseline, CoTKR. These outcomes demonstrate that aligning both modules with downstream QA feedback substantially progresses conversational QA.

REFERENCES

[1] Chen J, Lu Y, Zhang S, et al. Storysparkqa: Expert-annotated qa pairs with real-world knowledge for childrens story-based learning[C]. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024: 17351-17370.

[2] Alkawsi G, Al-Sharafi M A, Alajmi Q. Determinants of open educational resources adoption and their impact on educational sustainability: A cross-cultural comparison using SEM-fsQCA[J]. Education and Information Technologies, 2025, 30(8): 10039-10072.

[3] Saparina I, Lapata M. Disambiguate First Parse Later: Generating Interpretations for Ambiguity Resolution in Semantic Parsing[J]. arXiv preprint arXiv:2502.18448, 2025.

[4] Faustini P, Chen Z, Fetahu B, et al. Answering Unanswered Questions through Semantic Reformulations in Spoken QA[C]. The 61st Annual Meeting Of The Association For Computational Linguistics, 2023: 729-743.

[5] Chen Z, Zhao J, Fang A, et al. Reinforced Question Rewriting for Conversational Question Answering[C]. Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track, 2022: 357-370.

[6] Wu Y, Huang Y, Hu N, et al. CoTKR: Chain-of-Thought Enhanced Knowledge Rewriting for Complex Knowledge Graph Question Answering[C]. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, 2024: 3501-3520.

[7] Mao W, Wang R, Guo J, et al. Enhancing text-to-SQL parsing through question rewriting and execution-guided refinement[C]. Findings of the Association for Computational Linguistics ACL, 2024: 2009-2024.

[8] Kongyoung S, MacDonald C, Ounis I. Multi-Task Learning of Query Generation and Classification for Generative Conversational Question Rewriting[C]. The 2023 Conference on Empirical Methods in Natural Language Processing, 2023: 13667-13678.

[9] Kaiser M, Saha Roy R, Weikum G. Robust training for conversational question answering models with reinforced reformulation generation[C]. Proceedings of the 17th ACM International Conference on Web Search and Data Mining, 2024: 322-331.

[10] Chai L, Xiao D, Yan Z, et al. QURG: Question rewriting guided context-dependent text-to-SQL semantic parsing[C]. Pacific Rim International Conference on Artificial Intelligence. Singapore: Springer Nature Singapore, 2023: 275-286.

[11] Raposo G, Ribeiro R, Martins B, et al. Question rewriting? assessing its importance for conversational question answering[C]. European Conference on Information Retrieval, 2022: 199-206.

[12] Liu L, Hill B, Du B, et al. Conversational Question Answering with Language Models Generated Reformulations over Knowledge Graph[C]. Findings of the Association for Computational Linguistics ACL, 2024: 839-850.

[13] Ren G, Ni X, Malik M, et al. Conversational query understanding using sequence to sequence modeling[C]. Proceedings of the 2018 World Wide Web Conference, 2018: 1715-1724.

[14] Yu S, Liu J, Yang J, et al. Few-shot generative conversational query rewriting[C]. Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, 2020: 1933-1936.

[15] Elgohary A, Peskov D, Boyd-Graber J. Can You Unpack That? Learning to Rewrite Questions-in-Context[C]. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019: 5918-5924.

[16] Vakulenko S, Longpre S, Tu Z, et al. Question rewriting for conversational question answering[C]. Proceedings of the 14th ACM international conference on web search and data mining, 2021: 355-363.

[17] Fader A, Zettlemoyer L, Etzioni O. Open question answering over curated and extracted knowledge bases[C]. Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, 2014: 1156-1165.

[18] Li S, Jiang W, Si P, et al. Enhancing dialogue generation with conversational concept flows[C]. Findings of the Association for Computational Linguistics, 2023: 1514-1525.

[19] Raffel C, Shazeer N, Roberts A, et al. Exploring the limits of transfer learning with a unified text-to-text transformer[J]. Journal of machine learning research, 2020, 21(140): 1-67.

[20] Wang Z, Zhang J, Feng J, et al. Knowledge graph embedding by translating on hyperplanes[C]. Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014, 28(1): 1112-1119.

[21] Pennington J, Socher R, Manning C D. Glove: Global vectors for word representation[C]. Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), 2014: 1532-1543.

[22] GLM T, Zeng A, Xu B, et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools[J]. arXiv preprint arXiv:2406.12793, 2024.

[23] Yih W, Richardson M, Meek C, et al. The value of semantic parse labeling for knowledge base question answering[C]. Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, 2016: 201-206.

[24] Su Y, Sun H, Sadler B, et al. On generating characteristic-rich question sets for qa evaluation[C]. Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 2016: 562-572.

[25] Christmann P, Saha Roy R, Weikum G. Explainable conversational question answering over heterogeneous sources via iterative graph neural networks[C]. Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval, 2023: 643-653.