PolyG: Effective and Efficient GraphRAG with Adaptive Graph Traversal

Renjie Liu [1]   Haitian Jiang [2]   Xiao Yan [3]   Bo Tang [1]   Jinyang Li [2]

Abstract

GraphRAG enhances large language models (LLMs) to generate quality answers for user questions by retrieving related facts from external knowledge graphs. Existing GraphRAG methods adopt a *fixed* graph traversal strategy for fact retrieval but we observe that user questions come in different types and require different graph traversal strategies. As such, existing GraphRAG methods are limited in *effectiveness* (i.e., quality of the generated answers) and/or *efficiency* (i.e., response time or the number of used tokens). In this paper, we propose to classify the questions according to a *complete* four-class taxonomy and *adaptively* select the appropriate graph traversal strategy for each type of questions. Our system PolyG is essentially a query planner for GraphRAG and can handle diverse questions with an unified interface and execution engine. Compared with SOTA GraphRAG methods, PolyG achieves an overall win rate of 75% on generation quality and a speedup up to $4\times$ on response time.

1. Introduction

Large Language Models (LLMs) (Touvron et al., 2023; OpenAI et al., 2024; Brown et al., 2020) have achieved remarkable success for various natural language processing (NLP) tasks (Bang et al., 2023; Dam et al., 2024; Wang et al., 2023). However, they are susceptible to hallucinations when generating answers for questions that require information beyond their knowledge (Sahoo et al., 2024; Zhang et al., 2023). To tackle this problem, graph-based retrieval-augmented generation (GraphRAG) (Edge et al., 2024) has emerged as an effective approach, which grounds LLM's responses using external knowledge graphs (Wang et al., 2018; Weber et al., 2019; Garton et al., 1997). In particular, to answer a

[1]Southern University of Science and Technology [2]New York University [3]Centre for Perceptual and Interactive Intelligence. Email: Renjie Liu <liurj2023@mail.sustech.edu.cn>, Haitian Jiang <haitian.jiang@nyu.edu>, Xiao Yan <yanxiaosunny@gmail.com>, Bo Tang <tangb3@sustech.edu.cn>, Jinyang Li <jinyang@cs.nyu.edu>.

Table 1. Win rates (i.e., the chance for a method to produce the best answer among all the methods) for 4 types of graph questions.

Method	Type 1	Type 2	Type 3	Type 4
MS_GraphRAG, LightRAG	80.83%	12.92%	12.50%	5.00%
RoG	N/A	77.08%	N/A	N/A
Fast-graphrag	8.33%	11.67%	1.67%	7.92%
Graph-CoT, ToG	11.25%	6.25%	1.67%	2.92%
PolyG (ours)	**80.83%**	**77.08%**	**85.00%**	**71.67%**

user question, GraphRAG first retrieves the relevant entities and relations from the knowledge graph and then ultilizes this information along with the user question to prompt the LLM for answer generation.

Limitations of existing solutions. Several GraphRAG methods have been proposed, and they mainly differ in how they perform graph traversal based on entities extracted from the question. For example, Microsoft GraphRAG (Edge et al., 2024) (hereafter referred to as MS_GraphRAG) and LightRAG (Guo et al., 2024) adopt one-hop BFS to retrieve information from the neighborhood of the extracted entities; RoG (Luo et al., 2024) deduces concrete and faithful reasoning paths starting from the question's entities and use them for traversal; Fast-GraphRAG (fas, 2024) uses the personalized PageRank (PPR) score (Page et al., 1999) to select the top paths; while ToG (Sun et al., 2024) and Graph-CoT (Jin et al., 2024) invoke the LLM at each traversal step to select the next entity to visit in the knowledge graph.

Since existing GraphRAG methods adopt *a single fixed graph traversal strategy*, they suffer from two limitations.

- *Limited effectiveness*: Most methods work well for a specific type of questions but fall short for other types. For instance, using one-hop BFS, MS_GraphRAG and LightRAG suit questions that ask about general aspects of an entity but cannot handle questions involving entities that are more than one-hop apart. Relying on the question to deduce the reasoning paths, RoG works well when the question provides explicit relation constraints but
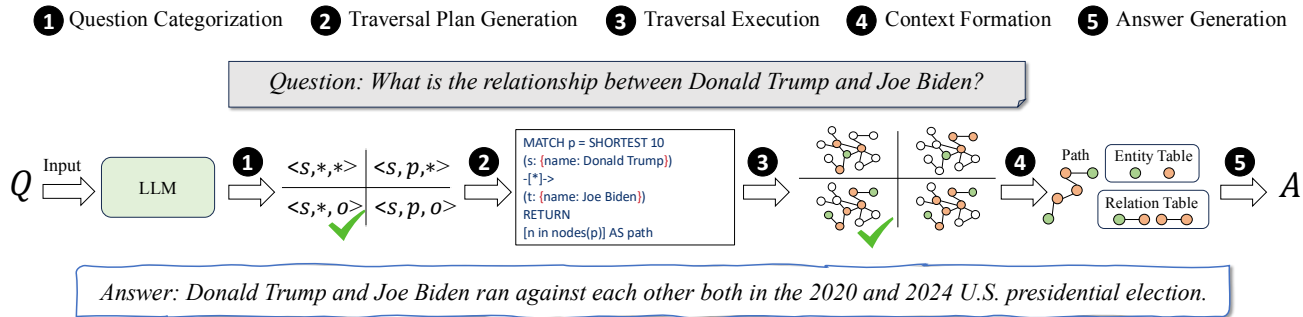
① Question Categorization ② Traversal Plan Generation ③ Traversal Execution ④ Context Formation ⑤ Answer Generation

Question: What is the relationship between Donald Trump and Joe Biden?

Figure 1. The workflow of our PolyG.

fails when these constraints are missing or vague. Fast-GraphRAG are good for questions without constraints by ranking possible paths using their PPR scores but the selected top-ranking paths may be irrelevant when the question has explicit constraints. As shown in Table 1, existing GraphRAG methods achieve high accuracy for some types of questions (e.g., MS_GraphRAG for Type 1 and RoG for Type 2) but they have low accuracy or are inapplicable for the other types of questions.

- *High cost*: By using an LLM to determine the next entity to explore at every traversal step, ToG and Graph-CoT can handle all types of questions. However, they incur significant overheads due to frequent LLM invocations, increasing the end-to-end response time by approximately 4× and the number of used tokens by 30× according to our profiling. Moreover, as shown in Table 1, they often do not achieve the best answer quality.

Our solution PolyG. To tackle the limitations of existing solutions, we propose PolyG, which achieves high accuracy for diverse graph questions while maintaining efficiency. As shown in Figure 1, PolyG categorizes the graph questions into four distinct types and adopts a suitable graph traversal strategy for each type of questions. Specifically, we abstract each graph question as a knowledge graph triple (i.e., *subject*, *predicate*, and *object*) and classify the questions *depending on the missing component* in the triple. That is, a question may inquire about specific or vague aspects of an entity, explore potential relations between two entities, or verify the existence of a relation. We show that such a question classification is *complete* and identify an appropriate graph traversal strategy for each type of questions that faithfully retrieves the related facts. Before execution, PolyG prompts an LLM to decide the question type and select the graph traversal method. To our knowledge, PolyG is the first query planner for GraphRAG, and by introducing an explicit query planning stage, PolyG also opens the room for other advanced query processing strategies, e.g., dividing a complex query into simple ones.

To systematically evaluate PolyG and the baselines, we introduce a new GraphRAG benchmark that encompasses all four types of graph questions generated from carefully crafted templates. We compare the GraphRAG methods in terms of answer quality, end-to-end latency, and token usage. Extensive experimental results show that our PolyG consistently outperforms the baselines across all four question types, achieving an average win rate of 75% for answer quality. Moreover, PolyG also reduces the response latency by up to 4× and the token consumption by up to 95%. Micro-benchmark results indicate that PolyG incurs a small overhead for query planning, requiring less than 10% of the processing time and 30% of the token usage.

To summarize, we make the following key contributions.

- We systematically study diverse graph questions and propose a four-class taxonomy for question classification.
- We design tailored graph traversal strategies for each question type, which effectively retrieve the related facts while reducing time and token consumption.
- We develop a general GraphRAG solution that automatically prompts an LLM to decide the graph traversal plan and executes with a unified interface and backend for different types of graph questions.
- We introduce a new GraphRAG benchmark that encompasses a diverse set of graph questions. We also analyze the pros and cons of different graph traversal strategies with extensive experiments.

2. GraphRAG Basics

Knowledge graph. A knowledge graph $KG = \{(s, p, o) \mid s, o \in \mathcal{E}, p \in \mathcal{R}\}$ consists of triples (facts) in the form $\langle s, p, o \rangle$, where s is the subject, p is the predicate, and o is the object. For example, a knowledge graph may contain the fact: "Donald Trump is the president of the United States," where "Donald Trump" is the subject, "is president of" is the predicate, and "United States" is the object. The set of

Table 2. The 4 types of graph questions in our taxonomy and their corresponding graph traversal strategy. In a knowledge graph triple $\langle s, p, o \rangle$, s stands for the subject, p refers to the predicates, and o is the object.

Question Type	Description	Graph Traversal Strategy
$\langle s, *, * \rangle$	Questions about an entity (the subject) with no specific relation constraints (the predicate) and target entity (the object). The task is to answer general question about the entity. Example: *"Who is Donald Trump?"*	BFS-style Neighbor Expansion
$\langle s, p, * \rangle$	Questions about an entity (the subject) with specific relation constraints (the predicate) but misses the target entity (the object). The task is to answer one specific aspect of the entity. Example: *"Which college did the current president of United States graduate from?"*	Meta-path Guided Walks
$\langle s, *, o \rangle$	Questions about any relations (the predicate) between two entities (the subject and object). The task is to provide the relations between two entities. Example: *"What is the relationship between Donald Trump and Joe Biden?"*	Top-k Shortest Paths
$\langle s, p, o \rangle$	Questions about specific relations (the predicate) between two entities (the subject and object). The task is to check the existence of a specific relationship between the two entities. Example: *"Have Donald Trump and Joe Biden ever run for the same presidential term, and if so, when is that?"*	Top-k Constrained Shortest Paths

all subjects and objects forms the entity set $\mathcal{E}$, while the set of all predicates constitutes the relation set $\mathcal{R}$.

Reasoning path. A reasoning path in a knowledge graph is a sequence of entities connected by relations. It is represented as $P = e_0 \rightarrow (r_0) \rightarrow e_1 \rightarrow (r_1) \rightarrow \cdots \rightarrow (r_n) \rightarrow e_n$, where $e_i \in \mathcal{E}$ denotes the i-th entity, and $r_i \in \mathcal{R}$ represents the i-th relation in an n-hop reasoning path. For example, the reasoning path for the fact "The current president of the United States graduated from UPenn." is:
`United States → (president) → Donald Trump → (graduate from) → UPenn`.

GraphRAG. Given a natural language question Q and a knowledge graph KG, the objective of GraphRAG is to develop a retriever ψ that extracts relevant entities, relations, and reasoning paths from the knowledge graph. The retrieved context is then provided to the LLM to generate the final answer A. This process is formalized as:

$$A = \mathcal{LLM}(Q, \psi(KG, Q)). \qquad (1)$$

Existing GraphRAG methods mainly differ in how to traverse the knowledge graph and extract the facts related to the user question. Detailed descriptions of their workflow are provided in Appendix A.

3. Question Taxonomy and Traversal Patterns

Graph questions primarily inquire about different components of triples, i.e., $\langle s, p, o \rangle$, in a knowledge graph, where s is the subject, p is the predicate, and o is the object. By masking each element of the triple as either known or unknown ($*$), a total of eight possible question types (2^3) can

be generated. For example, a question of the form $\langle s, *, * \rangle$ could be "Who is Donald Trump?" or "Tell me about Joe Biden." However, some combinations do not yield meaningful questions, while others are redundant. Specifically, the type $\langle *, *, * \rangle$ does not correspond to any valid query. Additionally, the types $\langle *, p, o \rangle$ and $\langle *, *, o \rangle$ are equivalent to $\langle s, p, * \rangle$ and $\langle s, *, * \rangle$, respectively, due to the existence of inverse relations in knowledge graphs. The type $\langle *, p, * \rangle$ necessitates a linear scan of all entities, making it unrelated to traversal-based operations and thus excluded from this work. Consequently, we identify four valid question types: $\langle s, *, * \rangle$, $\langle s, p, * \rangle$, $\langle s, *, o \rangle$, and $\langle s, p, o \rangle$.

In Table 2, we define each of the four question types, outline their corresponding tasks, and provide concrete examples. Additionally, we observe that each question type is best suited to a specific graph traversal method, which is also summarized in Table 2. In the following, we examine each question type in detail and explain its corresponding graph traversal pattern.

Question Type $\langle s, *, * \rangle$**.** This category of graph questions pertains to queries about entities in the knowledge graph without specific relation constraints. These questions take the form $\langle s, *, * \rangle$, where the subject s is provided, and the objective is to retrieve all associated predicates and objects. For example, questions like "Who is Donald Trump?" or "Tell me about Joe Biden." fall into this category. For such queries, BFS-style neighbor expansion (hereafter BFS), as employed by MS_GraphRAG and LightRAG, is well-suited. This approach efficiently captures all neighborhood knowledge of the subject s by exploring its one-hop neighbors and

retrieving directly connected relations.

Question Type $\langle s, p, * \rangle$**.** When a question specifies a subject s along with a specific relation constraint p, it falls into the $\langle s, p, * \rangle$ category. These questions seek information about a particular aspect of the subject entity. For example, the question "Which college did the current president of the United States graduate from?" specifically requests graduation-related information rather than general details about the president. Since the relation constraint is explicitly defined, the method proposed in RoG is well-suited for this task. RoG prompts the LLM to generate faithful reasoning paths that can be followed within the graph structure, leading from the subject s to the answer entity o. In addition to accurately retrieving the answer, this guided traversal minimizes extraneous context by focusing only on the relevant entities and relations.

Question Type $\langle s, *, o \rangle$**.** Graph questions of this type take the form $\langle s, *, o \rangle$, where both the subject s and the object o are specified, and the goal is to identify any existing relationship between them. For example, the question "What is the relationship between Donald Trump and Joe Biden?" falls into this category. Since no information is provided about how the two entities are interconnected or the number of hops between them, it is not feasible to construct faithful reasoning paths. Additionally, BFS is limited to capturing only short-distance knowledge within the separate neighborhoods of s and o. In this scenario, where no specific predicate p is given, the top-k shortest path method is a natural solution, as it identifies the shortest existing linkage between the two entities. Beyond shortest paths, more general random walk-based approaches, such as DeepWalk (Perozzi et al., 2014) and Node2Vec (Grover & Leskovec, 2016), can also effectively uncover reasoning paths between the entities.

Question Type $\langle s, p, o \rangle$**.** When a question specifies a subject s, an object o, and a specific relation constraint p, it falls into the $\langle s, p, o \rangle$ category. These questions either inquire about particular aspects of the relationship between the entities or verify whether a specific relationship exists. In this case, the top-k shortest path method may fail to retrieve the correct paths, as the shortest connections might represent unrelated relationships that do not match the query context. For example, in the question "Have Donald Trump and Joe Biden ever run for the same presidential term, and if so, when?", the correct answer is "yes", but the shortest paths might only capture their common role as U.S. presidents rather than their electoral competition. To address this, a relation filtering step can be applied to the top-k shortest paths, ensuring that only paths satisfying the specified constraint p are considered.

Exhaustiveness of the taxonomy. While we acknowledge that our four question categories do not encompass all possible graph questions—particularly due to the free-form nature of natural language and questions about global summaries (Edge et al., 2024)—we empirically validate that the majority of questions in widely used KBQA benchmarks fall within these four types. In particular, we analyze questions from these benchmarks and prompt an LLM to classify each question into one of our four categories or determine if it remains unclassified. In the WebQSP (Yih et al., 2016) dataset, 96.45% of the questions belong to the $\langle s, p, * \rangle$ type. For the remaining 3.55% of unclassified questions, we observe that they can be decomposed into multiple sub-tasks—some belonging to one of our four question types, while the others involve non-traversal operations such as attribute retrieval or sorting. For example, the question "What are the 5 biggest cities in the USA?" consists of an initial $\langle s, p, * \rangle$-type query to identify all cities in the USA, followed by a sorting operation based on population. Additional experimental results and case studies are provided in Appendix F. While edge cases undoubtedly exist, our empirical findings suggest that the four proposed question types cover the majority of graph-related queries encountered in practice. We leave the argentic workflow to handling question decomposition for more complex, compounded queries to future work.

4. The PolyG System

In this section, we first present the workflow of PolyG, detailing its process for automatically categorizing user queries and generating responses. We then discuss how PolyG constructs an executable traversal plan for each question type and organizes the retrieved information into a structured context during traversal execution.

Workflow. As shown in Figure 1, PolyG adaptively selects the appropriate graph traversal method based on the input question and executes the traversal operations through a unified interface to the backend execution engine. Specifically, PolyG follows these steps to answer a user query.

❶ *Question Categorization:* Given a graph-based question, PolyG prompts the judger LLM to classify it into one of the predefined categories by providing definitions and few-shot examples, as detailed in Table 2. The prompt used for classification is provided in Appendix E. The judger LLM outputs the question type, which is then used to generate concrete traversal operations in the next step.

❷ *Traversal Plan Generation:* PolyG extracts anchor entities and potential relation constraints based on the question type. It then fills the extracted keywords into predefined templates for $\langle s, *, * \rangle$ and $\langle s, *, o \rangle$ questions or invokes the generator LLM to construct concrete reasoning paths for $\langle s, p, * \rangle$ and $\langle s, p, o \rangle$ questions.

❸ *Traversal Execution:* With an executable traversal plan, PolyG interfaces with an underlying Neo4j (Neo4j, 2012)

Table 3. Statistics of the knowledge graphs used for evaluation

Attributes	Academia	Literature	E-commerce
Entity Types	3	4	2
Relation Types	6	7	5
Entities	2.7M	3.7M	8.1M
Relations	67M	28M	193M

graph database and submits the plan as a Cypher query. This query initiates graph traversal starting from the anchor entity (s) and, if applicable, reaching the target entity (o), depending on the question type.

❹ *Context Formation:* Once relevant graph entities and relations are retrieved, PolyG organizes the results into structured formats, such as entity and relation tables for $\langle s, *, * \rangle$ and $\langle s, p, * \rangle$ questions, and reasoning path lists for $\langle s, *, o \rangle$ and $\langle s, p, o \rangle$ questions.

❺ *Answer Generation:* Finally, the structured context is fed to the generator LLM, which synthesizes a final response grounded in the retrieved graph knowledge.

Traversal templates and output formats. For $\langle s, *, * \rangle$ questions, PolyG follows the approach used in MS_GraphRAG and LightRAG, locating the anchor entities in the graph, gathering their one-hop neighbors, and returning two tables (i.e., an entity table and a relation table) as the retrieved context. For $\langle s, p, * \rangle$ questions, PolyG provides the LLM with a detailed graph schema, including entity and relation definitions and how different entity types interconnect. It then prompts the LLM to generate reasoning paths based on the schema and the specific input question. PolyG follows the generated reasoning path and retrieves the traversal path from the anchor entity (s) to the answer (o), along with the entity and relation tables. Unlike RoG, which requires instruction-tuning on the LLM, PolyG avoids fine-tuning, offering greater ease of use. For $\langle s, *, o \rangle$ questions, PolyG extracts the two anchor entities (s and o) and executes the top-k shortest paths algorithm on the graph. The results are structured as k paths in a list format (entity→relation→...→entity). Finally, for $\langle s, p, o \rangle$ questions, PolyG extracts both anchor entities (s and o) along with the constraints for path matching (p). It then performs constrained path matching between the two entities and retrieves the top-k shortest paths that satisfy the constraints. The retrieved paths are presented as multiple path lists, illustrating examples of relationships under the specified constraints.

5. Evaluation

In this part, we conduct experiments to evaluate PolyG and compare it with existing baseline solutions. Through the experimental results, we tend to explore:

- *How well does PolyG generates responses to various types of graph questions compared to baseline methods?*
- *How efficient is PolyG when answering various types of graph questions compared to baseline methods?*
- *How many fewer tokens does PolyG consume in order to answer the question compared to baseline methods?*

5.1. Experiment Settings

Datasets. We utilize the publicly available GraphQA benchmark, GRBENCH, released in Graph-CoT (Jin et al., 2024). GRBENCH consists of 10 knowledge graphs spanning five domains: academia, e-commerce, literature, healthcare, and legal, collected from various public sources. For evaluation, we select three knowledge graphs from the academia (Leipzig, 2022), literature (Graph, 2022), and e-commerce (Benchmark, 2022) domains. The statistics of these three knowledge graphs are presented in Table 3.

Question generation. Existing GraphQA benchmarks (Yih et al., 2016; Talmor & Berant, 2018; Yang et al., 2018) primarily focus on the $\langle s, p, * \rangle$ question type, as only these questions have ground-truth answers available for evaluation. To comprehensively assess the performance of existing methods across a diverse range of graph questions, we introduce a new GraphRAG benchmark that encompasses all four question types discussed above, utilizing manually designed question templates. Specifically, we generate 80 questions for each question type by randomly selecting entities and populating the corresponding templates, resulting in a total of 320 questions per knowledge graph. The question templates used in this generation process are detailed in Appendix C.

Baselines and models. We compare PolyG against existing GraphRAG methods discussed in section 1. To ensure a fair comparison, we categorize baselines based on their adopted graph traversal methods and select one representative from each category, despite minor differences in other modules. For instance, since both MS_GraphRAG and LightRAG use BFS for one-hop neighbor expansion, we select MS_GraphRAG as the representative method. Similarly, for traversal methods that invoke LLMs at each step, we compare with Graph-CoT instead of ToG. Additionally, since RoG requires fine-tuning an LLM to generate faithful reasoning paths, we instead prompt the LLM with the graph schema to avoid instruction tuning. We also include the Top-k shortest paths method (Top-k SP) and the Top-k constrained shortest paths method (Top-k CSP) as baselines to evaluate their performance across different question types, setting $k = 10$ for both. For consistency, we use the Claude-3.5-sonnet API (Anthropic, 2024) as the generator LLM for

Table 4. Answer generation quality of our PolyG and the baselines on the *Academia* network. *SP* means Shortest Paths, and *CSP* refers to Constrained Shortest Paths. The win rates sum over 100% as we allow multiple winners. **Bold faces** denote the highest win rate.

Question Type	Criteria	MS_GraphRAG (BFS)	RoG (Guided Walk)	Fast-graphrag (PPR)	Graph-CoT (CoT)	Top-k SP	Top-k CSP	PolyG
$\langle s, *, * \rangle$	Comprehensiveness	**95.00%**	N/A	28.75%	21.25%	N/A	N/A	**95.00%**
	Diversity	**88.75%**	N/A	22.50%	15.00%	N/A	N/A	**88.75%**
	Empowerment	**85.00%**	N/A	32.50%	17.50%	N/A	N/A	**85.00%**
	Overall	**86.25%**	N/A	5.00%	8.75%	N/A	N/A	**86.25%**
$\langle s, p, * \rangle$	Comprehensiveness	31.25%	**80.00%**	31.25%	30.00%	N/A	N/A	**80.00%**
	Diversity	23.75%	**75.00%**	21.25%	7.50%	N/A	N/A	**75.00%**
	Empowerment	30.00%	**75.00%**	20.00%	7.50%	N/A	N/A	**75.00%**
	Overall	23.75%	**70.00%**	17.50%	6.25%	N/A	N/A	**70.00%**
	Precision	0.2	**0.78**	0.37	0.63	N/A	N/A	**0.78**
	Recall	0.17	**0.76**	0.29	0.44	N/A	N/A	**0.76**
	F1	0.18	**0.76**	0.3	0.48	N/A	N/A	**0.76**
$\langle s, *, o \rangle$	Comprehensiveness	36.25%	N/A	0.00%	12.50%	**93.75%**	18.75%	**93.75%**
	Diversity	10.00%	N/A	0.00%	6.25%	**92.50%**	16.25%	**92.50%**
	Empowerment	21.25%	N/A	0.00%	5.00%	**91.25%**	15.00%	**91.25%**
	Overall	2.50%	N/A	0.00%	1.25%	**88.75%**	10.00%	**88.75%**
$\langle s, p, o \rangle$	Comprehensiveness	16.25%	N/A	16.25%	21.25%	35.00%	**81.25%**	**81.25%**
	Diversity	16.25%	N/A	10.00%	6.25%	42.50%	**76.25%**	**76.25%**
	Empowerment	11.25%	N/A	11.25%	7.50%	35.00%	**83.75%**	**83.75%**
	Overall	8.75%	N/A	3.75%	3.75%	22.50%	**72.50%**	**72.50%**
Overall		30.31%	17.50%	6.56%	5.00%	27.81%	20.63%	**79.38%**

PolyG and all baselines.

Performance metric. Since many graph questions do not have a single definitive answer, we assess generation quality using win rates determined by LLM judgments. We evaluate responses from four perspectives: *Comprehensiveness*, *Diversity*, *Empowerment*, and *Overall Winner*, following the approach in (Edge et al., 2024; Guo et al., 2024). Detailed explanations of these criteria are provided in Appendix D. As each question is judged across multiple baselines, we prompt the LLM to allow multiple winners, leading to total win rates exceeding 100%. In addition to win rates, we report *Precision*, *Recall*, and *F1-score* for $\langle s, p, * \rangle$ questions, as they have ground-truth answers. For consistency, we use the Claude-3.5-sonnet API (Anthropic, 2024) as the judger LLM, the same as in question generation. Additionally, we measure end-to-end response latency and token usage to evaluate computational efficiency across methods.

5.2. Experimental Results

Generation quality. Table 4 and Table 5 present the win rates of different methods across all four question types on the academia and literature networks, respectively. As the observations on the e-commerce network are similar, we provide the corresponding results in Appendix B due to space constraints. Notably, Guided Walk is applicable only to the $\langle s, p, * \rangle$ question type, as it relies on concrete reasoning paths that cannot be inferred for other question

types where relation constraints are unspecified or where the question seeks to explore relationships between two given entities. Similarly, Top-k SP and Top-k CSP are inapplicable to both $\langle s, *, * \rangle$ and $\langle s, p, * \rangle$ questions, as these methods require two anchor entities to perform traversal, whereas no object entity o is provided in these types.

Results show that PolyG constantly achieves the highest win rates across the 4 criteria in all 4 question types and the highest precision, recall and f1-score on the question set of type $\langle s, p, * \rangle$. This demonstrates the necessity of categorizing diverse graph questions and the effectiveness of adaptive graph traversal for different types. The win rates of PolyG align with a baseline method in each question type as PolyG selects the same traversal method and organize outputs in the same format with that specific baseline. Note that although Graph-CoT and ToG can generally adapt to various task scenarios by leveraging the reasoning adaptability of LLMs, they have a low win rate on all 4 question types as the LLM typically explores a subset of relevant entities and then reaches to the answer summary, which produces less comprehensive and diverse responses than methods that retrieves all relevant entities, e.g., MS_GraphRAG (BFS) and RoG (Guided Walk).

End-to-end latency. Table 6 presents the end-to-end response latency of PolyG and baseline methods across all three knowledge graphs. "N/A" indicates that a method is not applicable to the corresponding question type. The latency of PolyG includes the execution time of the se-

Table 5. Answer generation quality of our PolyG and the baselines on the *Literature* network. *SP* means Shortest Paths, and *CSP* refers to Constrained Shortest Paths. The win rates sum over 100% as we allow multiple winners. **Bold faces** denote the highest win rate.

Question Type	Criteria	MS_GraphRAG (BFS)	RoG (Guided Walk)	Fast-graphrag (PPR)	Graph-CoT (CoT)	Top-k SP	Top-k CSP	PolyG
$\langle s,*,*\rangle$	Comprehensiveness	**97.50%**	N/A	16.25%	28.75%	N/A	N/A	**97.50%**
	Diversity	**75.00%**	N/A	26.25%	33.75%	N/A	N/A	**75.00%**
	Empowerment	**87.50%**	N/A	21.25%	18.75%	N/A	N/A	**87.50%**
	Overall	**83.75%**	N/A	5.00%	11.25%	N/A	N/A	**83.75%**
$\langle s,p,*\rangle$	Comprehensiveness	22.50%	**81.25%**	35.00%	27.50%	N/A	N/A	**81.25%**
	Diversity	10.00%	**82.50%**	12.50%	7.50%	N/A	N/A	**82.50%**
	Empowerment	5.00%	**81.25%**	16.25%	7.50%	N/A	N/A	**81.25%**
	Overall	2.50%	**78.75%**	11.25%	7.50%	N/A	N/A	**78.75%**
	Precision	0.27	**0.77**	0.53	0.51	N/A	N/A	**0.77**
	Recall	0.23	**0.73**	0.39	0.35	N/A	N/A	**0.73**
	F1	0.24	**0.74**	0.43	0.38	N/A	N/A	**0.74**
$\langle s,*,o\rangle$	Comprehensiveness	42.50%	N/A	8.75%	8.75%	**88.75%**	2.50%	**88.75%**
	Diversity	32.50%	N/A	8.75%	7.50%	**81.25%**	2.50%	**81.25%**
	Empowerment	35.00%	N/A	5.00%	7.50%	**91.25%**	1.25%	**91.25%**
	Overall	20.00%	N/A	5.00%	1.25%	**78.75%**	1.25%	**78.75%**
$\langle s,p,o\rangle$	Comprehensiveness	8.75%	N/A	30.00%	10.00%	26.25%	**82.50%**	**82.50%**
	Diversity	20.00%	N/A	26.25%	13.75%	35.00%	**70.00%**	**70.00%**
	Empowerment	10.00%	N/A	18.75%	8.75%	25.00%	**73.75%**	**73.75%**
	Overall	6.25%	N/A	13.75%	5.00%	20.00%	**67.50%**	**67.50%**
Overall		28.13%	19.69%	8.75%	6.25%	24.69%	17.19%	**77.19%**

Table 6. End-to-end generation latency (s) and total token usage of PolyG and the baselines on all 3 datasets. The results are in the form *time/tokens*. *SP* stands for Shortest Paths, and *CSP* refers to Constrained Shortest Paths.

Dataset	Question Type	MS_GraphRAG (BFS)	RoG (Guided Walk)	Fast-graphrag (PPR)	Graph-CoT (CoT)	Top-k SP	Top-k CSP	PolyG
Academia	$\langle s,*,*\rangle$	10.25/8,412	N/A	32.84/73,041	41.67/30,767	N/A	N/A	11.01/8,931
	$\langle s,p,*\rangle$	5.13/7,789	11.33/1,127	28.18/64,816	44.14/42,352	N/A	N/A	12.11/1,651
	$\langle s,*,o\rangle$	7.6/13,352	N/A	28.78/72,218	57.11/48,929	25.5/1,107	16.21/1,346	26.31/1,640
	$\langle s,p,o\rangle$	77.84/44,637	N/A	27.42/55,738	72.14/164,723	9.89/661	19.46/1,873	20.24/2,406
Literature	$\langle s,*,*\rangle$	9.06/4,209	N/A	19.89/25,427	50.5/43,983	N/A	N/A	9.86/4,723
	$\langle s,p,*\rangle$	4.73/1,682	13.05/1,517	18.39/28,706	56.66/68,211	N/A	N/A	13.88/2,037
	$\langle s,*,o\rangle$	8.63/13,277	N/A	17.54/28,853	49.21/43,498	6.43/1,021	13.52/1,225	7.24/1,548
	$\langle s,p,o\rangle$	29.32/58,673	N/A	19.5/38,960	108.01/448,029	6.48/607	16.31/1,733	17.10/2,269
E-commerce	$\langle s,*,*\rangle$	7.82/6,885	N/A	64.79/61,330	44.34/36,018	N/A	N/A	8.65/7,406
	$\langle s,p,*\rangle$	5.36/7,040	11.59/1,357	59.77/41,113	50.58/54,298	N/A	N/A	12.39/1,885
	$\langle s,*,o\rangle$	7.84/17,523	N/A	58.73/45,560	52.06/54,464	6.32/1,428	12.85/1,204	7.14/1,971
	$\langle s,p,o\rangle$	16.2/40,325	N/A	59.74/39,660	83.24/239,014	6.17/840	16.7/1,980	17.49/2,517
Overall		15.82/18,650	11.99/1,334	36.30/47,952	59.14/106,191	10.13/944	15.84/1,560	13.62/3,249

lected traversal method plus the overhead of LLM-based question categorization. In some cases, PolyG exhibits higher latency than BFS and Top-k SP, as these methods retrieve less relevant context and yield lower win rates (see Table 4 and Table 5). Consequently, PolyG selects more effective traversal methods, such as Guided Walk and Top-k CSP, to ensure better context retrieval. Additionally, PolyG consistently outperforms PPR and CoT, as PPR requires dynamically computing PPR scores for each node during retrieval, and CoT invokes the LLM at every step of graph traversal—both incurring high computational complexity and resource consumption. Furthermore, BFS can also be inefficient in certain cases of the $\langle s,p,o\rangle$ question type due to the inherent skewness of graphs, where expanding all

neighbors of a high-degree entity significantly increases computational cost.

Token consumption. Table 6 also reports the total number of tokens consumed by PolyG and each baseline method. The token consumption of PolyG includes both the tokens required for answering a specific question type and the additional overhead from prompting the LLM to classify the question into one of the four predefined types. Results show that PolyG achieves the best generation quality while requiring the least number of tokens in most cases. There are instances where certain baseline methods consume fewer tokens than PolyG, such as Top-k SP for the $\langle s,p,o\rangle$ question type. This occurs because Top-k SP retrieves the shortest paths between the two given entities, whereas PolyG selects
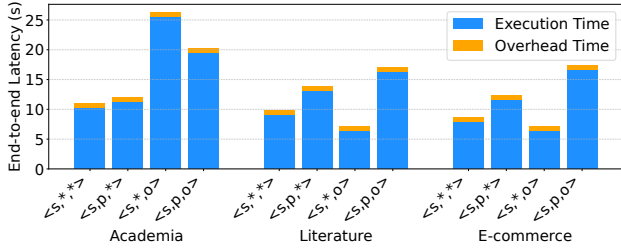
Figure 2. The time overhead of PolyG for question categorization in end-to-end latency across all 3 datasets and 4 question types.
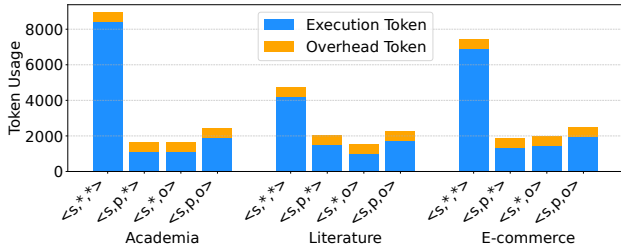

Figure 3. The token overhead of PolyG for question categorization in the total token usage across all 3 datasets and 4 question types.

the Top-k CSP method, which retrieves only the shortest reasoning paths that satisfy the relation constraint in the question. These paths are typically longer than those retrieved by Top-k SP. Despite its lower token usage, Top-k SP underperforms in generation quality, as it retrieves paths that focus on other relationships rather than those specified in the query, leading to less informative responses.

Question categorization overhead. Figure 2 and Figure 3 present the end-to-end latency and token usage overhead of PolyG for question categorization across the three knowledge graphs. Results indicate that classifying an input question incurs minimal overhead. Specifically, in terms of latency, PolyG introduces a constant overhead of approximately 0.8 seconds, which accounts for less than 10% of the total execution time for traversal and generation in all cases. Regarding token usage, the additional tokens required for question categorization never exceed 30% of the tokens used for the retrieved context. More importantly, even when factoring in the categorization overhead, the total token consumption of PolyG remains only 20% of that incurred by methods that retrieve excessive neighbors or frequently invoke the LLM, such as MS_GraphRAG, LightRAG, Fast-GraphRAG, and Graph-CoT.

Result summary. Figure 4 provides an overall comparison of PolyG and baseline methods in terms of both generation quality and system efficiency, averaged across the three
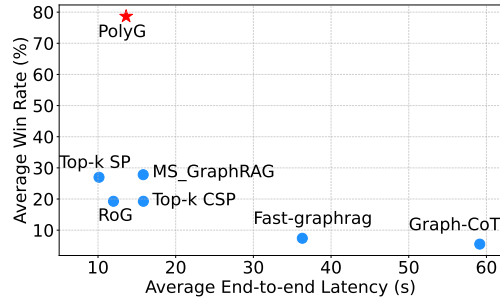

Figure 4. Overall comparison for PolyG and the baselines on generation quality (%) end end-to-end generation latency (s).

datasets. The results indicate that baseline methods exhibit low average win rates across all four question types, suggesting that relying on a fixed traversal method does not perform well across diverse graph questions. In particular, Fast-GraphRAG and Graph-CoT incur high computational costs in their traversal operations, leading to prolonged end-to-end response latency. In contrast, by dynamically selecting an appropriate traversal method for each question type, PolyG achieves the highest win rate while maintaining low end-to-end latency, demonstrating the necessity of question categorization and adaptive graph traversal.

6. Conclusions

Existing GraphRAG solutions either suffer from limited accuracy and applicability or incur high computation costs, as they adopt a single fixed graph traversal method for all graph questions or invoke the LLM at every traversal step. To address their limitations, we systematically analyze the graph questions and propose a taxonomy that categorizes the graph questions into four distinct types, based on both query target (i.e., entity or relation) and contents (i.e., specific or abstract). Building on this taxonomy, we design PolyG, a solution that dynamically selects the most suitable graph traversal method for each question type and executes the graph traversal operations with a unified interface and backend. Extensive experiments show that PolyG consistently achieves higher generation quality than SOTA GraphRAG methods while maintaining low end-to-end latency and token consumption across all four question types.

PolyG is essentially a query planner for GraphRAG, which is extensively studied in database but less explored for RAG and GraphRAG. We envision that holistic query categorization and planning may be promising directions for processing diverse and complex RAG and GraphRAG queries.

7. Impact Statement

While this work is focused solely on advancing the field of machine learning and is not expected to introduce any

societal risks, we acknowledge that there are certainly cases where the large language model (LLM) applications can be possibly developed to cause unintended social implications, whether deliberately or inadvertently. For instance, the improper use of GraphRAG could inadvertently expose confidential user data from internal company databases.

References

Streamlined and promptable fast graphrag framework designed for interpretable, high-precision, agent-driven retrieval workflows., 2024. URL https://github.com/circlemind-ai/fast-graphrag.

Anthropic. Claude 3.5 sonnet, 2024. URL https://www.anthropic.com/.

Bang, Y., Cahyawijaya, S., Lee, N., Dai, W., Su, D., Wilie, B., Lovenia, H., Ji, Z., Yu, T., Chung, W., Do, Q. V., Xu, Y., and Fung, P. A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity. In Park, J. C., Arase, Y., Hu, B., Lu, W., Wijaya, D., Purwarianti, A., and Krisnadhi, A. A. (eds.), *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 675–718, Nusa Dua, Bali, November 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.ijcnlp-main.45. URL https://aclanthology.org/2023.ijcnlp-main.45/.

Benchmark, O. G. Amazon product, 2022. URL https://ogb.stanford.edu/docs/nodeprop/#ogbn-products.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners, 2020. URL https://arxiv.org/abs/2005.14165.

Dam, S. K., Hong, C. S., Qiao, Y., and Zhang, C. A complete survey on llm-based ai chatbots, 2024. URL https://arxiv.org/abs/2406.16937.

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., and Larson, J. From local to global: A graph rag approach to query-focused summarization, 2024. URL https://arxiv.org/abs/2404.16130.

Garton, L., Haythornthwaite, C., and Wellman, B. Studying online social networks. *J. Comput. Mediat. Commun.*, 1997.

Graph, U. B. Goodreads, 2022. URL https://sites.google.com/eng.ucsd.edu/ucsdbookgraph/home.

Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pp. 855–864, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939754. URL https://doi.org/10.1145/2939672.2939754.

Guo, Z., Xia, L., Yu, Y., Ao, T., and Huang, C. Lightrag: Simple and fast retrieval-augmented generation, 2024. URL https://arxiv.org/abs/2410.05779.

Jin, B., Xie, C., Zhang, J., Roy, K., Zhang, Y., Li, Z., Li, R., Tang, X., Wang, S., Meng, Y., and Han, J. Graph chain-of-thought: Augmenting large language models by reasoning on graphs. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *62nd Annual Meeting of the Association for Computational Linguistics, ACL 2024 - Proceedings of the Conference*, Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 163–184. Association for Computational Linguistics (ACL), 2024.

Leipzig, D. G. Dblp-scholar, 2022. URL https://dbs.uni-leipzig.de/file/DBLP-Scholar.zip.

Luo, L., Li, Y.-F., Haf, R., and Pan, S. Reasoning on graphs: Faithful and interpretable large language model reasoning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ZGNWW7xZ6Q.

Neo4j. Neo4j - the world's leading graph database, 2012. URL http://neo4j.org/.

OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., and et al., S. A. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.

Page, L., Brin, S., Motwani, R., and Winograd, T. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999. URL http://ilpubs.stanford.edu:8090/422/. Previous number = SIDL-WP-1999-0120.

Perozzi, B., Al-Rfou, R., and Skiena, S. Deepwalk: online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD

'14, pp. 701–710, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450329569. doi: 10.1145/2623330.2623732. URL https://doi.org/10.1145/2623330.2623732.

Sahoo, P., Meharia, P., Ghosh, A., Saha, S., Jain, V., and Chadha, A. A comprehensive survey of hallucination in large language, image, video and audio foundation models. In Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 11709–11724, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.685. URL https://aclanthology.org/2024.findings-emnlp.685/.

Sun, J., Xu, C., Tang, L., Wang, S., Lin, C., Gong, Y., Ni, L., Shum, H.-Y., and Guo, J. Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=nnVO1PvbTv.

Talmor, A. and Berant, J. The web as a knowledge-base for answering complex questions. In Walker, M., Ji, H., and Stent, A. (eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 641–651, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1059. URL https://aclanthology.org/N18-1059/.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. Llama 2: Open foundation and fine-tuned chat models, 2023. URL https://arxiv.org/abs/2307.09288.

Wang, J., Huang, P., Zhao, H., Zhang, Z., Zhao, B., and Lee, D. L. Billion-scale commodity embedding for e-commerce recommendation in alibaba. In *KDD*, pp. 839–848, 2018.

Wang, L., Lyu, C., Ji, T., Zhang, Z., Yu, D., Shi, S., and Tu, Z. Document-level machine translation with large language models. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 16646–16661, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.1036. URL https://aclanthology.org/2023.emnlp-main.1036/.

Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., and Leiserson, C. E. Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *CoRR*, 2019.

Yang, Z., Qi, P., Zhang, S., Bengio, Y., Cohen, W., Salakhutdinov, R., and Manning, C. D. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL https://aclanthology.org/D18-1259/.

Yih, W.-t., Richardson, M., Meek, C., Chang, M.-W., and Suh, J. The value of semantic parse labeling for knowledge base question answering. In Erk, K. and Smith, N. A. (eds.), *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 201–206, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-2033. URL https://aclanthology.org/P16-2033/.

Zhang, Y., Li, Y., Cui, L., Cai, D., Liu, L., Fu, T., Huang, X., Zhao, E., Zhang, Y., Chen, Y., Wang, L., Luu, A. T., Bi, W., Shi, F., and Shi, S. Siren's song in the ai ocean: A survey on hallucination in large language models, 2023. URL https://arxiv.org/abs/2309.01219.

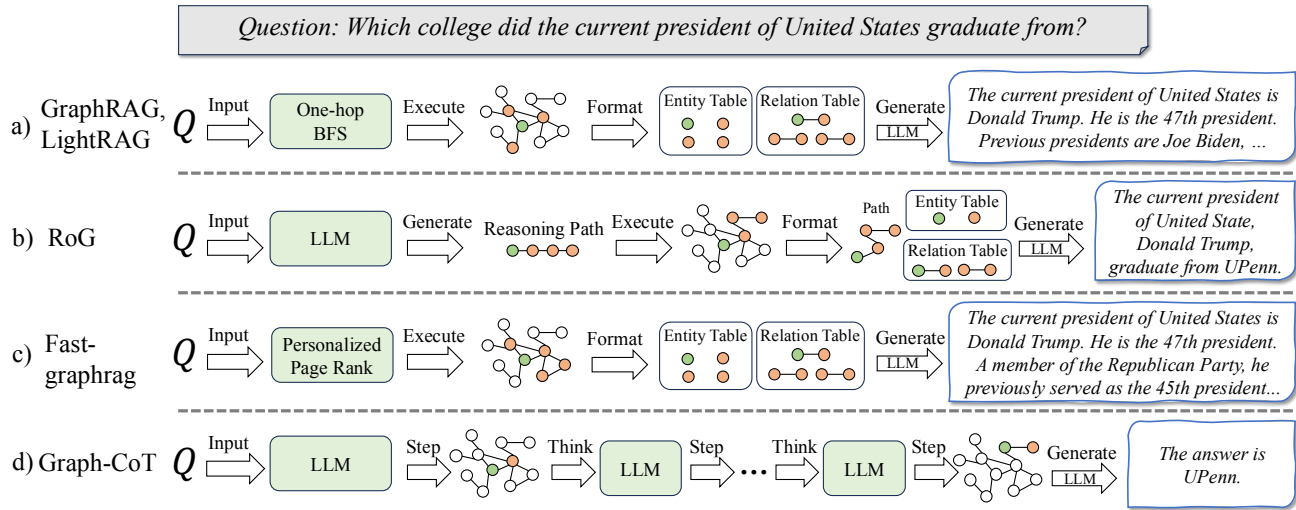# A. Overview of Existing GraphRAG Methods

Figure 5. Workflow demonstration and comparison of existing GraphRAG methods.

Many GraphRAG approaches are proposed to tackle GraphQA tasks or answer general user questions about the knowledge graph, and they differ in the adopted graph traversal methods. In the following, we introduce some representative GraphRAG methods and Figure 5 illustrates the workflow of them.

As shown in Figure 5a, GraphRAG (Edge et al., 2024) and LightRAG (Guo et al., 2024) use one-hop BFS to expand neighboring entities and relations of the anchor entities (entities that appear in the question), and return two tables as the retrieved context to the LLM as external grounding knowledge. While one-hop BFS is efficient in execution, this method fails to capture long-distance knowledge and thus cannot well answer questions that require reasoning paths beyond one-hop.

RoG (Luo et al., 2024), as depicted in Figure 5b, prompts the LLM to generate faithful reasoning paths that can be followed on the knowledge graph and returns the retrieved paths for generation. Faithful reasoning paths directly point from the anchor entity to the answer entities without involving useless noise in the context. However, it is only applicable when the question itself explicitly reveals concrete relation constraints. If the question is abstract in its predicates, e.g., asking about general information of an entity or relationships with another entity, no faithful reasoning paths could be given.

The workflow of Fast-graphrag (fas, 2024) is shown in Figure 5c. It retrieves entities and relations by running the well-studied random walk method, Personalized PageRank, and returns those with a relevance score greater than a threshold. Due to the diversity and complexity of graph questions, this fuzzing relevance matching that purely depends on the structure of knowledge graph might not accurately find what the question exactly requires.

The workflow of ToG (Sun et al., 2024) and Graph-CoT (Jin et al., 2024) is displayed in Figure 5d. They hand over the decision on every step of graph traversal to the LLM, prompting it with the current neighbors and all previous contexts to determine which neighbors to explore next. Traversal stops when the LLM thinks it reaches the answer, and then the answer entities are returned as results. This chain-of-thought manner is adaptable to various scenarios by leveraging the reasoning adaptability in LLMs. However, they end up at a high cost in both end-to-end latency and token consumption, which is expensive to use in real applications.

B. Results on E-commerce Network

Table 7 compares the generation quality of PolyG and baseline methods on the e-commerce network. Results show that PolyG selects the best-performing traversal strategy for each type of question and achieves the highest overall win rate (79.38%).

Table 7. Generation quality comparison between PolyG and baseline methods on the e-commerce network. SP means Shortest Paths and CSP refers to Constrained Shortest Paths. The win rates sum over 100% as we allow multiple winners. **Bold faces** denote the highest win rate.

Question Type	Criteria	GraphRAG (BFS)	RoG (Guided Walk)	Fast-graphrag (PPR)	Graph-CoT (CoT)	Top-k SP	Top-k CSP	PolyG
$\langle s, *, * \rangle$	Comprehensiveness	**83.75%**	N/A	41.25%	28.75%	N/A	N/A	**83.75%**
	Diversity	**77.50%**	N/A	30.00%	31.25%	N/A	N/A	**77.50%**
	Empowerment	**75.00%**	N/A	38.75%	27.50%	N/A	N/A	**75.00%**
	Overall	**72.50%**	N/A	15.00%	13.75%	N/A	N/A	**72.50%**
$\langle s, p, * \rangle$	Comprehensiveness	26.25%	**85.00%**	26.25%	26.25%	N/A	N/A	**85.00%**
	Diversity	21.25%	**77.50%**	16.25%	3.75%	N/A	N/A	**77.50%**
	Empowerment	17.50%	**85.00%**	20.00%	6.25%	N/A	N/A	**85.00%**
	Overall	12.50%	**82.50%**	6.25%	5.00%	N/A	N/A	**82.50%**
	Precision	0.3	**0.82**	0.4	0.49	N/A	N/A	**0.82**
	Recall	0.29	**0.83**	0.39	0.41	N/A	N/A	**0.83**
	F1	0.29	**0.82**	0.38	0.42	N/A	N/A	**0.82**
$\langle s, *, o \rangle$	Comprehensiveness	28.75%	N/A	2.50%	6.25%	**97.50%**	8.75%	**97.50%**
	Diversity	27.50%	N/A	5.00%	7.50%	**88.75%**	5.00%	**88.75%**
	Empowerment	26.25%	N/A	2.50%	5.00%	**92.50%**	7.50%	**92.50%**
	Overall	15.00%	N/A	0.00%	2.50%	**87.50%**	5.00%	**87.50%**
$\langle s, p, o \rangle$	Comprehensiveness	1.25%	N/A	21.25%	3.75%	46.25%	**90.00%**	**90.00%**
	Diversity	0.00%	N/A	21.25%	10.00%	56.25%	**72.50%**	**72.50%**
	Empowerment	1.25%	N/A	15.00%	2.50%	43.75%	**85.00%**	**85.00%**
	Overall	0.00%	N/A	6.25%	0.00%	26.25%	**75.00%**	**75.00%**
Overall		25.00%	20.63%	6.88%	5.31%	28.44%	20.00%	**79.38%**

C. Question Templates for Question Generation

C.1. Academia Network

This knowledge graph is a citation graph in physics, there are three types of nodes in this graph: paper, author and venue.

Graph Schema.

Node properties:

- type: author, properties: ["id", "name", "node_type"]
- type: paper, properties: ["id", "name", "label", "year", "node_type", "abstract"]
- type: venue, properties: ["id", "name", "node_type"]

Edge properties:

Author nodes are linked to their paper nodes by authorship.

- author → "paper" → paper

Paper nodes are linked to their author nodes, venue nodes, reference paper nodes and cited_by paper nodes.

- paper → "author" → author
- paper → "reference" → paper
- paper → "venue" → venue
- paper → "cited_by" → paper

Venue nodes are linked to their included paper nodes.

- venue → "paper" → paper

$\langle s, *, * \rangle$ **Question Template.**

- *Tell me about "[author name]".*

- *Tell me about "[paper name]".*

$\langle s, p, * \rangle$ **Question Template.**

- *What paper have the author "[author name]" published?*
- *What are the academic collaborators of "[author name]"?*
- *What venues have the author "[author name]" published in?*
- *Who are the authors of the paper "[paper name]"?*
- *Where is the paper "[paper name]" published?*
- *Who are the academic collaborators of the author who writes the paper "[paper name]"?*
- *What venues have the author of the paper "[paper name]" published in?*
- *What venues have the academic collaborators of the author who writes the paper "[paper name]" published in?*

$\langle s, *, o \rangle$ **Question Template.**

- *What is the relationship between "[author name]" and "[author name]"?*
- *What is the relationship between "[paper name]" and "[paper name]"?*
- *What is the relationship between "[author name]" and "[paper name]"?*

$\langle s, p, o \rangle$ **Question Template.**

- *Have the author "[author name]" cited or been cited by the work of the author "[author name]" and what are those works?*
- *Have authors "[author name]" and "[author name]" both collaborated with some other authors and who are they?*
- *Have the authors "[author name]" and "[author name]" ever published papers in the same venues? If so, tell me some examples.*
- *Do the venues "[venue name]" and "[venue name]" have the same authors publishing work in both of them and who are they?*

C.2. Literature Network

This knowledge graph is a literature graph names goodreads, there are four types of nodes in this graph: book, author, publisher and series.

Graph Schema.

Node properties:

- type: book, properties: ["id", "name", "node_type", "description", "publication_year", "genres"]
- type: author, properties: ["id", "name", "node_type"]
- type: publisher, properties: ["id", "name", "node_type"]
- type: series, properties: ["id", "name", "node_type", "description"]

Edge properties:

Book nodes are linked to neighboring book nodes, author nodes, publisher nodes and series nodes.

- book → "author" → author
- book → "publisher" → publisher
- book → "series" → series
- book → "similar_books" → book

Author nodes are linked to their neighboring book nodes.

- author → "book" → book

Publisher nodes are linked to their neighboring book nodes.

- publisher → "book" → book

Series nodes are linked to their neighboring book nodes. 1. series → "book" → book

⟨s, ∗, ∗⟩ **Question Template.**

- *Tell me about "[author name]".*
- *Tell me about "[book name]".*
- *Tell me about "[publisher name]".*
- *Tell me about "[series name]".*

⟨s, p, ∗⟩ **Question Template.**

- *Who are the authors of the book "[book name]"?*
- *What series have the author of the book "[book name]" published?*
- *What books has the author "[author name]" published?*
- *What books have the collaborators of the author "[author name]" published?*
- *What are the series published by the publishers that have published books of the author "[author name]"?*
- *What are the authors of the books published by the publisher "[publisher name]"?*
- *Where does the books of the series "[series name]" published in?*
- *What are the authors of the books that are published by the publishers that have published books of the series "[series name]"?*

⟨s, ∗, o⟩ **Question Template.**

- *What is the relationship between "[author name]" and "[author name]"?*
- *What is the relationship between "[book name]" and "[book name]"?*
- *What is the relationship between "[publisher name]" and "[publisher name]"?*
- *What is the relationship between "[series name]" and "[series name]"?*
- *What is the relationship between "[author name]" and "[book name]"?*

⟨s, p, o⟩ **Question Template.**

- *Have the authors "[author name]" and "[author name]" ever published books in the same publishers? If so, tell me some examples.*
- *Do the publishers "[publisher name]" and "[publisher name]" have any authors publishing books in both of them and what are the publications and authors?*
- *Do the publishers "[publisher name]" and "[publisher name]" have books that belong to the same series, and if so, what are those books?*
- *Do the series "[series name]" and "[series name]" contain books that are published by the same publisher? If so, tell me about them.*

C.3. E-commerce Network

This knowledge graph is an e-commerce graph in amazon, there are two types of nodes in this graph: item and brand.

Graph Schema.

Node properties:

- type: item, properties: ["id", "name", "node_type"]
- type: brand, properties: ["id", "name", "node_type"]

Edge properties: Item nodes are linked to neighboring item nodes and brand nodes.

- item → "also_viewed_item" → item
- item → "buy_after_viewing_item" → item
- item → "also_bought_item" → item
- item → "bought_together_item" → item
- item → "brand" → brand

Brand nodes are linked to their neighboring item nodes. Specific relations are:

- brand → "item" → item

$\langle s, *, * \rangle$ **Question Template.**

- *Tell me about "[brand name]".*
- *Tell me about "[item name]".*

$\langle s, p, * \rangle$ **Question Template.**

- *What is the brand of the item "[item name]"?*
- *What are the brands of the items that are also bought after viewing the item "[item name]"?*
- *What are the items that are also viewed when viewing items of the brand owning the item "[item name]"?*
- *What are the brands of the items that are also bought with items of the brand owning the item "[item name]"?*
- *What are the items of the brand "[brand name]"?*
- *What are the items that are bought together with items of the brand "[brand name]"?*
- *What are the brands of the items that are also bought with items of the brand "[brand name]"?*
- *What items does the brands of the items that are also viewed together with items of the brand "[brand name]" have?*

$\langle s, *, o \rangle$ **Question Template.**

- *What is the relationship between "[brand name]" and "[brand name]"?*
- *What is the relationship between "[item name]" and "[item name]"?*
- *What is the relationship between "[brand name]" and "[item name]"?*

$\langle s, p, o \rangle$ **Question Template.**

- *Have the items of the brands "[brand name]" and "[brand name]" ever both been also bought with some other items, and if so, what are those items?*
- *Have the items of the brands "[brand name]" and "[brand name]" ever both been bought after viewing some other items, and if so, what are those items?*
- *Have the items of the brands "[brand name]" and "[brand name]" ever been viewed together with some other items, and if so, what are those items?*
- *Have the items of the brands "[brand name]" and "[brand name]" ever been bought together with some other items, and if so, what are those items?*

D. Criteria Used in Evaluation

The detailed explanation of the evaluation metrics used in the experiments are as follows.

Comprehensiveness. This criterion how comprehensive an answer is. The prompt to the LLM is "How much detail the

answer provides to cover all aspects and details of the question?"

Diversity. This criterion evaluates diverse of the information provided in the answer is. The prompt to the LLM is "How varied and rich the answer is in providing different perspectives and insights on the question?"

Empowerment. This criterion evaluates how insightful an answer is. The prompt to the LLM is "How well the answer helps the reader understand and make informed judgments about the topic?"

Overall Winner. This is an overall judgment jointly considering the above 3 measures.

E. Prompts for Question Categorization.

Prompt:
You are an intelligent assistant tasked with classifying questions into four types based on the question target (entity or relation) and the specificity of the information being asked.

- **Ent-Abs(0):**
The question focuses on one entity and asks about general or abstract information about it (e.g., themes, concepts, or a general description).
Examples:
  - "Tell me about 'The Woman in Black: A Ghost Play'."
  - "Tell me about 'Frankenstein, or The Modern Prometheus'."
- **Ent-Con (1):**
The question focuses on one entity and asks about specific, concrete details related to it (e.g., names, authors, publishers, or other direct attributes).
Examples:
  - "Who are the authors of the book 'Sunshine for the Latter-Day Sa'?"
  - "What series have the author of the book 'Cookies for the Dragon (Saint Lakes, #2.1)' published?"
  - "What are the authors of the books that are published by the publishers that have published books of the series 'Shifter Justice'?"
- **Rel-Abs (2):**
The question involves multiple entities and focuses on their abstract relationships or interactions (e.g., conceptual connections, influences).
Examples:
  - "What is the relationship between 'Kizuna' and 'Kazuma Kodaka'?"
  - "What is the relationship between 'Monster House Press' and 'Matt Hart'?"
- **Rel-Con (3):**
The question involves multiple entities and focuses on concrete details related to their relationship (e.g., shared authors, publishers, or specific examples).
Examples:
  - "Have the authors 'Rubem Fonseca' and 'Lygia Fagundes Telles' ever published books in the same publishers? If so, tell me some examples."
  - "Do the publishers 'Scholastic Inc.' and 'Klutz' have any authors publishing books in both of them and what are the publications and authors?"

Instruction:
When given a question, analyze it based on the definitions above. Then, return **only a single number** corresponding to the type of the question:

- **0** for Ent-Abs
- **1** for Ent-Con
- **2** for Rel-Abs
- **3** for Rel-Con

Question: {}

Your Answer:

Question Categorization Prompt

Figure 6. The prompt for the LLM to categorize questions into one of the four types we defined.

In Figure 6, we provide the prompt we use for the LLM to judge which type a input question belongs to, including the definitions of each question type and some few-shot examples. For instance, the definition for the $\langle s, p, * \rangle$ type is "The object is missing in the facts and question focuses on finding the object, which is some specific, concrete details related to the subject (e.g., names, authors, publishers, or other direct attributes) with a concrete predicate p_c" and we give three examples "Who are the authors of the book 'Sunshine for the Latter-Day Sa'?", "What series have the author of the book 'Cookies for the Dragon (Saint Lakes, #2.1)' published?" and "What are the authors of the books that are published by the publishers that have published books of the series 'Shifter Justice'?" to the LLM to better understand how the questions in this type looks like.

F. Question Type in Public KBQA Benchmarks

We look into the popular public KBQA benchmarks WebQSP, and classify each of the question using the LLM (Claude-3.5-sonnet). In addtion to the prompt shown in Figure 6, we add an exception that tells the LLM the questions not necessarily fall into all those 4 types and output -1 for those unclassified ones with a possible decomposition into traversal-involved questions and non-traversal operators, as shown in Figure 7.

Prompt:
You are an intelligent assistant tasked with classifying questions into the following types based on the question target (entity or relation) and the specificity of the information being asked.

- **Ent-Abs(0):**
The question focuses on one entity and asks about general or abstract information about it (e.g., themes, concepts, or a general description).
Examples:
    - "Tell me about 'The Woman in Black: A Ghost Play'."
    - "Tell me about 'Frankenstein, or The Modern Prometheus'."
- **Ent-Con (1):**
The question focuses on one entity and asks about specific, concrete details related to it (e.g., names, authors, publishers, or other direct attributes).
Examples:
    - "Who are the authors of the book 'Sunshine for the Latter-Day Sa'?"
    - "What series have the author of the book 'Cookies for the Dragon (Saint Lakes, #2.1)' published?"
    - "What are the authors of the books that are published by the publishers that have published books of the series 'Shifter Justice'?"
- **Rel-Abs (2):**
The question involves multiple entities and focuses on their abstract relationships or interactions (e.g., conceptual connections, influences).
Examples:
    - "What is the relationship between 'Kizuna' and 'Kazuma Kodaka'?"
    - "What is the relationship between 'Monster House Press' and 'Matt Hart'?"
- **Rel-Con (3):**
The question involves multiple entities and focuses on concrete details related to their relationship (e.g., shared authors, publishers, or specific examples).
Examples:
    - "Have the authors 'Rubem Fonseca' and 'Lygia Fagundes Telles' ever published books in the same publishers? If so, tell me some examples."
    - "Do the publishers 'Scholastic Inc.' and 'Klutz' have any authors publishing books in both of them and what are the publications and authors?"
- **Others (-1):**
The question may not belong to any of the above 4 types. For those questions, if they can be decomposed into multiple step where some step the sub-question belongs to one of the four types above and other steps are non-traversal operators like "comparing entity attributes" or "doing a sort", then you should also provide a explanation for this on how it can be decomposed.
Examples:
    - "What are the 5 biggest cities in the USA?": 1. the first step is a [s,p_c,?] question, traverse from the USA and find all the cities. 2. non-traversal operator: get the population of the city and do a sort.
    - "Which city held the summer Olympics twice?": 1. the first step is a [s,p_c,?] question, traverse from summer Olympics to the cities. 2. non-traversal operator: count which cities appear twice in the results.

Instruction:
When given a question, analyze it based on the definitions above. Then, return **only a single number** corresponding to the type of the question:

- **0** for Ent-Abs
- **1** for Ent-Con
- **2** for Rel-Abs
- **3** for Rel-Con
- **-1, [decomposition]** for other unclassified questions

Question: {}

Question Categorization Prompt With Exception for Unclassified Questions

Your Answer:

Figure 7. The prompt for the LLM to categorize questions into one of the four types we defined with exceptions that questions not necessarily fall into the four types and the LLM is prompted to give possible question decomposition for those unclassified ones.

Results are that 96.45% questions in the WebQSP dataset fall into the $\langle s, p, * \rangle$ type of question while the other 3.55% are unclassified. There are no $\langle s, *, * \rangle$, $\langle s, *, o \rangle$ and $\langle s, p, o \rangle$ questions in WebQSP as these questions do not have a unique set of ground-truth. In the following, we demonstrate how the 3.55% unclassified questions can be decomposed into 1 out of the 4 types we defined in this work and some other non-traversal operators by some concrete examples.

❶ *"What are the 5 biggest cities in the USA?"*

Decomposed Steps:

- The first step is to answer a $\langle s, p, * \rangle$ question, which requires to traverse from the USA and find all the cities in there.
- The second step is a non-traversal operator to get the population attribute of each retrieved city.
- The third step is also a non-traversal operator to do sorting on the population and return the top-5 cities with the largest population.

❷ *"During what war did Abraham Lincoln serve as president?"*

Decomposed Steps:

- The first step is to answer a $\langle s, p, o \rangle$ question, which requires to traverse the links between Abraham Lincoln and the president.
- The second step is a non-traversal operator to get the time slot attribute of the edges linking Abraham Lincoln and the

president.

- The third step is also a non-traversal operator to scan over all wars, get the time slot of them and return those within the time slot when Abraham Lincoln serves as the president.

3 *"Which city held the summer Olympics twice?"*

Decomposed Steps:

- The first step is to answer a $\langle s, p, * \rangle$ question, which requires to traverse from the summer Olympics to the cities which have ever held it.
- The second step is a non-traversal operator to scan through the retrieved results and return those cities which appear twice in the retrieved entities.

4 *"What year did the Florida Marlins win their 2nd world series title?"*

Decomposed Steps:

- The first step is to answer a $\langle s, p, * \rangle$ question, which requires to traverse from the Florida Marlins to the World Series titles won by the them.
- The second step is a non-traversal operator to order these titles chronologically and return the year of the second title.

5 *"when was the last time the Kansas City Chiefs won a Superbowl?"*

- The first step is to answer a $\langle s, p, * \rangle$ question, which requires to traverse from the Kansas City Chiefs to all the Superbowls won by the them.
- The second step is a non-traversal operator to sort these wins by date and select the most recent one.

G. Case Study

In the following, we demonstrate how the LLM judges the generation quality of different traversal methods for each of the 4 question types by a concrete example using the academia network. We don't display PolyG in the following examples as PolyG is the same as the best performing method.

Table 8 shows the responses of different methods to an example question of $\langle s, *, * \rangle$ type. For the question "Tell me about 'superfluid transitions in bosonic atom molecule mixtures near a feshbach resonance'", the BFS method (which PolyG selects) adopted by GraphRAG achieves the best performance in all three criteria and wins the overall winner, by providing the most comprehensive, diverse and insightful response to this question. While GraphCoT also has the best diversity in its response, it falls short in the other two criteria.

Table 9 shows the responses of different methods to an example question of $\langle s, p, * \rangle$ type. For the question "What venues have the academic collaborators of the author who writes the paper 'dark world and baryon asymmetry from a common source' published in?", the Guided Walk method (which PolyG selects) adopted by RoG achieves the best performance in all three criteria and wins the overall winner, by providing the most comprehensive, diverse and insightful response to this question. Especially, GraphRAG does not give any valid venue names as BFS can only explores one-hop neighbors while the answer entities are 5 hops away in this question. Fast-graphrag acknowledges that it cannot determine the venues as the PPR score does not necessarily pick the 5-hop reasoning path between the paper and venues. While Graph-CoT can give some required venue names, the answer is not comprehensive and also does not provide any insights for the user.

Table 10 shows the responses of different methods to an example question of $\langle s, *, o \rangle$ type. For the question "What is the relationship between 'laser assisted pion decay' and 'A. Chahboune'?", the Top-k SP method (which PolyG selects) achieves the best performance in two of the three criteria and wins the overall winner, providing the most comprehensive and insightful response to this question. Especially, GraphRAG and Fast-graphrag fail to retrieve the relationship between them as GraphRAG only captures one-hop neighbors and Fast-graphrag uses the PPR score which does not necessarily pick the reasoning paths between the paper and the author in the question. Though Graph-CoT gives a thematic connection between the paper and the author's work, it does not find a relation connecting these two entities. While Top-k CSP also successfully points out the indirect relation through paper citing, it falls short in giving further insights to the user.

Table 11 shows the responses of different methods to an example question of $\langle s, p, o \rangle$ type. For the question "Do the venues

'progress of theoretical physics' and 'journal of geophysical research' have the same authors publishing work in both of them and who are they?", the Top-k CSP method (which PolyG selects) achieves the best performance in all the three criteria and wins the overall winner, providing the most comprehensive, diverse and insightful response to this question. Especially, GraphRAG, Fast-graphrag and Grpah-CoT all acknowledge that they do not have sufficient evidence to answer this question because they do not successfully retrieve the reasoning path between the two venues. Top-k SP also fails to answer the question because the shortest paths between these two venues reveal a paper-reference connection, rather than the common authorship relation required by the question.

Table 8. An example $\langle s, *, * \rangle$ question on the academia network with responses of different methods and the LLM decision.

Question	Tell me about 'superfluid transitions in bosonic atom molecule mixtures near a feshbach resonance'.
GraphRAG (BFS)	"Superfluid transitions in bosonic atom-molecule mixtures near a Feshbach resonance" is a scientific paper published in Physical Review Letters. Here are the key points about this paper: 1. Authors: The paper was co-authored by Leo Radzihovsky and Peter B. Weichman. 2. Topic: It studies the behavior of bosonic atoms near a Feshbach resonance, which is a quantum mechanical phenomenon used to control interactions between atoms in ultracold gases. 3. Key findings: The paper predicts that in addition to standard normal and atomic superfluid phases, this system can exhibit a distinct phase of matter called a molecular superfluid. In this phase, molecules are superfluid while atoms are not. 4. Phases: The study explores zero and finite temperature properties of three phases: Normal phase, Atomic superfluid and Molecular superfluid 5. Transitions: The paper examines quantum and classical phase transitions between these states. 6. Significance: The molecular superfluid is described as a bosonic strong coupling analog of a BCS superconductor, which is an important concept in condensed matter physics. 7. Methodology: While specific details aren't provided, the study likely involves theoretical modeling and analysis of bosonic systems near Feshbach resonances. This paper appears to be significant in the field of ultracold atomic physics, particularly in understanding novel quantum phases and transitions in bosonic systems. It builds upon and relates to numerous other studies in this area, as evidenced by the many references to other papers on related topics.
RoG (Guided Walk)	N/A
Fast-graphrag (PPR)	Based on the information provided in the input data, "superfluid transitions in bosonic atom molecule mixtures near a feshbach resonance" refers to a scientific paper that studies the behavior of bosonic atoms near a Feshbach resonance. The paper predicts that in addition to standard normal and atomic superfluid phases, this system exhibits a distinct phase of matter called a molecular superfluid, where molecules are superfluid while atoms are not. The paper explores zero and finite temperature properties of the molecular superfluid, which is described as a bosonic strong coupling analog of a BCS superconductor. It also studies quantum and classical phase transitions between the normal, molecular superfluid, and atomic superfluid states. This research provides insights into the complex behavior of bosonic atom-molecule mixtures near Feshbach resonances, particularly focusing on the emergence of a novel molecular superfluid phase and the transitions between superfluid states.
Graph-CoT (CoT)	The paper "Superfluid transitions in bosonic atom molecule mixtures near a feshbach resonance" was published in Physical Review Letters in 2004. It was authored by Peter B. Weichman, Jae Park, and Leo Radzihovsky. The paper studies bosonic atoms near a Feshbach resonance and predicts the existence of a distinct phase of matter called a molecular superfluid, where molecules are superfluid while atoms are not. This is in addition to the standard normal and atomic superfluid phases. The authors explore zero and finite temperature properties of the molecular superfluid, which they describe as a bosonic strong coupling analog of a BCS superconductor. They also study quantum and classical phase transitions between the normal, molecular superfluid, and atomic superfluid states. This research contributes to our understanding of quantum states of matter and phase transitions in bosonic systems, particularly in the context of atom-molecule mixtures near Feshbach resonances.
Top-K SP	N/A
Top-K CSP	N/A
LLM Decision	**Comprehensiveness: Winner=Method BFS** Method BFS provides the most comprehensive answer, covering multiple aspects of the paper including authors, topic, key findings, phases, transitions, significance, and methodology. It also mentions that the paper relates to numerous other studies in the field, giving a broader context. **Diversity: Winner=Method BFS, Method GraphCoT** Method BFS offers the most diverse information, touching on various aspects of the paper. Method GraphCoT adds some unique information not mentioned in other responses, such as the publication year (2004) and an additional author (Jae Park), contributing to the diversity of information provided. **Empowerment: Winner=Method BFS** Method BFS empowers the reader the most by providing a structured breakdown of the paper's content, explaining key concepts like Feshbach resonance, and highlighting the significance of the study in the field of ultracold atomic physics. This approach helps the reader understand the topic more thoroughly and its importance in the broader scientific context. **Overall Winner: Winner=Method BFS** Method BFS is the overall winner as it excels in all three criteria. It provides the most comprehensive overview of the paper, covers diverse aspects of the research, and empowers the reader with a well-structured and informative response. While other methods also provide valuable information, Method BFS stands out for its thoroughness and ability to contextualize the research within the broader field of study.

Table 9. An example $\langle s, p, * \rangle$ question on the academia network with responses of different methods and the LLM decision.

Question	What venues have the academic collaborators of the author who writes the paper 'dark world and baryon asymmetry from a common source' published in?
GraphRAG (BFS)	Based on the information provided in the data tables, I can only identify one author for the paper "dark world and baryon asymmetry from a common source": Daijiro Suematsu. The data does not provide information about other academic collaborators of Daijiro Suematsu or their publication venues. The only venue mentioned in relation to this paper is: Journal of Cosmology and Astroparticle Physics. This is the venue where dark world and baryon asymmetry from a common source was published. However, we don't have information about other venues where Daijiro Suematsu's collaborators may have published their work. To properly answer this question, we would need additional data about Daijiro Suematsu's collaborators and their publication records, which is not provided in the given information.
RoG (Guided Walk)	Based on the information provided in the table, the academic collaborators of the author who wrote Dark World and Baryon Asymmetry from a Common Source have published in several prestigious physics journals. These venues include: 1. Physical Review D 2. International Journal of Modern Physics A 3. Progress of Theoretical Physics 4. Journal of High Energy Physics 5. European Physical Journal C 6. Physics Letters B 7. Nuclear Physics 8. Astroparticle Physics 9. Journal of Cosmology and Astroparticle Physics This diverse list of publication venues suggests that the author's collaborators have a broad range of expertise in theoretical physics, particle physics, cosmology, and related fields. These journals are well-respected in the physics community and cover topics ranging from high-energy physics to cosmology and astrophysics, which aligns with the themes present in the paper "Dark World and Baryon Asymmetry from a Common Source".
Fast-graphrag (PPR)	Based on the provided input data, I cannot determine the venues that academic collaborators of the author of "Dark World and Baryon Asymmetry from a Common Source" have published in. The information given does not include details about collaborators of the paper's author or their publication venues. The only venue information provided is for the paper itself, which was published in the Journal of Cosmology and Astroparticle Physics. There is no additional information about other authors or their publications in the given data.
Graph-CoT (CoT)	Physical Review D, International Journal of Modern Physics A
Top-K SP	N/A
Top-K CSP	N/A
LLM Decision	**Comprehensiveness: Winner=Method Guided Walk** Method Guided Walk provides the most comprehensive answer, listing nine specific publication venues that align with the ground truth. It covers a wide range of journals in the field of physics, which matches closely with the expected answer. Other methods either fail to provide specific venues or acknowledge a lack of information. **Diversity: Winner=Method Guided Walk** Method Guided Walk offers the most diverse response, mentioning various types of physics journals ranging from general physics to specialized fields like cosmology and particle physics. This diversity aligns well with the ground truth, which includes a wide array of physics-related journals. **Empowerment: Winner=Method Guided Walk** Method Guided Walk empowers the reader by providing a comprehensive list of publication venues, giving context about their prestige and relevance to the field. It also explains the significance of these venues in relation to the original paper's topic, helping the reader understand the breadth of the collaborators' expertise. **Overall Winner: Winner=Method Guided Walk** Method Guided Walk is the clear overall winner as it excels in all three criteria. It provides the most comprehensive and diverse list of publication venues, closely matching the ground truth. The response also empowers the reader by offering context and explanations about the significance of these venues. While it doesn't capture every journal mentioned in the ground truth, it provides the most accurate and informative answer among all the responses.

Table 10. An example $\langle s, *, o \rangle$ question on the academia network with responses of different methods and the LLM decision.

Question	What is the relationship between 'laser assisted pion decay' and 'A. Chahboune'?
GraphRAG (BFS)	Based on the information provided in the data tables, there is no direct relationship mentioned between 'laser assisted pion decay' and 'A. Chahboune'. The tables show that 'A. Chahboune' is an author associated with the paper "relativistic proton impact excitation of hydrogen atom in the presence of intense laser field". However, this paper is not listed as a reference for or cited by the 'laser assisted pion decay' paper. The 'laser assisted pion decay' paper has several authors listed (Bouzid Manaut, S. Taj, E. Siher, S. Mouslih, and M. Jakha), but A. Chahboune is not among them. Without additional information, I cannot establish a direct relationship between 'laser assisted pion decay' and 'A. Chahboune' based on the given data.
RoG (Guided Walk)	N/A
Fast-graphrag (PPR)	Based on the information provided in the Relationships table, A. Chahboune is not directly related to the paper "laser assisted pion decay". There is no relationship listed between A. Chahboune and this paper. The authors listed for "laser assisted pion decay"are M. Jakha, Bouzid Manaut, S. Taj, E. Siher, and S. Mouslih, but A. Chahboune is not mentioned among them.
Graph-CoT (CoT)	A. Chahboune is indirectly connected to the paper "laser assisted pion decay" through shared co-authors. While A. Chahboune did not author or directly cite this paper, they have collaborated with Bouzid Manaut and S. Taj on a related paper titled "Relativistic proton impact excitation of hydrogen atom in the presence of intense laser field. Both papers involve research on particle interactions with intense laser fields, suggesting a thematic connection in their research areas.
Top-K SP	Based on the provided information, there is an indirect relationship between 'laser assisted pion decay' and 'A. Chahboune'. The paper "laser assisted pion decay" is connected to A. Chahboune through multiple paths of citations and references involving other papers and authors. Specifically: 1. "Laser assisted pion decay" is cited by or references several papers on Mott scattering in laser fields. 2. These Mott scattering papers are then cited by or referenced in a paper titled "relativistic proton impact excitation of hydrogen atom in the presence of intense laser field". 3. A. Chahboune is listed as an author of this "relativistic proton impact excitation" paper. So while A. Chahboune did not directly author "laser assisted pion decay", their work is connected through a chain of related research in the field of particle interactions in the presence of laser fields. This suggests that A. Chahboune's research may have been influenced by or built upon concepts discussed in "laser assisted pion decay", or that both works contribute to a broader area of study in laser-assisted particle physics.
Top-K CSP	Based on the provided information, there is an indirect relationship between 'laser assisted pion decay' and 'A. Chahboune'. The relationship can be traced through multiple hops: 'Laser assisted pion decay' is cited by papers on Mott scattering (either "in the presence of a linearly polarized laser field" or "in strong laser fields"). These Mott scattering papers, in turn, cite or reference a paper titled "relativistic proton impact excitation of hydrogen atom in the presence of intense laser field". A. Chahboune is listed as an author of this last paper.So while there isn't a direct connection, A. Chahboune's work on relativistic proton impact excitation is indirectly linked to research on laser assisted pion decay through intermediate papers on Mott scattering in laser fields.
LLM Decision	**Comprehensiveness: Winner=Method Top-k CSP, Method Top-k SP** Both Method Top-k CSP and Method Top-k SP provide the most comprehensive answers. They describe the indirect relationship between 'laser assisted pion decay' and 'A. Chahboune' through intermediate papers on Mott scattering. They explain the chain of connections from the 'laser assisted pion decay' paper to A. Chahboune's work, showing a detailed understanding of the research network. **Diversity: Winner=Method GraphCoT** Method GraphCoT offers a unique perspective by highlighting the indirect connection through shared co-authors. While other methods focus on citation paths, this response introduces the idea of collaborative networks, adding a different dimension to the relationship between the two entities. **Empowerment: Winner=Method Top-k SP** Method Top-k SP provides the most empowering answer. It not only explains the indirect relationship but also contextualizes it within the broader field of laser-assisted particle physics. This response helps the reader understand the significance of the connection and how it fits into the larger research landscape, enabling them to make informed judgments about the relevance of these works to each other. **Overall Winner: Winner=Method Top-k SP** Method Top-k SP emerges as the overall winner by excelling in all three criteria. It provides a comprehensive explanation of the indirect relationship, offers diverse insights by detailing the citation path, and empowers the reader by contextualizing the connection within the broader field of study. This response gives the most complete and nuanced understanding of the relationship between 'laser assisted pion decay' and 'A. Chahboune', allowing readers to grasp both the specific connection and its wider implications in the research area.

Table 11. An example $\langle s, p, o \rangle$ question on the academia network with responses of different methods and the LLM decision.

Question	**Do the venues 'progress of theoretical physics' and 'journal of geophysical research' have the same authors publishing work in both of them and who are they?**
GraphRAG (BFS)	Based on the information provided in the data tables, I don't have enough information to determine if there are any authors who have published work in both "Progress of Theoretical Physics" and "Journal of Geophysical Research", or who those authors might be. The data does not contain information about specific authors or their publications across different venues. To answer this question, we would need additional data that links authors to their publications in each venue.
RoG (Guided Walk)	N/A
Fast-graphrag (PPR)	Based on the information provided in the input data, there is no evidence to determine if 'progress of theoretical physics' and 'journal of geophysical research' have the same authors publishing work in both of them. The data only shows that both are venues, but does not provide any information about authors or their publications across these venues. Therefore, I cannot answer the user's query about shared authors or identify who they might be.
Graph-CoT (CoT)	Sorry, I cannot determine if there are any authors publishing work in both of the venues 'progress of theoretical physics' and 'journal of geophysical research'.
Top-K SP	Based on the information provided in the data tables, I cannot definitively answer whether the venues 'Progress of Theoretical Physics' and 'Journal of Geophysical Research' have the same authors publishing work in both of them. The data does not include information about specific authors or their publications across these venues. However, the data does show that there are several connections between these two venues through paper references. Multiple papers published in 'Progress of Theoretical Physics' reference papers published in the 'Journal of Geophysical Research', and vice versa. This suggests that there is some overlap in the research areas covered by these two venues, and it's possible that some researchers may publish in both. Without more specific information about authors and their publications, I cannot identify who these authors might be. To answer this question definitively, we would need additional data about the authors of papers in these venues.
Top-K CSP	Yes, based on the provided information, there are several authors who have published work in both 'Progress of Theoretical Physics' and 'Journal of Geophysical Research'. These authors are: 1. S. Yoshida, 2. Y. J. Kwon, 3. Hidemi Ito and 4. Hiroyuki Sagawa. These authors have papers published in 'Progress of Theoretical Physics', and they are also connected to papers published in the 'Journal of Geophysical Research'. Specifically, they are all authors of the paper "Gamma ray showers observed at ground level in coincidence with downward lightning leaders" which appears in the Journal of Geophysical Research. It's worth noting that the data doesn't provide direct information about all of their publications, but rather shows connections between these authors, their papers in Progress of Theoretical Physics, and at least one paper in the Journal of Geophysical Research.
LLM Decision	**Comprehensiveness: Winner=Method Top-k CSP** This method provides the most comprehensive answer by clearly stating that there are authors who have published in both venues and identifying four specific authors. It also mentions a specific paper title that connects these authors to the Journal of Geophysical Research. **Diversity: Winner=Method Top-k CSP, Method Top-k SP** Method Top-k CSP offers a direct answer with specific author names. Method Top-k SP, while not providing author names, adds diversity by mentioning the connections between the two venues through paper references, suggesting overlap in research areas. This provides a different perspective on the relationship between the two venues. **Empowerment: Winner=Method Top-k CSP** This method empowers the reader the most by providing concrete information about authors who have published in both venues. It gives specific names and even mentions a paper title, allowing the reader to potentially follow up on this information for further research. **Overall Winner: Winner=Method Top-k CSP** Method Top-k CSP is the overall winner as it directly answers the question with specific author names and a paper title. It provides the most comprehensive information, empowering the reader with concrete details. While it doesn't offer as much diversity in perspective as some other answers, the specificity and relevance of its information make it the most valuable response to the given question.