

Beyond Chains: Bridging Large Language Models and Knowledge Bases in Complex Question Answering

Anonymous ACL submission

Abstract

Knowledge Base Question Answering (KBQA) aims to answer natural language questions using structured knowledge from KBs. While LLM-only approaches offer generalization, they suffer from outdated knowledge, hallucinations, and lack of transparency. Chain-based KG-RAG methods address these issues by incorporating external KBs, but are limited to simple chain-structured questions due to the absence of planning and logical structuring. Inspired by semantic parsing methods, we propose **PDRR**: a four-stage framework consisting of **Predict, Decompose, Retrieve, and Reason**. Our method first **predicts** the question type and **decomposes** the question into structured triples. Then **retrieves** relevant information from KBs and guides the LLM as an agent to **reason** over and complete the decomposed triples. Experimental results show that our proposed KBQA model, PDRR, consistently outperforms existing methods across different LLM backbones and achieves superior performance on various types of questions.

1 Introduction

Knowledge bases (KBs) offer rich, structured repositories of world knowledge, where facts are organized as (subject, relation, object) triples. Large-scale KBs such as Freebase (Bollacker et al., 2008), DBpedia (Lehmann et al., 2015), Wikidata (Pelissier Tanon et al., 2016), and YAGO (Suchanek et al., 2007) provide useful resources for downstream applications. Knowledge base question answering (KBQA) leverages this structured information to translate natural language question queries into precise, verifiable answers, which poses challenges since it requires accurate multi-hop reason-

The anonymous GitHub repository is available at <https://anonymous.4open.science/r/conferencesubmission/README.md>.

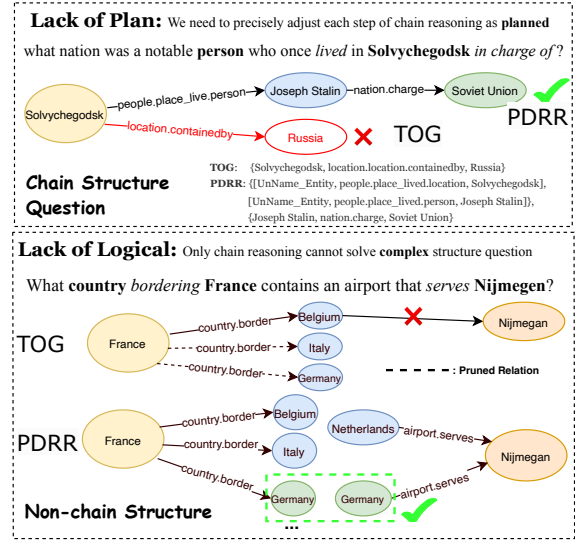


Figure 1: Drawbacks of ToG (a chain-based KG-RAG approach). ToG and similar chain-based KG-RAG methods lack a planning module for explicit reasoning control and are limited to chain-type questions due to their insufficient logical structuring. Our PDRR framework resolves both issues.

ing. This capability is essential for domains that require accurate and verifiable retrieval of facts.

In response to the costly human effort required to annotate training data, recent studies have turned to leveraging large language models (LLMs) based methods for KBQA, for example, IO (Brown et al., 2020), CoT (Wei et al., 2022), and SC (Wang et al., 2022). These methods harness broad, pre-trained knowledge of LLMs and exhibit strong generalization across datasets. Nonetheless, LLM’s knowledge may not reflect recent or domain-specific facts; they could also exhibit hallucinations, which is difficult to audit or verify, leading to unfaithful and unsupported response. To address these limitations, knowledge-graph retrieval-augmented generation (KG-RAG) methods further extend LLMs with retrieved structural knowledge. Various studies such as ToG (Sun et al., 2023), GoG (Xu et al.,

2024), and CoK (Li et al., 2023) leverage chain-based KG-RAG approaches that performs reasoning over retrieved chains of KG facts. Specifically, they infer intermediate “bridge” entities at each hop, which repeatedly fetches triplets linked to the current entities from the KG, and uses the LLM to choose which relation and entity to retain for the next step. These chain-based KG-RAG methods ground multi-step reasoning in explicit KG paths, thereby mitigating outdated knowledge, reducing hallucinations through factual grounding and rendering each inference step to be transparent.

Despite their strengths, existing chain-based KG-RAG methods exhibit two principal limitations as shown in Figure 1. First, they reason over the question in a holistic manner and lack a mechanism to structure the inference into targeted sub-tasks. In the upper half “Lack of Plan” example, one should first identify the notable person and then determine the nation they governed. Instead, ToG treats the query as an indivisible whole, leading it to select the superficially matching *location.containedby* relation rather than the correct *people.place_live.person*.

Second, chain-based approaches cannot handle richer logical structures beyond simple, linear hops, such as conjunction questions that require the intersection of multiple constraint sets. Logic structure refers to the structural form of reasoning required to answer a question. In the lower half “Lack of Logical” example, the answer hinges on intersecting (a) countries bordering France and (b) countries with airports serving Nijmegen. ToG’s linear, single-path search prunes away additional candidates after the first hop, making it incapable of resolving the conjunction and thus yielding an incorrect result.

Motivated by these challenges, we propose the **Predict–Decompose–Retrieve–Reason (PDRR)** framework. Inspired by pre-LLM semantic-parsing (SP-based) methods (Lan et al., 2022) that translate questions into executable KB logic forms, PDRR introduces a planning module (Predict, Decompose) for structuring multi-step inference, followed by a retrieval and reasoning module (Retrieve, Reason) that grounds the plan in KB facts and guides the LLM to execute it step by step, effectively simulating logical form execution over KGs. Finally, the question answering module utilizes the reasoning triples provided by the previous step to generate an answer.

First, our Predict stage classifies each query by type (i.e., chain or parallel) to determine an appro-

priate reasoning strategy and plan structure before retrieval begins. Next, the Decompose stage converts the question into a set of partial KG triples that reflect the plan, which breaks the overall query into manageable inference units, ensuring each reasoning step is focused and auditable. The Retrieve stage issues targeted KB lookups to fill in missing triple elements, which grounds the planned inference steps in factual knowledge. Finally, the Reason stage leverages the LLM to verify and complete each triple in accordance with the plan. By unifying logical-form-style planning with LLM-driven execution, PDRR supports a variety of reasoning patterns while preserving transparency.

We evaluate PDRR on two standard KBQA benchmarks: CWQ (Talmor and Berant, 2018) and WebQSP (Berant et al., 2013). Experimental results show that our proposed KBQA model, PDRR, consistently outperforms existing methods across different LLM backbones and achieves superior performance on various types of questions, demonstrating the effectiveness of PDRR’s explicit planning module and the precise control over retrieval and reasoning modules.

2 Related Work

Sementic Parsing-Based Methods Before LLMs, KBQA was dominated by SP-based approaches, which generate logical forms to query structured KBs. CBR-KBQA (Das et al., 2021) combines a non-parametric memory of question–logical form pairs with a parametric retriever to guide logical form generation. TemplateQA (Zheng et al., 2018) bypasses complex parsing by matching questions to a large set of binary templates. SPARQA (Sun et al., 2020) uses a skeleton grammar and a BERT-based coarse-to-fine parser to handle complex questions. While interpretable, SP-based methods are limited by incomplete KBs and the need for additional model training.

LLM retrieval augmented methods The emergence of LLMs has led to LLM-only approaches that require no additional training and leverage internalized general knowledge to mitigate KB incompleteness, such as IO (Brown et al., 2020), CoT (Wei et al., 2022), and SC (Wang et al., 2022). Despite their simplicity, these methods struggle with outdated knowledge, limited reasoning transparency, and hallucinations. To mitigate these issues, LLM+KG approaches have been developed.

ToG (Sun et al., 2023) uses an LLM to guide reasoning over KGs by dynamically selecting relations and entities. CoK (Li et al., 2023) adaptively selects KGs and performs stepwise retrieval. GoG (Xu et al., 2024) augments KG coverage by letting LLMs infer missing links.

However, most chain-based KG-RAG methods lack explicit planning and are confined to chain reasoning, limiting reasoning control. Recent methods like chatKBQA (Luo et al., 2023a) and RoG (Luo et al., 2023b) introduce planning by fine-tuning LLMs to generate logic forms or reasoning paths, but they remain training-intensive. To address this, we propose PDRR: a training-free framework with explicit planning, enabling precise reasoning control and handling of complex question structures.

3 Preliminary

In this section, we first introduce Knowledge Graphs (KGs). Then, we use notations of KGs to describe reasoning triples and Knowledge Base Question Answering (KBQA).

Knowledge Graphs A KG \mathcal{G} consists of factual triples $(e^h, r, e^t) \in \mathcal{G} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, where e^h, e^t represent head and tail entities, and \mathcal{E} and \mathcal{R} denote the sets of entities and relations, respectively.

Reasoning Triples The Reasoning Triples \mathcal{T}^q denote the set of triples retrieved from the KG \mathcal{G} to answer the question q , formally defined as:

$$\mathcal{T}^q = \{\mathcal{T}_n^q = (e_n^h, r_n, e_n^t) \mid n = 1, 2, \dots, m\} \subseteq \mathcal{G}$$

Here, \mathcal{T}_n^q denotes the n -th reasoning triple.

Knowledge Base Question Answering KBQA involves reasoning over a KG to answer natural language questions. Formally, given a question q and a knowledge graph \mathcal{G} , the objective is to learn a function f that returns answers $a \in \mathcal{A}_q$ based on the information in \mathcal{G} , i.e., $a = f(q, \mathcal{G})$. In line with previous studies (Sun et al., 2019; Jiang et al., 2022), we assume that the entities $e_q \in \mathcal{Q}_q$ mentioned in the question and the gold answers $a \in \mathcal{A}_q$ are annotated and linked to entities in the KG, such that $\mathcal{Q}_q, \mathcal{A}_q \subseteq \mathcal{E}$, where \mathcal{E} denotes the entity set of \mathcal{G} .

4 Approach

We propose a training-free method to address the lack of planning and logical structure handling in existing chain-based KG-RAG approaches. Our

approach consists of three main modules: Plan, Retrieval and Reasoning, and Question Answering. The Plan module predicts the question type (i.e., chain or parallel) and decomposes the question accordingly. The Retrieval and Reasoning module retrieves the KB for relevant facts and employs the LLM to reason over and complete the decomposed triples. Finally, the Question Answering Module generates the final answer. Figure 2 illustrates the overall framework of our approach. All prompts used in the approach are provided in Appendix A.3.

4.1 Plan Module

Our Plan module resembles logical form generation in SP-based methods. It first predicts the question type (i.e., chain or parallel) and then decomposes the question into decomposition triples, which serve a similar role to logical forms by guiding subsequent retrieval and reasoning. The motivation is to support not only chain-structured questions but also more complex reasoning patterns, while enabling precise downstream control through the generated decomposition triples.

4.1.1 Question Type Prediction

Given a question q , we first use few-shot learning with LLMs predict its structural type, which we classify into two main categories: chain structure q_c and parallel structure q_p . The chain structure is the most common type in chain-based KG-RAG methods. Its reasoning process is sequential, where each step depends on the bridge entity e_n^B identified in each step of reasoning triple \mathcal{T}_n^q . Therefore, the reasoning steps are interdependent and must be performed in order. In contrast, the parallel structure comprises multiple logically independent sub-steps that can be executed concurrently. More detail refers to section 4.2.3.

4.1.2 Question Decomposition

Based on the predicted question type, question q is decomposed into KG-style triples $\mathcal{T}^{q,D}$, providing precise control for downstream processing.

For chain-structured questions q_c , we identify all bridge entities e_n^B (denoted as *entity#index*) to distinguish them in multi-hop reasoning. Each decomposition triple $\mathcal{T}_n^{q_c,D} = (e_n^h, r_n, e_n^t)$ connects entities via relation r_n , with adjacent triples linked through e_n^B such that $e_{n+1}^h = e_n^t = e_n^B$, forming a reasoning chain. For example, in Figure 2, the bridge entity *artist#1* connects two steps: identifying the *artist#1* of Country Nation World Tour,

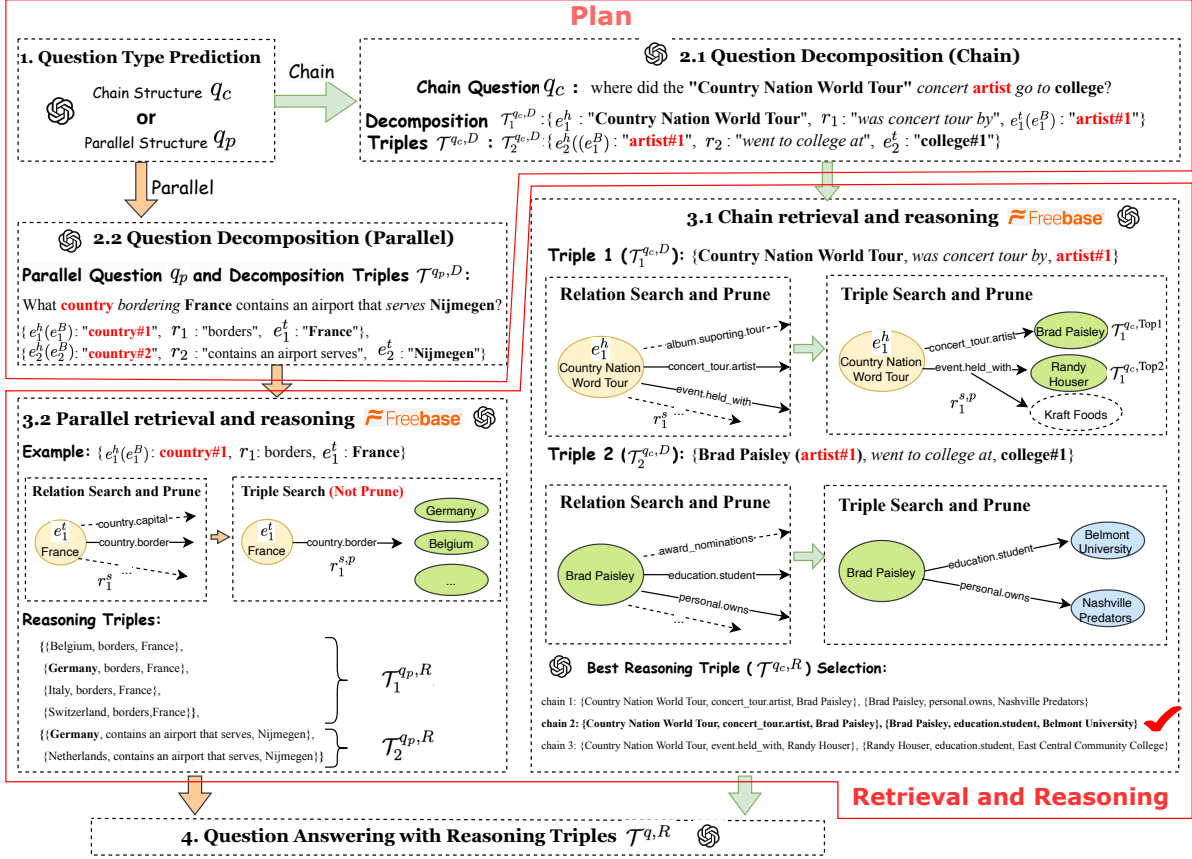


Figure 2: The framework of the PDRR method. The process follows the Predict-Decompose-Retrieve-Reason pipeline. Dashed lines and circles indicate pruned components with low relevance to the specific decomposed triple $\mathcal{T}_n^{q,D}$. Red-labeled *entity#index* (e.g., **artist#1**) elements denote key bridge entities e_n^B that are essential for reasoning.

and determining the college that the *artist#1* attended. The corresponding decomposition triples are illustrated in the figure 2.

For parallel-structured questions, we similarly identify bridge entities and construct KG-style triples along independent reasoning paths. Unlike chain structures, these steps are mutually independent and can be executed in parallel. For example, in Figure 2, the reasoning involves two conditions: identifying *country#1* bordering France and *country#2* with an airport serving Nijmegen, as reflected in the decomposition triples.

4.2 Retrieval and Reasoning Module

In this module, we adopt different reasoning strategies based on the identified question type. Similar to how SP-based methods use logic forms to retrieve information from KBs, we leverages KB knowledge and employs the LLM as an agent to complete missing information in the decomposition triples $\mathcal{T}^{q,D}$ generated by the plan module. The SPARQL code used for Freebase retrieval in this

module is presented in Appendix A.4.

4.2.1 Chain Retrieval and Reasoning

Given a chain-type question q_c and its decomposition triples $\mathcal{T}^{q_c,D}$, we sequentially complete them by identifying bridge entities step by step, enabling accurate construction of reasoning triples and final answer derivation.

Complete first Decomposition triples We first extract $\mathcal{T}_1^{q_c,D} = (e_1^h, r_1, e_1^t(e_1^B))$ and retrieve the bridge entity e_1^B using the method illustrated in Figure 2.

- **Relation Search and Prune** We first apply SPARQL fuzzy matching to obtain the Freebase entity ID of the non-bridge entity in the triple thought its string. In the **Search** phase, we query the KB with this entity ID to obtain all connected relations r_1^s . In the case of Figure 2, we first identify the entity ID of e_1^h (*Country Nation World Tour*), and then retrieve all connected relations r_1^s .

We then **prune** the retrieved relation set r_1^s . Unlike ToG (Sun et al., 2023), which ranks relations $r_{1,i}^s$ by their similarity $Sim(r_{1,i}^s, q)$ to the entire question q and may overlook local information, we rank them by their similarity $Sim(r_{1,i}^s, \mathcal{T}_1^{q_c, D})$ to the specific decomposition triple $\mathcal{T}_1^{q_c, D}$. This yields the pruned set $r_1^{s, p}$. This helps avoid errors such as the one in Figure 1, where ToG selects *contained by*, which aligns with the global semantics of “What nation” instead of the correct local relation *people.place_live.person*.

By aligning pruning with the decomposition triple, our approach ensures precise control and preserves step-level semantics. As shown in the Figure 2, we obtain a pruned relation set $r_1^{s, p} = \{concert_tour.artist, event.held_with\}$, which is most similar to the decomposition triple $\mathcal{T}_1^{q_c, D}$.

- **Triple Search and Prune** Given ID of e_1^h and the pruned relation set $r_1^{s, p}$, we use SPARQL to **search** all tail entities (bridge entities) $e_1^t(e_1^B)$ to construct the candidate triple set $\mathcal{T}_1^{q_c, Cand}$. In the Triple **Prune** stage, instead of pruning only entities as in ToG, which inherits the same limitations as its relation pruning, we rank candidate triples $\mathcal{T}_1^{q_c, Cand_i}$ based on their similarity $Sim(\mathcal{T}_1^{q_c, Cand_i}, \mathcal{T}_1^{q_c, D})$ to the decomposition triple $\mathcal{T}_1^{q_c, D}$ and retain the top two: $\mathcal{T}_1^{q_c, Top1} = \{Country\ Nation\ World\ Tour, concert_tour.artist, Brad\ Paisley\}$, and $\mathcal{T}_1^{q_c, Top2}$.

Complete rest Decomposition triples Through Relation Search and Prune and Triple Search and Prune, we identify the top two bridge entities in first decomposition triple: $e_1^{B, top1}$ and $e_1^{B, top2}$. We then replace the bridge entity in next decomposition triple $\mathcal{T}_2^{q_c, D}$ with the bridge entities from the previous step and repeat the same procedure until all decomposition triples in $\mathcal{T}^{q_c, D}$ are completed.

Choose Best Reasoning Triples We apply beam search to retain the top-2 triples at each hop, preserving sufficient information. For example, for a 2-hop question, this yields 4 reasoning triples. The LLM then selects the most relevant triple: $\mathcal{T}^{q_c, R}$ for answering the question.

4.2.2 Parallel Retrieval and Reasoning

Parallel-structured questions q_p follow a non-sequential reasoning logic, where the decomposi-

tion triples $\mathcal{T}^{q_p, D}$ are mutually independent, allowing all retrieval and reasoning steps to be executed concurrently.

For example, given the first decomposition triple $\mathcal{T}_1^{q_p, D}$ in Figure 2, we perform relation search and pruning, followed by triple search, using the same procedure as in the chain setting. Unlike the chain case, we skip triple pruning, as all reasoning triples are needed for the final answer. This process results in a set of reasoning triples $\mathcal{T}_1^{q_p, R}$.

We repeat this process for all decomposition triples $\mathcal{T}_k^{q_p, D}$ to obtain the corresponding reasoning triples $\mathcal{T}_k^{q_p, R}$, where k indexes each decomposition triple.

4.2.3 Discussion of Chain and Parallel Reasoning

While our question type classification is based on reasoning strategy, another key criterion is the number of bridge entities that the model should memorize during the reasoning process. Chain-structured questions require exactly one bridge entity between each pair of triples to support step-by-step reasoning, whereas parallel-structured questions allow multiple bridge entities without such constraints.

When applying chain reasoning to parallel-structured questions, retaining all retrieved triples without pruning can still lead to correct answers. For example, in Figure 2, all *countries* (bridge entities) bordering France are considered and individually checked for containing an airport that serves Nijmegen, which can lead to the right answer. However, this significantly increases computational cost. Introducing pruning in chain reasoning mitigates this cost but risks discarding correct answers.

Therefore, parallel reasoning can be regarded as a complementary strategy to chain reasoning, particularly in cases with numerous bridge entities where computational efficiency becomes critical.

In practice, the logical structures or types of questions are not limited to just chain and parallel. For example, SP-based methods often train models using the training set to generate logic forms that can handle more complex logical structures. In our work, we focus only on chain and parallel logic structures because a well-designed parallel structure that complements the chain structure, combined with the general knowledge and reasoning capabilities of LLMs, is sufficient to address the majority of KBQA questions. This also allows our method to remain training-free, avoiding additional

| Question type | Question Example | Canonical Reasoning Logic | Expected Reasoning Type |
|---------------|--|--|-------------------------|
| Composition | Where did the "Country Nation World Tour" concert artist go to college? | First, identify the artist of the concert tour "Country Nation World Tour," and then determine which college this artist attended. | Chain |
| Conjunction | What country bordering France contains an airport that serves Nijmegen? | Find the intersection of two sets: (1) all countries that border France, and (2) all countries that contain an airport that serves Nijmegen. | Parallel |
| Comparative | Which of the countries bordering Mexico have an army size of less than 1050? | First, identify all countries bordering Mexico, and then select those with an army size of less than 1050. | Parallel |
| Superlative | What movies does Taylor Lautner play in and is the film was released earliest? | First, identify all films in which Taylor Lautner appeared, and then find the one that was released the earliest. | Parallel |

Table 1: Examples of the four question types in the CWQ dataset, their corresponding canonical reasoning logic, and the expected reasoning type for each.

| Method | CWQ | | | | | WebQSP |
|---|-------------|-------------|-------------|-------------|-------------|-------------|
| | All | Composition | Conjunction | Comparative | Superlative | All |
| <i>Without KB Knowledge, Without Training</i> | | | | | | |
| IO (Brown et al., 2020) | 44.3 | 41.3 | 50.6 | 33.8 | 27.9 | 70.8 |
| CoT (Wei et al., 2022) | 44.8 | 44.2 | 49.5 | 30.5 | 27.4 | 72.1 |
| PDR (ours) | 45.7 | 49.3 | 46.4 | 31.5 | 26.4 | 68.9 |
| <i>With KB Knowledge, With Training</i> | | | | | | |
| SPARQA (Sun et al., 2020) ^α | 31.6 | - | - | - | - | - |
| UniKGQA (Jiang et al., 2022) ^β | 51.2 | - | - | - | - | 77.2 |
| RoG (Luo et al., 2023b) ^γ | 62.6 | - | - | - | - | 85.7 |
| CBR-KBQA (Das et al., 2021) ^σ | 67.1 | - | - | - | - | 69.9 |
| <i>With KB Knowledge, Without Training</i> | | | | | | |
| ToG (Sun et al., 2023) | 48.9 | 49.9 | 50.1 | <u>42.7</u> | 37.1 | 80.2 |
| PDRR (ours) | <u>59.6</u> | 66.2 | <u>59.1</u> | 38.5 | 34.5 | <u>79.2</u> |
| PDRR(gold type)(ours) | 62.2 | <u>65.5</u> | 64.6 | 43.7 | <u>36.6</u> | - |

Table 2: Hit@1 accuracy results with different baselines on KGQA datasets. **Bold** denotes the best performance among the training-free methods, and underline indicates the second-best. Results marked with superscripts $\alpha, \beta, \gamma, \sigma$ are taken directly from the original papers. All other results are reproduced using GPT-4o as the backbone model. Our proposed methods include PDR, PDRR, and PDRR (gold type). In the gold type setting (an ablation setup), ground-truth question types from CWQ are used instead of LLM predictions. Composition-type questions are handled with chain reasoning, while all others use parallel reasoning, as we explained in Appendix A.1.3.

computational costs.

4.3 Question Answering Module

Given the original question q , the decomposition triples $\mathcal{T}^{q,D}$, and the reasoning triples $\mathcal{T}^{q,R}$. The LLM is guided to answer the question q by leveraging the information in retrieved reasoning triples $\mathcal{T}^{q,R}$ in accordance with the reasoning logic encoded in $\mathcal{T}^{q,D}$.

5 Experiment

5.1 Experiment Setup

Dataset To evaluate performance beyond simple chain-structured QA, we evaluate the model on the CWQ (Talmor and Berant, 2018) test set (3,531 questions), which includes four question types: composition (45%), conjunction (45%), compar-

ative (5%), and superlative (5%). Composition questions are expected to follow chain reasoning, while the other three suit parallel reasoning. Table 1 shows examples, reasoning logic, and expected types, with further explanations in Appendix A.1.3. We also evaluate the model on WebQSP (Berant et al., 2013) (license: CC License) test set (1,639 questions) for additional validation. The questions in the WebQSP dataset are all of chain structure, with most being one- or two-hop questions.

Evaluation metrics Consistent with prior studies (Luo et al., 2023b; Sun et al., 2023), we adopt Hits@1 as the evaluation metric, which reflects the percentage of questions for which the top-1 rank predicted answer is correct.

Baseline We divide the baselines into three groups. The first group includes LLM-only meth-

ods including IO (Brown et al., 2020) and CoT (Wei et al., 2022), and our ablated variant PDR (Predict-Decompose-Reason), which removes the retrieve stage from PDRR. The second group incorporates external KBs and require additional training, including UniKBQA (Jiang et al., 2022), ROG (Luo et al., 2023b), CBR-KBQA (Das et al., 2021), and SPARQA (Sun et al., 2020). The third group leverages external KBs without additional training, such as ToG (Sun et al., 2023) and our model, PDRR.

Implementation For all intermediate steps, the max token length is 256, and 1024 for final answer generation. A temperature of 0.1 is used to reduce hallucinations and improve control. We adopt 5-shot prompting for answer generation and question type classification, and 3-shot for all other components. More detail refers to Appendix A.1.2

5.2 Result

5.2.1 Main Results

In this section, we compare PDRR with various baselines across KGQA datasets as shown in Table 2. On the more challenging CWQ dataset, which features diverse question structures beyond simple chains, PDRR performs competitively with training-based methods and outperforms training-free ones like ToG by nearly 10% in accuracy.

By question type, PDRR significantly surpasses ToG on composition questions (66.2 vs. 49.9), highlighting the effectiveness of the Plan Module in guiding step-by-step reasoning. It also achieves higher accuracy on conjunction questions and remains competitive on comparative and superlative types, showing that augmenting the Retrieval and Reasoning Module with complementary parallel reasoning enhances performance on complex, diverse questions.

On the simpler WebQSP dataset, which includes only composition-type questions, PDRR slightly underperforms ToG due to occasional errors in question type prediction.

5.2.2 Different Backbone Models

To evaluate the robustness of PDRR, we test its performance on the CWQ dataset using different LLM backbones to assess its effectiveness beyond GPT-4o. Specifically, we compare CoT, ToG, PDRR, and PDRR (gt) across four models: GPT-3.5-turbo, DeepSeek-V3, and GPT-4o. Additionally, we evaluate CoT, PDRR, and PDRR (gt) on LLaMA-3.3B-Instruct. As shown in Table 3, PDRR consistently

| Method | CWQ | | | | |
|---------------------------|-------------|-------------|-------------|-------------|-------------|
| | All | Compo | Conju | Compa | Super |
| <i>GPT-3.5-turbo-0125</i> | | | | | |
| CoT | 36.8 | 33.8 | 43.0 | <u>26.8</u> | 20.8 |
| ToG | 30.9 | 30.1 | 35.6 | 16.4 | 15.8 |
| PDRR | <u>37.7</u> | <u>34.5</u> | <u>43.9</u> | 26.3 | <u>25.9</u> |
| PDRR(gt) | 40.7 | 36.1 | 48.6 | 29.6 | 26.4 |
| <i>Llama3.3-Instruct</i> | | | | | |
| CoT | 46.0 | 41.5 | 52.4 | <u>42.7</u> | 33.5 |
| PDRR | <u>54.1</u> | <u>54.9</u> | <u>57.7</u> | 36.2 | <u>38.1</u> |
| PDRR(gt) | 58.7 | 56.4 | 65.0 | 46.5 | 39.1 |
| <i>Deepseek-V3</i> | | | | | |
| CoT | 44.3 | 43.0 | 48.5 | 35.7 | 29.9 |
| ToG | 48.0 | 48.5 | 51.7 | 35.7 | 27.4 |
| PDRR | <u>56.0</u> | 60.3 | <u>57.0</u> | <u>41.3</u> | <u>30.5</u> |
| PDRR(gt) | 57.1 | 60.3 | 57.5 | 43.7 | 42.6 |
| <i>GPT-4o-2024-11-20</i> | | | | | |
| CoT | 44.8 | 44.2 | 49.5 | 30.5 | 27.4 |
| ToG | 48.9 | 49.9 | 50.1 | <u>42.7</u> | 37.1 |
| PDRR | <u>59.6</u> | 66.2 | <u>59.1</u> | 38.5 | 34.5 |
| PDRR(gt) | 62.2 | <u>65.5</u> | 64.6 | 43.7 | <u>36.6</u> |

Table 3: Performance of PDRR with different backbone models on overall accuracy and by question type in the CWQ dataset. ‘gt’ denotes the gold type setting.

outperforms across all models. The performance gap is modest on GPT-3.5-turbo but becomes more significant on the stronger LLMs. These results confirm that PDRR is robust and not dependent on any specific language model.

5.2.3 Discussion of Question Structure

A core component of PDRR is the use of the Plan Module to predict the question structure type and apply corresponding decomposition and reasoning strategies. Thus, investigating this process is essential.

Question Type Prediction We analyze how different LLMs predict the structure type (chain or parallel) for various CWQ question types, as shown in Figure 3.

For composition-type questions, most LLMs correctly predict a chain structure, consistent with expectations. For conjunction-type questions, which are best handled by parallel reasoning, GPT-3.5-turbo and DeepSeek-V3 predict the correct structure in most cases, while Llama3.3-Instruct and GPT-4o achieve a much lower rate. Despite this, the results in Table 2 reveal that many misclassified conjunction questions still lead to correct answers, as they can be effectively addressed by either chain or parallel reasoning.




Figure 3: Predicted question structure types (chain or parallel) by different LLMs on various question types in the CWQ dataset. Blue indicates that the LLM predicts the question as a chain structure, while orange indicates a parallel structure.




Figure 4: Hits@1 accuracy of different question types in the CWQ dataset under chain and parallel reasoning strategies. The evaluation includes 500 chain, 500 intersection, 213 comparative, and 197 superlative questions. GPT-4o is used as the LLM backbone.

This phenomenon can be explained by the number of bridge entities: when only one is involved, both chain and parallel reasoning are generally effective; with more, parallel reasoning becomes notably more robust. This supports the discussion in Section 4.2.3. Two detailed case examples in Appendix A.2.1 further illustrate and explain how both strategies succeed when only a single bridge entity is present.

Additional case studies and analysis of comparative and superlative questions are provided in Appendix A.2.2.

Different reasoning strategies We evaluate the effectiveness of chain and parallel reasoning across four CWQ question types using Hits@1 accuracy, as shown in Figure 4. Chain reasoning performs significantly better than parallel on composition-type questions (67.2% vs. 57.0%), while parallel reasoning clearly outperforms chain




Figure 5: Hits@1 accuracy on the first 500 composition-type questions in CWQ using chain reasoning with different number of retained triples.

on conjunction-type questions (71.2% vs. 56.0%) and slightly outperforms it on comparative and superlative questions.

These results align with our hypothesis: composition-type questions favor chain reasoning, whereas the others benefit more from parallel reasoning. In contrast, methods like ToG, which rely solely on chain reasoning, underperform on non-chain questions. This highlights the need to adapt reasoning strategies to question type.

5.2.4 Ablation Study

We aim to explore the impact of the Number of Retained Triples within Chain Reasoning on the experimental results. As shown in Figure 5, increasing the number of retained triples from top-1 to top-2 significantly improves performance (61.2% to 67.2%). Although top-3 offers a minor additional gain (+0.6%), we adopt top-2 to balance accuracy and computational cost. More ablation study refers to Appendix A.2.3. And we present some case study in Appendix A.2.2 and A.2.4.

6 Conclusion

To address the limitation of chain-based KG-RAG methods, which are restricted to simple chain-structured questions due to the lack of planning and logical structuring. We propose PDRR, which first **predicts** the question type and **decomposes** the question into structured triples. Then **retrieves** relevant information from KBs and guides the LLM as an agent to **reason** over and complete the decomposed triples. Experimental results show that PDRR consistently outperforms existing methods across various LLM backbones, with up to a 10% gain on CWQ using GPT-4o. It also performs robustly across diverse question types. Additionally, our in-depth analysis of question structures reveals that LLMs rely not only on the structural form of the question but also on the number of bridging entities they must retain during reasoning.

Limitations

Unnecessary Decomposition on Simple Dataset

When applied to simpler datasets like WebQSP, PDRR tends to perform unnecessary decomposition by converting inherently 1-hop questions into 2-hop structures. This introduces superfluous retrieval and reasoning steps for questions that could have been answered in the first step, ultimately leading to incorrect results.

Limitations on questions with highly complex logical structures

Even though a well-designed parallel structure that complements the chain structure, combined with the general knowledge and reasoning capabilities of LLMs, is sufficient to handle most KBQA questions, it remains inadequate for highly complex cases, particularly during the question decomposition phase in the plan module. Although we incorporate question type prediction prior to decomposition to improve accuracy, the current method struggles with highly complex logical structure. In these cases, relying solely on the LLM for decomposition leads to suboptimal retrieval and reasoning performance. Due to the absence of benchmark datasets for hybrid questions, we leave this challenge for future work. In the next stage, we plan to construct a dedicated hybrid dataset and design new methods tailored to such complex questions.

Acknowledgements

Ethics Statement

This study complies with the [ACL Ethics Policy](#).

References

Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.

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*, pages 1247–1250.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot

learners. *Advances in neural information processing systems*, 33:1877–1901.

Rajarshi Das, Manzil Zaheer, Dung Thai, Ameya Godbole, Ethan Perez, Jay-Yoon Lee, Lizhen Tan, Lazaros Polymenakos, and Andrew McCallum. 2021. Case-based reasoning for natural language queries over knowledge bases. *arXiv preprint arXiv:2104.08762*.

Jinhao Jiang, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Unikgqa: Unified retrieval and reasoning for solving multi-hop question answering over knowledge graph. *arXiv preprint arXiv:2212.00959*.

Yunshi Lan, Gaole He, Jinhao Jiang, Jing Jiang, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Complex knowledge base question answering: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 35(11):11196–11215.

Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleeef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web*, 6(2):167–195.

Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing. 2023. Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources. *arXiv preprint arXiv:2305.13269*.

Haoran Luo, Zichen Tang, Shiyao Peng, Yikai Guo, Wentai Zhang, Chenghao Ma, Guanting Dong, Meina Song, Wei Lin, Yifan Zhu, et al. 2023a. Chatkbqa: A generate-then-retrieve framework for knowledge base question answering with fine-tuned large language models. *arXiv preprint arXiv:2310.08975*.

Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023b. Reasoning on graphs: Faithful and interpretable large language model reasoning. *arXiv preprint arXiv:2310.01061*.

Thomas Pellissier Tanon, Denny Vrandečić, Sebastian Schaffert, Thomas Steiner, and Lydia Pintscher. 2016. From freebase to wikidata: The great migration. In *Proceedings of the 25th international conference on world wide web*, pages 1419–1428.

Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706.

Haitian Sun, Tania Bedrax-Weiss, and William W Cohen. 2019. Pullnet: Open domain question answering with iterative retrieval on knowledge bases and text. *arXiv preprint arXiv:1904.09537*.

- 659 Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo
660 Wang, Chen Lin, Yeyun Gong, Lionel M Ni, Heung-
661 Yeung Shum, and Jian Guo. 2023. Think-on-
662 graph: Deep and responsible reasoning of large lan-
663 guage model on knowledge graph. *arXiv preprint*
664 *arXiv:2307.07697*.
- 665 Yawei Sun, Lingling Zhang, Gong Cheng, and Yuzhong
666 Qu. 2020. Sparqa: skeleton-based semantic pars-
667 ing for complex questions over knowledge bases. In
668 *Proceedings of the AAAI conference on artificial in-*
669 *telligence*, volume 34, pages 8952–8959.
- 670 Alon Talmor and Jonathan Berant. 2018. [The web as](#)
671 [a knowledge-base for answering complex questions](#).
672 In *Proceedings of the 2018 Conference of the North*
673 *American Chapter of the Association for Computa-*
674 *tional Linguistics: Human Language Technologies,*
675 *Volume 1 (Long Papers)*, pages 641–651, New Or-
676 leans, Louisiana. Association for Computational Lin-
677 guistics.
- 678 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le,
679 Ed Chi, Sharan Narang, Aakanksha Chowdhery, and
680 Denny Zhou. 2022. Self-consistency improves chain
681 of thought reasoning in language models. *arXiv*
682 *preprint arXiv:2203.11171*.
- 683 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten
684 Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,
685 et al. 2022. Chain-of-thought prompting elicits rea-
686 soning in large language models. *Advances in neural*
687 *information processing systems*, 35:24824–24837.
- 688 Yao Xu, Shizhu He, Jiabei Chen, Zihao Wang, Yangqiu
689 Song, Hanghang Tong, Guang Liu, Kang Liu,
690 and Jun Zhao. 2024. Generate-on-graph: Treat
691 llm as both agent and kg in incomplete knowl-
692 edge graph question answering. *arXiv preprint*
693 *arXiv:2404.14741*.
- 694 Weiguo Zheng, Jeffrey Xu Yu, Lei Zou, and Hong
695 Cheng. 2018. Question answering over knowledge
696 graphs: question understanding via template decom-
697 position. *Proceedings of the VLDB Endowment*,
698 11(11):1373–1386.

A Appendix

A.1 Experiment Detail

A.1.1 Experiment Dataset Detail

CWQ (Talmor and Berant, 2018) and WebQSP (Berant et al., 2013) were designed for KBQA task, and we employ them for KBQA tasks, and all two datasets have no individual people or offensive content.

A.1.2 Experiment Implementation Detail

We use four LLMs as backbones in our experiments: GPT-3.5-turbo (gpt-3.5-turbo-0125), GPT-4o (gpt-4o-2024-11-20), LLaMA3.3-70B-Instruct¹, and DeepSeek-V3. GPT-3.5-turbo and GPT-4o are accessed via the OpenAI API², while DeepSeek-V3 is accessed through the DeepSeek API³. All experiments and datasets are conducted using Freebase (Bollacker et al., 2008) as the underlying knowledge base.

Furthermore, the training of models was carried out on four A100 GPUs for the Llama3.3-Instruct inference task. Specifically, for the PDRR model, the running durations were roughly 120 hours for the CWQ dataset, 30 hours for WebQSP. Our experiments were facilitated by leveraging PyTorch, Huggingface, and Numpy as essential tools. Furthermore, We use ChatGPT in our paper writing and programming. Finally, we obtain results by using a single run for all results.

A.1.3 question type details of CWQ dataset

Table 1 provides examples of the four question types in the CWQ dataset, along with their corresponding canonical reasoning logic and the expected reasoning type for each.

From the reasoning logic, we observe that Composition questions align with standard chain reasoning, which can also be effectively handled by chain-based methods. In contrast, Conjunction questions are expected to be addressed using parallel reasoning.

The main challenge lies in classifying Comparative and Superlative questions. Logically, they follow a chain reasoning pattern, as identifying bridge entity is a prerequisite for the next reasoning step, and chain reasoning typically retains only one such entity. However, given the often large

number of bridge entities involved, parallel reasoning becomes necessary to reduce computational cost while preserving relevant information—an issue we previously discussed in the Comparison of Chain and Parallel Reasoning section 4.2.3. Considering that our approach leverages LLMs instead of relying solely on KBs as in traditional SP-based methods, we opt for parallel reasoning to retain as many bridge entities as possible. The subsequent reasoning is then delegated to the LLM, which can effectively handle this complexity.

A.2 Result Detail

A.2.1 Detail case example of Chain and Parallel reasoning type for Conjunction question

Table 4 presents two representative examples to illustrate under what conditions conjunction questions can be successfully answered using either chain or parallel reasoning.

In the first example, the correct reasoning logic is to find the intersection of two sets: (1) regions where *William Morris* serves as the religious head, and (2) regions that are part of *the United Kingdom*. This clearly aligns with a parallel reasoning structure. However, when using chain reasoning, even though the second set contains multiple bridge entities, the first set includes only a single bridge entity—*Wales*. As shown in the reasoning triples in the table, chain reasoning can still correctly answer the question.

In the second example, although the second set of bridge entities—countries in *Eastern Europe*—includes many candidates, the first set, referring to the countries that appointed *Mikheil Saakashvili* to a governmental position, contains only one entity: *Georgia*, which is also the correct answer. Thus, chain reasoning can be used to answer the question as well.

These two examples demonstrate that for conjunction-type questions, if one reasoning chain (independent of the others) contains only a single bridge entity, the question can be answered using either parallel or chain reasoning. This supports the explanation of experimental results discussed in Section 5.2.3.

A.2.2 Experiment explanation and case study on Comparative and Superlative type question

First, as shown by the main accuracy results for each question type in Table 2, both baseline meth-

¹<https://www.llama.com>.

²<https://openai.com>.

³<https://www.deepseek.com>

| | | |
|------------------------|--|--|
| Question | William Morris is religions head in which region that is part of the United Kingdom? (question type: Conjunction) | |
| Reasoning Logic | Find the intersection of two sets:(1) regions where William Morris serves as the religious head, and (2) regions that are part of the United Kingdom. | |
| Answer | Wales | |
| Reasoning type | Reasoning Triples | Result |
| Chain | $\mathcal{T}_1^{qc,R}$: [UnName Entity, religion.religious-organization-leadership.leader, William Morris], [UnName Entity, religion.religious-organization-leadership.jurisdiction, Wales], $\mathcal{T}_2^{qc,R}$: [Wales , location.location.containedby, United Kingdom] | Step 1: Identify the region where William Morris is a religious head. According to the retrieved reasoning chains, William Morris is associated with an unnamed religious entity that has jurisdiction over Wales . Step 2: Confirm that Wales is part of the United Kingdom. The retrieved reasoning chains confirm that Wales is contained within the United Kingdom. |
| Parallel | $\mathcal{T}_1^{qp,R}$: [William Morris, religion.religious-leader.religious-leadership, UnName Entity], [Wales , religion.religious-leadership-jurisdiction.leader, UnName Entity] $\mathcal{T}_2^{qp,R}$: [Wales , location.location.containedby, United Kingdom], [England, location.location.containedby, United Kingdom], [Northern Ireland, location.location.containedby, United Kingdom], [Scotland, location.location.containedby, United Kingdom] | Step 1: Identify the region where William Morris is a religious head. based on the retrieved triples, Wales is the region linked to religious leadership. Step 2: Identify regions that are part of the United Kingdom. Based on the retrieved triples, the United Kingdom consists of England, Scotland, Wales , and Northern Ireland. Step 3: Combine the information. Since Wales is part of the United Kingdom and is associated with religious leadership, it is the region where William Morris is likely the religious head. |
| Question | What location that appointed Mikheil Saakashvili to governmental position is a country in Eastern Europe? | |
| Reasoning Logic | Find the intersection the following two sets: (1) the countries that appointed Mikheil Saakashvili to governmental position, (2) the countries in Eastern Europe | |
| Answer | Georgia | |
| Reasoning type | Reasoning Triples | Result |
| Chain | Reasoning step 1: [UnName Entity, government.government-position-held.appointed-by, Mikheil Saakashvili], [UnName Entity, government.government-position-held.jurisdiction-of-office, Georgia] Reasoning step 2: [Georgia , location.location.containedby, Eastern Europe] | Step 1: Identify the location that appointed Mikheil Saakashvili to a governmental position. According to the retrieved reasoning chain, the jurisdiction of the office held by Mikheil Saakashvili is Georgia . Step 2: Determine if Georgia is a country in Eastern Europe. According to the retrieved reasoning chain, Georgia is contained within Eastern Europe. |
| Parallel | Triple group 1: [UnName Entity, government.government-position-held.appointed-by, Mikheil Saakashvili], [UnName Entity, government.government-position-held.jurisdiction-of-office, Georgia] Triple group 2: [Eastern Europe, location.location.partially-contains, Atyrau Region], [Eastern Europe, location.location.partially-contains, Georgia], ...(other triples) | Step 1: Identify the location that appointed Mikheil Saakashvili to a governmental position. According to the retrieved triples, Mikheil Saakashvili was appointed by an unnamed entity, and the jurisdiction of office is Georgia . Step 2: Determine the region that is located in Eastern Europe. According to the retrieved triples, Georgia , the Atyrau Region in Kazakhstan, and many other regions are geographically located in Eastern Europe. Step 3: Find the intersection of the two sets. The final answer is Georgia . |

Table 4: Case study of Conjunction-type question that can be correctly answered using either chain or parallel reasoning. "Reasoning Logic" in the table denotes standard reasoning logic for the question.

ods and our PDRR perform significantly worse on comparative and superlative questions compared to composition and conjunction questions. This is due to the higher structural complexity of the former two types. As indicated by the canonical reasoning logic in Table 1, comparative and superlative questions require not only the sequential reasoning process of chain reasoning, but also the ability to handle a larger number of bridge entities, as provided by parallel reasoning.

Nevertheless, we favor using parallel reasoning for these two question types, as it enables maximal retention of the retrieved bridge entities, which are essential for subsequent reasoning steps. In contrast, chain reasoning reduces computational

cost by retaining only one bridge entity at each step, often at the expense of losing key information. This hypothesis is supported by the results in Figure 4, which show that parallel reasoning outperforms chain reasoning on both comparative and superlative questions.

To further validate this, we provide case studies in Table 6. For the comparative question, parallel reasoning successfully retrieves and retains all relevant bridge entities (Italy, Lazio, Province of Roma), allowing for accurate comparison and judgment in the next step. However, chain reasoning retains only one entity (Italy), losing the other two critical bridge entities. A similar issue arises in the superlative question, where chain reasoning

| Methods | All | Compo | Conju | Compa | Super |
|----------|------|-------|-------|-------|-------|
| Sentence | 63.4 | 64.1 | 68.8 | 30.0 | 41.0 |
| Triple | 66.8 | 67.3 | 73.2 | 34.0 | 35.9 |

Table 5: Hits@1 accuracy on the CWQ dataset when generating answers using the LLM by either directly inputting the reasoning triples or first converting them into natural language sentences. The evaluation is conducted on the first 1000 questions.

discards key information, while parallel reasoning preserves the complete context.

In summary, for comparative and superlative question types, parallel reasoning is the preferred approach when choosing between chain and parallel reasoning.

A.2.3 Ablation study

We aim to investigate whether using triples or natural language sentences as input to the LLM in the final question answering stage yields better performance. As shown in Table 5, directly inputting reasoning triples into the LLM yields higher Hits@1 accuracy than converting them into natural language. Thus, PDRR adopts triples directly without transformation.

A.2.4 Case Study on Composition and Conjunction type question

We illustrate the answering process with concrete examples by analyzing the final retrieved reasoning triples. ToG, PDRR (chain), and PDRR (parallel) are applied to composition and conjunction questions for comparison.

Composition type question From Table 7, we observe that PPDR(C) retrieves triples that strictly follow the reasoning logic and successfully derives the final answer. In contrast, when using PPDR(P) with parallel reasoning, only the first step is completed, and the second step leads to logical confusion. With ToG, although the first step correctly identifies the artist as Taylor Swift, the lack of a concrete planning process in the second step results in selecting relations or entities that merely resemble the overall question. As a consequence, all relevant triples are pruned, making it impossible to answer the question correctly.

Conjunction type question From Table 8, we observe that PPDR(P) retrieves triples that strictly follow the reasoning logic and successfully leads to the correct answer. In contrast, PPDR with chain reasoning falls into logical confusion and fails to

produce the correct answer, even after multiple reasoning steps. Similarly, ToG also fails to follow the parallel reasoning logic and cannot arrive at the correct answer.

A.3 Prompt

Table 9 and 10 present prompts for question type prediction, question decomposition, relation pruning, triple pruning, best chain selection, chain question answering, parallel question answering, IO, CoT, and PDR (for both chain and parallel question answering).

A.4 SPARQL code

Table 11 presents four core sparql functions: Entity Match with Freebase, Head Relation Search, Tail Relation Search, Head Entity Search, and Tail Entity Search.

| | | |
|------------------------|--|--|
| Question | Where with a population once of less than 5732212 is Rome, Italy located? | |
| Type | Comparative | |
| Reasoning Logic | First, identify all regions that contains Roma, and then select those with a population once less than 5732212. | |
| Answer | Province of Rome | |
| Reasoning type | Reasoning Triples | Result |
| Chain | Reasoning step 1: [Rome, location.location.containedby, Italy] Reasoning step 2: [Italy, location.location.containedby, Europe] | Step 1: Identify where Rome, Italy, is located. According to the retrieved reasoning chain, Rome is located in Italy. Additionally, Rome is part of the Lazio region, which is located on the Italian Peninsula and in Southern Europe. Step 2: Determine if the location had a population of less than 5,732,212. Based on the reasoning chain, Italy, Lazio, and the Italian Peninsula have populations that could meet this criterion at certain points in history. However, the question does not specify a time frame , so the most relevant answer is the broader location of Italy. {Italy} |
| Parallel | Triple group 1: [Rome, location.location.containedby, Italy], [Rome, location.location.containedby, Lazio], [Rome, location.location.containedby, Province of Rome] | Following the question decomposition triples: Step 1: Identify where Rome, Italy is located. According to the retrieved triples, Rome is located in Italy, Lazio, and the Province of Rome. Step 2: Identify which of these locations had a population of less than 5,732,212. Since the retrieved triples do not provide population data, I will rely on my own knowledge. The population of Italy is significantly higher than 5,732,212, so it does not qualify. Lazio, the region containing Rome, has a population higher than 5,732,212 as well. However, the Province of Rome (now part of the Metropolitan City of Rome) had a population of less than 5,732,212 at one point. Step 3: Based on this reasoning, the location with a population once less than 5,732,212 where Rome is located is the Province of Rome. {Province of Rome} |
| Question | Which college attended by Tennessee Williams has the largest population of postgraduates? | |
| Type | Superlative | |
| Reasoning Logic | First, identify all the colleges that Tennessee Williams attended, and then find the one that has the largest number of postgraduates. | |
| Answer | University of Iowa | |
| Reasoning type | Reasoning Triples | Result |
| Chain | Reasoning Chain 1: [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, University of Missouri], [University of Missouri, education.educational_institution.campuses, University of Missouri] Reasoning Chain 2: [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, Washington University in St. Louis], [Washington University in St. Louis, education.educational_institution.campuses, Washington University in St. Louis] | Step 1: Identify the colleges attended by Tennessee Williams. According to the retrieved reasoning chain, Tennessee Williams attended Washington University in St. Louis and the University of Missouri. Step 2: Determine the postgraduate population of each college. The reasoning chain does not provide explicit information about the postgraduate populations of these institutions. Using my own knowledge: - Washington University in St. Louis is known for having a significant postgraduate population due to its strong graduate programs. - The University of Missouri also has a notable postgraduate population, but it is generally smaller compared to Washington University in St. Louis. Step 3: Compare the postgraduate populations of the two colleges. Based on my knowledge, Washington University in St. Louis has a larger postgraduate population than the University of Missouri. {Washington University in St. Louis} |
| Parallel | Triple Groups: [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, Washington University in St. Louis] [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, University of Missouri] [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, University of Iowa] [UnName_Entity, education.education.student, Tennessee Williams], [UnName_Entity, education.education.institution, The New School] | Step 1: Identify the colleges attended by Tennessee Williams. According to the retrieved triples, Tennessee Williams attended Washington University in St. Louis, University of Missouri, University of Iowa, and The New School. Step 2: Determine the population of postgraduates for each college. Since the retrieved triples do not provide this information, I will rely on my own knowledge: - Washington University in St. Louis has approximately 7,000 postgraduate students. - University of Missouri has approximately 7,000 postgraduate students. - University of Iowa has approximately 8,000 postgraduate students. - The New School has approximately 3,000 postgraduate students. Step 3: Compare the postgraduate populations of these colleges. The University of Iowa has the largest population of postgraduates. {University of Iowa} |

Table 6: Case study of comparative and superlative type questions. Red-highlighted text indicates incorrect parts.

| | | |
|------------------------|--|--|
| Question | What movies did the artist that had the concert tour called the Taylor Swift Fears 2009 Tour play in? | |
| Reasoning Logic | First, identify the artist who had the concert tour called the Taylor Swift Fears 2009 Tour, and then determine which movies this artist appeared in. | |
| Answer | The Lorax | |
| Method | Retrieved Triples | Result |
| ToG | ["Taylor Swift Fearless 2009 Tour", "music.concert-tour.artist", " Taylor Swift "] <i>Prune all other relations during prune phase.</i> | Based on the given knowledge triplet, the artist associated with the Taylor Swift Fearless 2009 Tour is Taylor Swift. However, the triplet does not provide information about the movies Taylor Swift has acted in. |
| PDRR(P) | $\mathcal{T}_1^{q_c, R}$: [Taylor Swift Fearless 2009 Tour, music.concert-tour.artist, Taylor Swift] $\mathcal{T}_2^{q_p, R}$: [Movies 1, music.recording.artist, Honey] | Step 1: Identify the artist associated with the Taylor Swift Fearless 2009 Tour. According to the retrieved triples, the artist is Taylor Swift. Step 2: Identify the movies in which Taylor Swift acted. The retrieved triples do not provide this information. |
| PDRR(C) | $\mathcal{T}_1^{q_c, R}$: [Taylor Swift Fearless 2009 Tour, music.concert-tour.artist, Taylor Swift] $\mathcal{T}_2^{q_c, R}$: [UnName Entity, film.performance.actor, Taylor Swift], [UnName Entity, film.performance.film, The Lorax] | Step 1: Identify the artist who had the concert tour called the Taylor Swift Fearless 2009 Tour. According to the retrieved reasoning chain, the artist is Taylor Swift . Step 2: Determine the movies Taylor Swift acted in. Based on the retrieved reasoning chain, Taylor Swift acted in " The Lorax " |

Table 7: Case study on handling a Composition-type question using ToG, PPDR with chain reasoning (denoted as PPDR(C) in the table), and PPDR with parallel reasoning (denoted as PPDR(P)). Red-highlighted text indicates incorrect parts.

| | | |
|------------------------|--|--|
| Question | What is there to see in Mountain Time Zone near the Grand Canyon? | |
| Reasoning Logic | Find the intersection of two sets: (1) regions that belong to the Mountain Time Zone, and (2) regions near the Grand Canyon. | |
| Answer | Phoenix | |
| Method | Retrieved Triples | Result |
| ToG | ["Grand Canyon", "location.location.time-zones", "Mountain Time Zone"], ["Grand Canyon", "travel.tourist-attraction.near-travel-destination", "Grand Canyon National Park"], ["Grand Canyon", "travel.tourist-attraction.near-travel-destination", "Lake Powell"] | Based on the given knowledge triplets, the Grand Canyon is in the Mountain Time Zone, and nearby travel destinations include Grand Canyon National Park and Lake Powell. Therefore, the answer to the question is Grand Canyon National Park and Lake Powell |
| PDRR(C) | $\mathcal{T}_1^{q_c, R}$: [Grand Canyon, location.location.nearby-airports, Grand Canyon National Park Airport] $\mathcal{T}_2^{q_c, R}$: [Grand Canyon National Park Airport, location.location.containedby, Coconino County] $\mathcal{T}_3^{q_c, R}$: [Coconino County, location.location.contains, Grand Canyon National Park] | Step 1: Identify locations near the Grand Canyon that are in the Mountain Time Zone. According to the retrieved reasoning chain, Grand Canyon National Park Airport is near the Grand Canyon. Step 2: Determine attractions in these locations. Based on my own knowledge, Grand Canyon National Park contains attractions such as the Tonto Trail. |
| PDRR(P) | $\mathcal{T}_1^{q_p, R}$: [Colorado, location.location.time-zones, Mountain Time Zone], [Colorado Springs, location.location.time-zones, Mountain Time Zone], [Phoenix , location.location.time-zones, Mountain Time Zone], ... (<i>more than 50 triples</i>) $\mathcal{T}_2^{q_p, R}$: [Grand Canyon, travel.tourist-attraction.near-travel-destination, Phoenix], [Grand Canyon, travel.tourist-attraction.near-travel-destination, Grand Canyon National Park], [Grand Canyon, travel.tourist-attraction.near-travel-destination, Lake Powell] | Step 1: Identify attractions located in the Mountain Time Zone. According to the retrieved triples, there are numerous locations in the Mountain Time Zone, including Phoenix , Flagstaff, Tucson, Salt Lake City, Denver, and others. Step 2: Identify attractions near the Grand Canyon. According to the retrieved triples, attractions near the Grand Canyon include Phoenix , Grand Canyon National Park, and Lake Powell. Step 3: Find the intersection of the two sets. Phoenix is both in the Mountain Time Zone and near the Grand Canyon. |

Table 8: Case study on handling a Conjunction-type question using ToG, PPDR with chain reasoning (denoted as PPDR(C) in the table), and PPDR with parallel reasoning (denoted as PPDR(P)). Red-highlighted text indicates incorrect parts.

| Function | Prompt | few-shot example |
|--------------------------|---|---|
| Question type Prediction | Please analyze the following question and determine its type. Question Type: 1. Chain Structure 2. Parallel Structure Question: {question} Output the question type with "{ 'question type' }", and provide explanation. Do NOT format into markdown or use headers. | Question: {Who is the coach of the team owned by Steve Bisciotti?} Answer: The type of this question is {Chain Structure}, the bridge entity is "team". We should first find the team owned by Steve Bisciotti. And then find the coach of the team. |
| Question Decomposition | Please first determine the reasoning process of the question. Then decompose the question into triples following the reasoning process. Each triple should contain concise head entity, relation, and tail entity. The entity with "#number" is what we need to find. Question: {question} Question Type: {question type} | Question: {Who is the coach of the team owned by Steve Bisciotti?} Answer: Given the question type is chain structure, the sequence of the triples is important. The bridge entity is "team". We should first find the team owned by Steve Bisciotti. And then find the coach of the team. The output triples are: { "head": "Steve Bisciotti", "relation": "owns", "tail": "team#1" }, { "head": "team#1", "relation": "is coached by", "tail": "coach#1" } |
| Relation Prune | Please retrieve relations that relative to the triple and rate their relative on a scale from 0 to 1 (the sum of the scores of relations is 1). Do NOT format into markdown or use headers. triple: {triple} Relations: {relations_text} | Triple: {Van Andel Institute, founded in part by, American businessman#1} Relations: {1. affiliation 2. country 3. donations 4. educated_at 5. employer 6. headquarters_location 7. legal_form 8. located_in_the_administrative_territorial_entity 9. total_revenue} Answer: 1. {affiliation (Score: 0.4)}: This relation is relevant because it can provide information about the individuals or organizations associated with the Van Andel Institute, including the American businessman who co-founded the Amway Corporation. 2. {donations (Score: 0.3)}: This relation is relevant because it can provide information about the financial contributions made to the Van Andel Institute, which may include donations from the American businessman in question. 3. {educated_at (Score: 0.3)}: This relation is relevant because it can provide information about the educational background of the American businessman, which may have influenced his involvement in founding the Van Andel Institute. |
| Triple Prune | Please identify the triples that are relevant to the given filter-triple and rate their relevance on a scale from 0 to 1 (the sum of the scores of triples is 1). Do NOT include irrelevant triples. Do NOT format into markdown or use headers. You should choose at least 1 triple from the triples. Filter Triple: {filter_triple} Triples: {triples_text} | Filter Triple: {Rift Valley Province, is located in, nation#1} Triples: {1. Rift Valley Province, is located in, Kenya 2. Kenya, location.country.currency_used, Kenyan shilling 3. San Antonio Spurs, home venue, AT&T Center 4. Rift Valley Province, is located in, UnName_Entity 5. UnName_Entity, education.education.institution, Castlemont High School 6. Rift Valley Province, location.contains, Baringo County 7. Rift Valley Province, location.contained_by, Kenya} Answer: 1. {Rift Valley Province, is located in, Kenya. (Score: 0.5)}: This triple provides significant information about Kenya's location, which relates to the filter-triple. 2. {Rift Valley Province, location.contained_by, Kenya. (Score: 0.4)}: This triple provides significant information about Kenya's location, which relates to the filter-triple. 3. {Rift Valley Province, location.contains, Baringo County. (Score: 0.1)}: This triple provides information cannot show us the location of it, so it is irrelevant. |
| Best Chain Selection | Please select the best reasoning chain to answer the question from the following chains: Reasoning Chains: {reasoning chain str} Question: {question} | Reasoning Chains: chain 1: {Country Nation World Tour, music.concert-tour.artist, Brad Paisley}, {Brad Paisley, owns, Nashville Predators} chain 2: {Country Nation World Tour, music.concert-tour.artist, Brad Paisley}, {Brad Paisley, attended, Belmont University} chain 3: {Country Nation World Tour, is hold by, Steve Bisciotti}, {Steve Bisciotti, attended, University of Alabama at Birmingham} Question: Where did the "Country Nation World Tour" concert artist go to college? Answer: The best reasoning chain is chain 2: {Country Nation World Tour, music.concert-tour.artist, Brad Paisley}, {Brad Paisley, attended, Belmont University}. It successfully finds the bridge entity "artist", which refers to Brad Paisley, the artist of the Country Nation World Tour, and then finds the college he attended: Belmont University. |

Table 9: Prompt List 1. The prompt list includes prompts for question type prediction, question decomposition, relation pruning, triple pruning, and best chain selection.

| Function | Prompt | few-shot example |
|-----------------------------------|--|---|
| Chain Question Answering | <p>Given a question and the associated information, you are asked to answer the question using the retrieved reasoning chain and your own knowledge. Please think setep by step and follow the Question Decomposition Triples carefully. Do NOT output answer without reasoning steps. Do NOT format into markdown or use headers. At the end, output the final answer in this format: "{answer}"</p> <p>Question: {question}</p> <p>Question Decomposition Triples: {question decomposition triples str}</p> <p>Retrieved Reasoning Chain: {best reasoning chain str}</p> | <p>Question: {Rift Valley Province is located in a nation that uses which form of currency?}</p> <p>Question Decomposition Triples: {Rift Valley Province, is located in, nation#1}, {nation#1, uses currency, currency#1}</p> <p>Retrieved Reasoning Chain: {Rift Valley Province, location.administrative_division.country, Kenya}, {Kenya, location.country.currency_used, Kenyan shilling}</p> <p>Answer: Following the question decomposition triples: Step 1: Identify the nation in which Rift Valley Province is located. According to the retrieved reasoning chain, Rift Valley Province is located in Kenya. Step 2: Determine the currency used by Kenya. The retrieved reasoning chain indicates that Kenya uses the Kenyan shilling. {Kenyan shilling}</p> |
| Parallel Question Answering | <p>Given a question and the associated information, you are asked to answer the question with these Retrieved Triples and your own knowledge. Please think setep by step and follow the Question Decomposition Triples carefully. Do NOT output answer without reasoning steps. Do NOT format into markdown or use headers. At the end, output the final answer in this format: "{answer}"</p> <p>Question: question</p> <p>Question Decomposition Triples: question decomposition triples str</p> <p>Retrieved Triples: formatted triples</p> | <p>Question: {What country bordering France contains an airport that serves Nijmegen?}</p> <p>Question Decomposition Triples: {country#1, borders, France}, {country#1, contains an airport that serves, Nijmegen}</p> <p>Retrieved Triples: {{{Belgium, borders, France}, {Germany, borders, France}, {Italy, borders, France}, {Switzerland, borders, France}}, {{Germany, contains an airport that serves, Nijmegen}, {Netherlands, contains an airport that serves, Nijmegen}}}</p> <p>Answer: Following the question decomposition Triples: Step 1: Identify the country that borders France. According to the retrieved triples, the country are Belgium, Germany, Italy, and Switzerland. Step 2: Identify the country that contains an airport that serves Nijmegen. According to the retrieved triples, the country is Netherlands. Step 3: Find the intersection of the two sets, which is Germany. {Germany}</p> |
| IO | <p>Please answer the question, and output the answer in this format: "{answer}". Do NOT format into markdown or use headers</p> <p>Question: {question}</p> | <p>Question: {What state is home to the university that is represented in sports by George Washington Colonials men's basketball?}</p> <p>Answer: {Washington, D.C.}</p> |
| COT | <p>Please think setep by step and answer the question. Output the answer in this format: "{answer}". Do NOT format into markdown or use headers</p> <p>Question: {question}</p> | <p>Question: {What state is home to the university that is represented in sports by George Washington Colonials men's basketball?}</p> <p>Answer: First, the education institution has a sports team named George Washington Colonials men's basketball in is George Washington University , Second, George Washington University is in Washington D.C. The answer is {Washington, D.C.}.</p> |
| PDR (Chain question answering) | <p>Answer the question using the provided decomposition triples and your own knowledge. Think step by step, and strictly follow the triples. Do not skip reasoning, use markdown or headers. At the end, output the final answer as: "{answer}"</p> <p>Question: {question}</p> <p>Question Decomposition Triples: {question decomposition triples str}</p> | <p>Question: {Where did the "Country Nation World Tour" concert artist go to college?}</p> <p>Question Decomposition Triples: {{Country Nation World Tour, is concert tour by, artist#1}, {artist#1, attended, college#1}}</p> <p>Answer: Following the question decomposition triples: Step 1: Identify the artist of the "Country Nation World Tour" concert. Based on my knowledge, the artist is Brad Paisley. Step 2: Determine the college that Brad Paisley attended. Based on my knowledge, he attended Belmont University. {Belmont University}</p> |
| PDR (Parallel question answering) | <p>Answer the question using the provided decomposition triples and your own knowledge. Think step by step, and strictly follow the triples. Do not skip reasoning, use markdown or headers. At the end, output the final answer as: "{answer}"</p> <p>Question: {question}</p> <p>Question Decomposition Triples: {question decomposition triples str}</p> | <p>Question: {What country bordering France contains an airport that serves Nijmegen?}</p> <p>Question Decomposition Triples: {{country#1, borders, France}, {country#1, contains an airport that serves, Nijmegen}}</p> <p>Answer: Following the question decomposition Triples: Step 1: Identify the country that borders France. Based on my own knowledge, the country are Belgium, Germany, Italy, and Switzerland. Step 2: Identify the country that contains an airport that serves Nijmegen. Based on my knowledge, the country which contains an airport that serves Nijmegen are Germany and Netherlands. Step 3: Find the intersection of the two sets, which is Germany. {Germany}</p> |

Table 10: Prompt List 2. The prompt list includes prompts for chain question answering, parallel question answering, IO, COT, PDR(chain question answering), PDR(parallel question answering).

| Function | Sparql code |
|--|---|
| Entity Match with Freebase: Given the entity string mentioned in text, use fuzzy matching to directly search for the corresponding Freebase entity and its entity ID. | <pre> PREFIX ns: <http://rdf.freebase.com/ns/> SELECT DISTINCT ?entity ?label WHERE { ?entity ns:type.object.name ?label . FILTER(LANG(?label) = "en") . FILTER(bif:contains(?label, "entity_string")) . } </pre> |
| Head Relation Search: Retrieve all relations where the head entity is the subject. | <pre> PREFIX ns: <http://rdf.freebase.com/ns/> SELECT ?relation WHERE { ns:head_entity_id ?relation ?x . } </pre> |
| Tail Relation Search: Retrieve all relations where the tail entity is the object. | <pre> PREFIX ns: <http://rdf.freebase.com/ns/> SELECT ?relation WHERE { ?x ?relation ns:tail_entity_id . } </pre> |
| Head Entity Search: Given a tail entity and the relations where it serves as the object, retrieve all corresponding head entities. | <pre> PREFIX ns: <http://rdf.freebase.com/ns/> SELECT ?headEntity WHERE { ?headEntity ns:relation ns:tail_entity_id . } </pre> |
| Tail Entity Search: Given a head entity and the relations where it serves as the subject, retrieve all corresponding tail entities. | <pre> PREFIX ns: <http://rdf.freebase.com/ns/> SELECT ?tailEntity WHERE { ns:head_entity_id ns:relation ?tailEntity . } </pre> |

Table 11: SPARQL Code for PDRR. This table presents four core functions: Entity Match with Freebase, Head Relation Search, Tail Relation Search, Head Entity Search, and Tail Entity Search. Inputs to each function are highlighted in red, and the outputs represent the desired search results.