

Graph Chain-of-Thought: Augmenting Large Language Models by Reasoning on Graphs

Bowen Jin¹, Chulin Xie¹, Jiawei Zhang¹, Kashob Kumar Roy¹, Yu Zhang¹
Zheng Li², Ruirui Li², Xianfeng Tang², Suhang Wang³, Yu Meng⁴, Jiawei Han¹

¹University of Illinois at Urbana-Champaign, ²Amazon

³Pennsylvania State University, ⁴University of Virginia

bowenj4@illinois.edu

Abstract

Large language models (LLMs), while exhibiting exceptional performance, suffer from hallucinations, especially on knowledge-intensive tasks. Existing works propose to augment LLMs with individual text units retrieved from external knowledge corpora to alleviate the issue. However, in many domains, texts are interconnected (*e.g.*, academic papers in a bibliographic graph are linked by citations and co-authorships) which form a (text-attributed) graph. The knowledge in such graphs is encoded not only in single texts/nodes but also in their associated connections. To facilitate the research of augmenting LLMs with graphs, we manually construct a **Graph Reasoning Benchmark** dataset called GRBENCH, containing 1,740 questions that can be answered with the knowledge from 10 domain graphs. Then, we propose a simple and effective framework called Graph Chain-of-thought (GRAPH-COT) to augment LLMs with graphs by encouraging LLMs to reason on the graph iteratively. Each GRAPH-COT iteration consists of three sub-steps: LLM reasoning, LLM-graph interaction, and graph execution. We conduct systematic experiments with three LLM backbones on GRBENCH, where GRAPH-COT outperforms the baselines consistently. The code is available at <https://github.com/PeterGriffinJin/Graph-CoT>.

1 Introduction

Large language models (LLMs) (Touvron et al., 2023; Jiang et al., 2024) have demonstrated their exceptional language understanding and text generation capability in real-world scenarios (Zhao et al., 2023). However, LLMs suffer from hallucination problems and sometimes tend to generate content that appears plausible but is ungrounded (Tonmoy et al., 2024). This is because they memorize world knowledge parametrically and fail to refer to concrete knowledge sources (Zhang et al.,

2023b). To alleviate the hallucination issues, existing works propose to augment LLMs with external text corpora as knowledge sources (Shuster et al., 2021; Wu et al., 2023) and treat every single document as a knowledge unit. Retrieval augmentation (RAG) (Lewis et al., 2020; Gao et al., 2023) is then proposed to enable LLMs to interact with external knowledge sources, where relevant texts are retrieved and serve as contexts to improve the factuality of LLMs (shown in Figure 1 (a)). However, retrieval augmentation assumes that knowledge is well represented in individual text units and ignores the correlations among multiple text units.

In real-world scenarios, text units are generally interconnected, forming a (text-attributed) graph. The knowledge of such graphs is reflected not only in the form of texts but also in the structure of their connections. For example, academic papers in a bibliographic graph are linked by citation links (Wang et al., 2020). We can trace the source of a research direction (Bai et al., 2019) by traversing such a graph. Cases and opinions in a legal graph are interconnected by reference edges (Sadeghian et al., 2018). We can verify the judgment for a case by looking up its citations on such a graph (Chen et al., 2019).

Although widely used for text corpora as external knowledge sources, retrieval-augmentation cannot be readily used to augment LLMs with graphs for two reasons: 1) *Structure Context*: Retrieval augmentation can find individual nodes/texts from the graphs which can serve as context to augment the LLMs. However, knowledge on the graph also lies in the structure which can not be captured by single nodes/texts. 2) *Graph Size Explosion*: Although it is feasible to convert local subgraph structures into text descriptions as the input contexts to LLMs, the size of the local subgraph increases exponentially as the hop number increases, resulting in an excessively long context sequence. This could cause LLMs to be lost in the middle (Liu

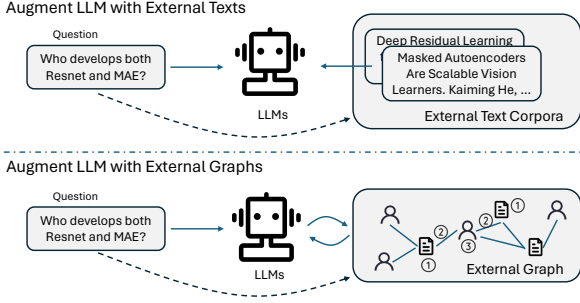


Figure 1: Augmenting LLMs with external text corpora or external text-attributed graph.

et al., 2023) given a plethora of irrelevant information in the context. In addition, the long sequence could potentially exceed the input length limitations of LLMs (Zhao et al., 2023).

Therefore, it is an important research topic to augment LLMs with such graph information. Unfortunately, there has been a lack of benchmark datasets to support the development of methodology and facilitate the evaluation of the proposed models. To this end, we first construct a **Graph Reasoning** benchmark dataset called **GRBENCH**. **GRBENCH** includes ten real-world graphs that can serve as external knowledge sources for LLMs from five domains including academic, e-commerce, literature, healthcare, and legal domains. Each sample in **GRBENCH** consists of a manually designed question and an answer, which can be directly answered by referring to the graphs or retrieving the information from the graphs as context. To make the dataset comprehensive, we include samples of different difficulty levels: easy questions (which can be answered with single-hop reasoning on graphs), medium questions (which necessitate multi-hop reasoning on graphs), and hard questions (which call for inductive reasoning with information on graphs as context).

We propose a simple and effective framework called **Graph Chain-of-thought (GRAPH-COT)**. The main idea is to enable LLMs to traverse the graph step-by-step to figure out the key information needed, rather than directly feeding the whole subgraph as context into the LLMs (shown in Figure 1 (b)). **GRAPH-COT** is an iterative framework, where one iteration corresponds to one step on the graph. Each iteration in **GRAPH-COT** consists of three sub-steps: 1) *Reasoning*: LLMs propose what conclusion we can make with the current information and what further information is needed from the graph; 2) *Interaction*: LLMs generate the interactions needed to fetch information from the

graph (e.g., finding the nodes, checking the neighbors, etc); 3) *Execution*: The requests from the interaction step are executed on the graph and the corresponding information is returned. In this way, LLMs can conduct chain-based reasoning on the graph and find the key information on the graph. This process will be iterated until LLMs conclude the final answer in the reasoning sub-step.

In summary, our contributions are as follows:

- We propose the problem of augmenting LLMs with external graphs and introduce a comprehensive benchmark dataset called **GRBENCH**.
- We develop a straightforward and effective framework **GRAPH-COT** to encourage the LLMs to reason on the graph iteratively.
- We conduct extensive experiments on **GRBENCH** to demonstrate the effectiveness of **GRAPH-COT** and analyze its performance across different demonstration settings, backbone LLMs, and questions difficulties. Furthermore, we explore its failure cases with future directions outlined.

2 Preliminaries

Definition 2.1. Graph. A graph can be denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} and \mathcal{E} are node set and edge set, respectively. Each $v_i \in \mathcal{V}$ can be associated with some feature information \mathcal{X}_{v_i} . For example, in an e-commerce item graph, $v \in \mathcal{V}$ are items, $e \in \mathcal{E}$ are co-purchase edges, and \mathcal{X}_v include features such as item title, description, price, and category. In this work, we formulate all the features as texts and the graph is also called a text-attributed graph.

Definition 2.2. Neighbors and Degree. The neighbors of a node v_i refer to nodes which are linked to v_i on the graph, denoted as $N(v_i) = \{v_j | e_{v_i, v_j} \in \mathcal{E}\}$. The degree of a node v_i refers to the number of v_i 's neighbors, denoted as $D(v_i) = |N(v_i)|$.

3 GRBENCH Dataset

3.1 Dataset Overview

We create the **GRBENCH** dataset to evaluate how effectively LLMs can interact with domain-specific graphs containing rich knowledge to solve the desired problem. **GRBENCH** contains 10 graphs from 5 general domains (academia, e-commerce, literature, healthcare, and legal). Each data sample in **GRBENCH** is a question-answer pair. The questions are designed to simulate the real-world use cases in specific domains. However, it is hard for LLMs to answer those questions directly with their internal knowledge stored in model parameters;

Table 1: Dataset Statistics of GRBENCH.

Domain	Topic	Graph Statistics		Data	
		# Nodes	# Edges	# Templates	# Questions
Academic	CS	~8M	~52M	15	150
	Biology	~4M	~39M	14	140
	Chemistry	~4M	~30M	14	140
	Material Science	~3M	~22M	14	140
	Medicine	~6M	~30M	14	140
	Physics	~2M	~33M	14	140
E-commerce	Amazon	~9M	~313M	20	200
Literature	Goodreads	~3M	~22M	24	240
Healthcare	Disease	~47K	~4M	27	270
Legal	Freelaw	~84M	~114M	18	180
SUM	-	-	-	174	1740

they need to interact with external domain-specific graphs. The overall statistics of GRBENCH are in Table 1.

To curate high-quality and diverse data without heavy human effort, the construction of GRBENCH contains four steps: 1) We first collect large *reference graph data* from real-world scenarios which can serve as the context for data generation. 2) Then, we manually *design question templates* which can be answered on the reference graph data. 3) After that, we call GPT-4 to generate diverse *question expressions* for each question template. 4) Finally, we *automatically generate ground truth answers* from the domain-specific graphs.

3.2 Reference Graph Data

We collect data from five domains where the knowledge lies in the format of graphs: academia, e-commerce, literature, healthcare, and legal. The detailed statistics of the graphs can be found in Appendix Table 5.

In the **academic domain**, papers, authors, and venues are naturally interconnected by citation, “written-by”, and “publish-in” relations. We obtain academic graphs across six disciplines including Biology, Computer Science, Chemistry, Material Science, Medicine, and Physics from DBLP¹ (Tang et al., 2008) and Microsoft Academic Graph (MAG)² (Wang et al., 2020; Zhang et al., 2023a). Nodes on such graphs are papers, authors, and venues, while edges include citation edges, authorship edges, and venueship edges.

In the **e-commerce domain**, a single product is assigned a brand, and different products are interlinked through “also-viewed” or “also-bought” relationships, which naturally embody graph-like

structures. We use Amazon product datasets³ (He and McAuley, 2016), which provides the metadata information of items across a myriad of product categories. Nodes on this graph are items and brands, while edges include “also-viewed”, “also-bought”, “buy-after-viewing”, “bought-together”, and “item-brand”.

In the **literature domain**, the inherent graph structure exists with interconnections between books, authors, publishers, and series. The Goodreads dataset⁴ (Wan and McAuley, 2018) offers an extensive collection of books with their metadata. Nodes on this graph are books, authors, publishers, and series, while edges include “written-by”, “publish-in”, “book-series” and so on.

In the **healthcare domain**, we can construct a graph by considering the diseases with their associated properties. We adopt the biological disease graph Hetionet⁵ (Himmelstein et al., 2017), which comprehensively summarizes existing disease and their symptoms, with the aim of repurposing drugs. Nodes on this graph include diseases, symptoms, side effects, compounds, and so on, while edges include “disease-present-symptom”, “compound-cause-side effect” and so on.

In the **legal domain**, there are rich citation links between cases and opinions (since judges rely on citing opinions from previous cases to write for the current case) which naturally form a graph. We use the data from CourtListener⁶. Nodes on this graph are opinion, opinion-cluster, docket, and court, while edges include “opinion-citation”, “opinion-cluster”, “cluster-docket”, and “docket-court”.

3.3 Manually Designed Question Templates

The question generation phase aims to generate questions that can be answered by LLMs after referring to the domain graphs. Considering that the generated questions should be accurate and meaningful, we ask four well-trained computer science Ph.D. students to write potential questions that can be answered given the graphs as context.

To comprehensively evaluate the LLMs and their capability to interact with graphs, we ask the annotators to design question templates of three different difficulties:

- **Easy:** These questions can be answered by

³<https://cseweb.ucsd.edu/~jmcauley/datasets/amazon/links.html>

⁴<https://mengtingwan.github.io/data/goodreads>

⁵<https://github.com/hetio/hetionet>

⁶<https://www.courtlistener.com/>

¹https://originalfileserver.aminer.cn/misc/dblp_v14.tar.gz

²<https://zenodo.org/records/7611544>

looking up the feature/degree of only one node or travel on the graph within one hop. For example, “What is the price of the {item}?” or “Who are the authors of {paper}?”

- **Medium:** These questions require reasoning on the graphs for more than one hop and involve returning the feature/degree of nodes. For example, “Who is the closest collaborator with {author} in {year}?”.
- **Hard:** These questions cannot be directly answered by looking up the graph, but the graph can be useful by providing informative context. For example, “What is the complementary item given this {query}?”

It is worth noting that the easy-level and medium-level questions can be answered from the given graph, while the ground truth for hard questions cannot be directly found in the graph. All the question templates can be found in Appendix B.

Once the question templates are manually designed, we extract values from the graph to transform the templates into actual questions. For example, given the question template “How many citations did {paper} have in {year}?”, we can refer to the academic graphs and sample “Language Models are Unsupervised Multitask Learners” as the “paper” value and “2021” as the “year” value. This will result in a real question: “How many citations did Language Models are Unsupervised Multitask Learners have in 2021?”

3.4 Diverse Question Expression with GPT-4

Following the previous steps, we obtain question samples for each graph. However, all samples pertaining to the same template will share the same expressions. For example, inquiring about the price of an item will consistently yield the question “What is the price of the {item}?”. This limits the diversity of the data samples and may lead to a partially comprehensive evaluation.

To this end, we propose to use GPT-4 to paraphrase each question template into five different expressions so that we can have more diverse question samples regarding the same type of question. The prompts for paraphrasing can be found in Appendix C.

3.5 Automatic Answer Generation

The final step is to obtain the ground truth answer from the graph for each generated question. To achieve this goal, we first implement *graph functions* (e.g. neighbor check, degree check), which

can be utilized to reason on the graph. Then we implement *function chains* which can serve as a combination of graph functions in order to fetch the ground truth answer from the graph. The function chains are manually written by annotators for each type of question. Examples can be found in Appendix D.

4 Graph Chain-of-Thought

The straightforward solution to let LLMs interact with the graph is through retrieval-augmentation generation (RAG) (Lewis et al., 2020; Gao et al., 2023), where a retriever fetches related information from graphs as context for LLM generation. However, different from text corpus as the external knowledge source, the information in graphs also lies in the complex interconnection between the text units, which poses a potential requirement for traversing and reasoning on graphs. To enable LLMs to reason, Chain-of-thought (Wei et al., 2022) is proposed to encourage LLMs to decompose complex tasks into several steps. However, it is designed for reasoning on texts and leaves reasoning on graphs with LLMs an open question.

To this end, we design a simple solution named Graph Chain-of-Thought (GRAPH-COT) to tackle the complex graph reasoning problem with LLMs (shown in Figure 2). GRAPH-COT is an iterative framework, with three steps in each iteration: **reasoning**, **interaction**, and **execution**. We delve into each step as follows:

Reasoning with LLMs. Given the question or the previous iteration context, the first step is to let the LLMs conduct reasoning on what further external information from graphs is needed to answer the question, or if the question is answerable with the current contexts from graphs. For example, given the question “Who are the authors of Language Models are Unsupervised Multitask Learners?”. The LLMs are expected to reason “We need to first find the paper node {Language Models are Unsupervised Multitask Learners} on the graph.”

Interaction between LLMs and Graphs. Based on the output results from the previous LLM reasoning step, the next step is to let LLMs know how to interact with the graphs and fetch relevant information from the graphs. Inspired by (Yao et al., 2022), we pre-define four graph functions to cover both the semantic information and structure information on the graphs:




Figure 2: The workflow of GRAPH-COT, an iterative framework with three steps in each iteration: reasoning with LLMs, interaction between LLMs and graphs, and execution on graphs.

- *RetrieveNode*(Text): Identify related nodes in the graph with semantic search.
- *NodeFeature*(NodeID, FeatureName): Extract the textual feature information from the graph for a specific node.
- *NeighborCheck*(NodeID, NeighborType): Return the neighboring information in the graph for a specific node.
- *NodeDegree*(NodeID, NeighborType): Return the degree of a specific neighbor type for a specific node in the graph.

The task at hand requires LLMs to generate accurate graph function calls, based on their previous reasoning results, to effectively interact with the graph. In the given example, the LLMs are expected to generate “*RetrieveNode*(Language Models are Unsupervised Multitask Learners)”.

Execution on Graphs. The final step is to call those functions given by the previous step and fetch the relevant information from the graph. For the previous example, the graph will execute the *RetrieveNode*(\cdot) function and return “The ID of the most relevant paper node is p-4123”. Then, the process for the current iteration is over, and we start the new iteration from “reasoning with LLMs”. The whole framework will be iterated until the LLM finishes the reasoning and outputs the final answer. In this work, we enable LLMs to learn how to conduct GRAPH-COT with in-context learning (Dong et al., 2022). The prompts and demonstrations can be found in Appendix E.

Connection to LLM agents. It is worth mentioning that GRAPH-COT can be seen as an agent framework (Xi et al., 2023), where the LLM backbones are the agents and the graphs are the environments. The agents (LLMs) can interact with the

environment (graphs) with some predefined functions (defined in this section above). The goal of the agents is to explore the graph environment and conduct question-answering.

5 Experiments

5.1 Experimental Setup

Baselines. We compare our proposed GRAPH-COT with three types of baseline methods: standard LLMs (Base LLMs), text retrieval-augmented LLMs (Text RAG LLMs), and graph retrieval-augmented LLMs (Graph RAG LLMs):

- **Base LLMs:** We test if the LLMs can answer the given question with their knowledge without interacting with external data. We adopt the standard prompting, which involves providing simple instructions and letting LLMs generate an answer for the question.
- **Text RAG LLMs (Gao et al., 2023):** We treat the external graphs as pure text corpora and utilize a retriever to retrieve relevant text information from them. Subsequently, the retrieved text serves as context to augment the LLM for question answering.
- **Graph RAG LLMs:** This is an extension of text RAG, where not only the retrieved text/node but also the subgraph associated with it is linearized into a text sequence (Ye et al., 2023) and serves as the context. In the main result, we use 1-hop ego-graphs.

For all categories of baselines, we explore three LLM backbones, including LLaMA-2-13b-chat (Touvron et al., 2023), Mixtral-8x7b-Instruct (Jiang et al., 2024), and GPT-3.5-turbo (Ouyang et al., 2022). R-L

Table 2: Model performance on GRBENCH comparing standard LLMs, text retrieval augmented LLMs (Text RAG), graph retrieval augmented LLMs (Graph RAG), and GRAPH-COT. We showcase their performance based on Rouge-L (R-L) and GPT4score. We adopt GPT-3.5-turbo as the backbone for GRAPH-COT.

	Model	Academic		E-commerce		Literature		Healthcare		Legal	
		R-L	GPT4score	R-L	GPT4score	R-L	GPT4score	R-L	GPT4score	R-L	GPT4score
Base	LLaMA-2-13b-chat	8.13	8.03	7.01	12.00	5.32	20.83	5.25	13.70	15.97	16.11
	Mixtral-8x7b	9.02	8.14	12.54	18.00	7.50	22.50	3.88	20.00	12.74	16.11
	GPT-3.5-turbo	6.05	12.80	9.18	23.50	10.43	26.67	5.83	14.44	10.51	20.00
Text RAG	LLaMA-2-13b-chat	8.69	8.52	9.23	12.50	7.61	20.00	1.44	5.93	15.37	16.67
	Mixtral-8x7b	8.44	8.02	23.14	29.50	13.35	27.92	3.22	16.67	19.69	25.00
	GPT-3.5-turbo	5.83	9.91	14.06	20.00	10.04	20.83	4.57	8.52	18.14	23.89
Graph RAG	LLaMA-2-13b	22.01	22.97	12.48	20.00	9.25	20.00	2.97	4.81	17.98	17.22
	Mixtral-8x7b	27.77	31.20	32.87	37.00	20.08	33.33	8.66	15.19	23.48	25.56
	GPT-3.5-turbo	18.45	26.98	17.52	28.00	14.94	24.17	8.69	14.07	18.66	22.22
	GRAPH-COT	31.89	33.48	42.40	44.50	41.59	46.25	22.33	28.89	30.52	28.33

Evaluation Metrics. We use both rule-based metrics and model-based metrics to comprehensively evaluate the model results. For the former, we use Rouge-L(R-L), which measures the longest common subsequence of words between the responses and the ground truth answers. For the latter, we call GPT-4 to measure if the model output and ground truth are the same. We calculate the percentage of “correct” predicted by GPT-4 as GPT4score.

Implementation Settings. All experiments are conducted on NVIDIA GeForce RTX A6000 GPUs with Python 3.8 and Huggingface 4.36.2. We use Mpnet-v2⁷ as the retriever for all the baselines and our method and implement the indexing with FAISS (Johnson et al., 2019). In GRAPH-COT, we adopt GPT-3.5-turbo-16k (Jan 2024) as the backbone LLM in the main results and set the temperature t to 0 for consistent responses. We provide demonstrations for GRAPH-COT on how to conduct reasoning in Appendix E.

5.2 Overall Performance

The main results are shown in Table 2. From the results, we can find that: 1) GRAPH-COT outperforms all the baselines consistently and significantly. 2) Base LLMs are exhibiting fairly poor performance, typically because the LLMs may not contain the knowledge needed to answer those questions. 3) Graph RAG LLMs outperform text RAG LLMs in most cases since the former can provide more structure-aware context, which is helpful for problem-solving. 4) While GRAPH-COT performs the best, the absolute score is not high, leaving a great space to improve.

⁷<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>




Figure 3: Ablation study of GRAPH-COT. It performs well with in-domain demonstrations and remains generally robust to domain shifts in demonstrations.

5.3 Ablation Study

How Important are the Demonstrations for GRAPH-COT? To answer this question, we conduct experiments from two aspects: zero-shot study (no demonstrations) and cross-domain study (demonstrations from other domains (Ding et al., 2018)). The results are shown in Figure 3, where the columns and rows correspond to the source domain and target domain respectively. For the zero-shot study, no demonstrations are given (rightmost column in Figure 3). We empirically find that given no reasoning demonstrations, GRAPH-COT cannot work in all the datasets (nearly 0 performance). This implies that the LLMs suffer if given insufficient instructions (only graph definition and interaction function definitions). For the cross-domain study, we provide demonstrations from the source domain graphs and test on the target domain graphs. From the result (left five columns in Figure 3), in-domain demonstrations (diagonal) perform quite well and GRAPH-COT is overall robust to demonstration domain-shift. This observation

Table 3: Results of GRAPH-COT with different LLM backbones.

Model	GPT4score
GRAPH-COT	
w. LLaMA-2-13b-chat	16.04
w. Mixtral-8x7b	36.46
w. GPT-3.5-turbo	36.63
w. GPT-4	46.28

underscores the adaptability and effectiveness of GRAPH-COT in capturing the key steps of graph chain-reasoning through in-context learning, despite the diverse demonstration domains.

How Different LLMs Perform in GRAPH-COT?

In the main results, we adopt GPT-3.5-turbo as the LLM backbone for GRAPH-COT. In this section, we explore GRAPH-COT with other LLM backbones including LLaMA-2-13b-chat, Mixtral-8x7b-Instruct, GPT-3.5-turbo, and GPT-4. We randomly extract a subset from GRBENCH (one sample for each question template) to experiment and the results are shown in Table 3. From the result, we find that the LLM backbone matters. An LLM with more advanced instruction following ability and reasoning ability (*i.e.*, GPT-4) can contribute to better performance in GRAPH-COT.

5.4 RAG vs GRAPH-COT

Is the Retrieval-Augmented LLM a Good Choice on Graphs? We study how graph retrieval-augmented LLMs work by setting the retrieved subgraph to be just one node, 1-hop ego-graphs, and 2-hop ego-graphs. For all the settings, the ego-graphs are linearized into text sequences and serve as context. The averaged results over all the datasets are shown in Table 4. From the results, retrieving 1-hop ego-graph performs the best, but still underperforms GRAPH-COT. The reason is that when doing subgraph retrieval, the number of nodes/texts will grow exponentially as the hop number grows linearly. Even though the bigger the subgraph is, the more information it contains, a large-hop ego-graph will lead to a super long context which is even over the maximum input length of LLMs and will cause LLMs to lose in the middle. In this case, GRAPH-COT can serve as a better way to extract more useful information from the graph.

Table 4: Results of LLM with different retrieval-augmentation methods on GRBENCH.

Model	GPT4score
GPT-3.5-turbo	19.48
+ node retrieval	16.63
+ 1-hop subgraph retrieval	23.09
+ 2-hop subgraph retrieval	22.12
+ GRAPH-COT	36.29




Figure 4: Results of GRAPH-COT on samples of different difficulties in GRBENCH.

5.5 GRAPH-COT on Questions of Different Difficulty Levels in GRBENCH

In this section, we analyze the performance of GRAPH-COT on questions of different difficulty levels. The results are shown in Figure 4, where we find that GRAPH-COT performs relatively high on easy questions (the reasoning chains for those questions are simple) while having worse performance on medium/hard questions (complex/inductive reasoning).

5.6 Case Studies of GRAPH-COT

We conduct case studies to understand the weakness of GRAPH-COT. The results of two failure cases are shown in Figure 5. For the left case, we can find that despite using the most advanced LLM backbone (*i.e.*, GPT-4), the framework sometimes refers to the occurrence of the word rather than understanding its semantic meaning, leading to the wrong interaction function calls. For the right case, we can find that the framework sometimes misunderstands the structure of the graph, resulting in interaction failures.

Although GRAPH-COT achieves relatively good performance on GRBENCH, there is still quite

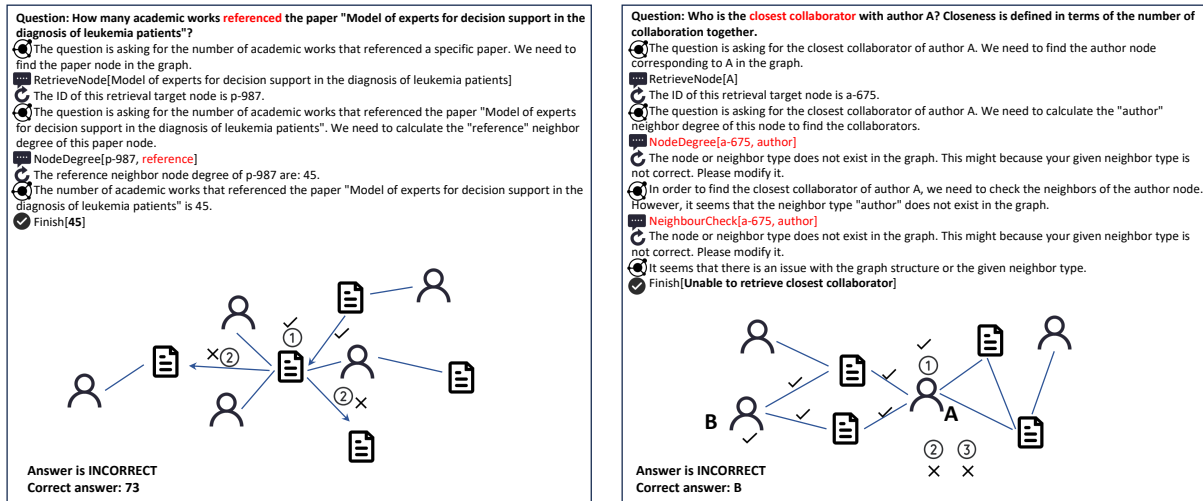


Figure 5: Failure cases of GRAPH-COT. The key information in the question and the wrong interaction of GRAPH-COT are colored in red. The author names in the second example are anonymized as A & B.

some room to improve. The two most promising directions to enhance LLMs’ graph reasoning abilities are exploring *how to let LLMs better understand the graphs* and *how to let LLMs conduct more complex reasoning*. For the former, in GRAPH-COT, we mainly use natural language to describe the graph for LLMs. Given that graphs are more structured rather than sequential, more structure-aware languages (e.g., graphXML (Herman and Marshall, 2000)) can be better choices. For the latter, given that reasoning problems on graphs are not only chain-reasoning problems, some more advanced reasoning paradigms such as tree-based reasoning (Yao et al., 2023) and graph-based reasoning (Besta et al., 2023) can be good directions.

6 Related Work

6.1 LLMs on graphs

Inspired by the recent success of LLMs on natural language processing tasks, researchers are exploring solving graph tasks with LLMs (Jin et al., 2023a). The main idea is to serve LLMs as the feature extractor (Chen et al., 2023) or final predictor (Jin et al., 2023b). For the former, many methods adopt a LLM-GNN cascaded structure (Chien et al., 2021), where LLMs extract node features for graph neural networks (GNNs) (Wu et al., 2020). For example, SimTeG (Duan et al., 2023) proposes to first warm up the LLM feature extractor before training the whole pipeline. GLEM (Zhao et al., 2022) introduces an iterative pipeline where GNNs can provide feedback for LLM feature extractors. For the latter, existing works transfer the structure in-

formation into a sequence to feed into LLMs (Tian et al., 2023; Xiong et al., 2024) or design advanced graph-empowered LLMs (Yang et al., 2021). For example, InstructGLM (Ye et al., 2023) utilizes natural language to describe graph structure. Heterformer (Jin et al., 2023c) proposes a graph-nested language model architecture. However, most existing works mainly focus on traditional graph tasks such as node classification (Xiao et al., 2022) and link prediction (Zhang and Chen, 2018). On the other hand, Graph-of-thought (Besta et al., 2023) proposes to conduct LLM reasoning with graph-structured thinking. Nevertheless, it mainly focuses on text-based reasoning rather than referring to external graphs. In our work, we research the question of augmenting LLMs with external graphs by conducting graph reasoning with LLMs.

6.2 Augmenting LLMs with external knowledge

Although LLMs (Touvron et al., 2023; Jiang et al., 2024) have shown their superb language understanding and generation capability (Zhao et al., 2023), they encounter issues with generating misleading information that seems credible but lacks factual basis, a phenomenon known as hallucination (Tonmoy et al., 2024; Rawte et al., 2023). To alleviate such an issue, existing works (Shuster et al., 2021) propose to augment LLMs with text corpora as external knowledge sources, with the retrieval-augmentation framework proposed (Lewis et al., 2020; Gao et al., 2023). Before LLMs’ inference, relevant text units are retrieved from the cor-

pora (Karpukhin et al., 2020) and serve as the context for LLMs to help reduce hallucination (Dong et al., 2022). Lewis et al. (2020) proposes to train the whole framework with a retriever and a generator end-to-end. Izacard and Grave (2020) introduces a fusion-in-decoder architecture to jointly consider all retrieved contexts in the generation. However, most existing works are designed to utilize external text corpora to augment LLMs. In our work, we explore how to augment LLMs with external text-attributed graphs and propose a benchmark for evaluation.

7 Conclusions

In this work, we study the problem of augmenting LLMs with (text-attributed) graphs as external knowledge sources. We first manually construct a benchmark dataset called GRBENCH, which contains 1,740 questions and 10 graphs from 5 domains. Each question in GRBENCH can be answered by referring to the graphs. We further propose a simple and effective framework called GRAPH-COT, which can augment LLMs with graphs by letting LLMs conduct iterative reasoning on graphs. GRAPH-COT contains three sub-steps in each iteration: LLM reasoning, LLM-graph interaction, and graph execution. We then conduct experiments with three backbone LLMs on GRBENCH and demonstrate the effectiveness of GRAPH-COT. Future works can explore how to let LLMs better understand the graphs and how to let LLMs conduct more complex reasoning.

Limitations

In this work, we mainly focus on augmenting LLMs with external graphs as knowledge sources by reasoning on the graphs, with a comprehensive benchmark dataset proposed. For GRBENCH construction, although we used GPT-4 to paraphrase the question templates, they are still mostly designed manually, so there might be room for improvement in terms of question diversity and difficulty. For GRAPH-COT, the LLM backbone used is an API model that cannot be fine-tuned (or is very costly to fine-tune). Future methods might need to consider how to train the LLMs explicitly to navigate on graphs.

Ethics Statement

Research has demonstrated the proficiency of Large Language Models (LLMs) (Touvron et al., 2023;

Jiang et al., 2024) in mastering language processing and generation. However, investigations have also pointed out their limitations, including social biases (Liang et al., 2021) and the propagation of false information (Abid et al., 2021). Our study aims to enhance LLMs by integrating external graphs as knowledge sources, proposing this approach as a potential solution to reduce bias and eradicate misinformation.

Acknowledgements

We thank Chen Yan (J.D.) for providing legal domain knowledge to help the authors construct the legal graph. Research was supported in part by US DARPA KAIROS Program No. FA8750-19-2-1004 and INCAS Program No. HR001121C0165, National Science Foundation IIS-19-56151, and the Molecule Maker Lab Institute: An AI Research Institutes program supported by NSF under Award No. 2019897, and the Institute for Geospatial Understanding through an Integrative Discovery Environment (I-GUIDE) by NSF under Award No. 2118329. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily represent the views, either expressed or implied, of DARPA or the U.S. Government. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Persistent anti-muslim bias in large language models. In *AIES*.
- Xiaomei Bai, Mengyang Wang, Ivan Lee, Zhuo Yang, Xiangjie Kong, and Feng Xia. 2019. Scientific paper recommendation: A survey. *Ieee Access*, 7:9324–9339.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, et al. 2023. Graph of thoughts: Solving elaborate problems with large language models. *arXiv preprint arXiv:2308.09687*.
- Si Chen, Pengfei Wang, Wei Fang, Xingchen Deng, and Feng Zhang. 2019. Learning to predict charges for judgment with legal graph. In *Artificial Neural Networks and Machine Learning–ICANN 2019: Text and Time Series: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part IV* 28, pages 240–252. Springer.

- Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Hongzhi Wen, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, Wenqi Fan, Hui Liu, et al. 2023. Exploring the potential of large language models (llms) in learning on graphs. *arXiv preprint arXiv:2307.03393*.
- Eli Chien, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Jiong Zhang, Olgica Milenkovic, and Inderjit S Dhillon. 2021. Node feature extraction by self-supervised multi-scale neighborhood prediction. *arXiv preprint arXiv:2111.00064*.
- Zhengming Ding, Sheng Li, Ming Shao, and Yun Fu. 2018. Graph adaptive knowledge transfer for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 37–52.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Keyu Duan, Qian Liu, Tat-Seng Chua, Shuicheng Yan, Wei Tsang Ooi, Qizhe Xie, and Junxian He. 2023. Simteg: A frustratingly simple approach improves textual graph learning. *arXiv preprint arXiv:2308.02565*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*, pages 507–517.
- Ivan Herman and M Scott Marshall. 2000. Graphxml—an xml-based graph description format. In *International Symposium on Graph Drawing*, pages 52–62. Springer.
- Daniel Scott Himmelstein, Antoine Lizee, Christine Hessler, Leo Brueggeman, Sabrina L Chen, Dexter Hadley, Ari Green, Pouya Khankhanian, and Sergio E Baranzini. 2017. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife*, 6:e26726.
- Gautier Izacard and Edouard Grave. 2020. [Leveraging passage retrieval with generative models for open domain question answering](#).
- Albert Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Bowen Jin, Gang Liu, Chi Han, Meng Jiang, Heng Ji, and Jiawei Han. 2023a. Large language models on graphs: A comprehensive survey. *arXiv preprint arXiv:2312.02783*.
- Bowen Jin, Wentao Zhang, Yu Zhang, Yu Meng, Xinyang Zhang, Qi Zhu, and Jiawei Han. 2023b. Patton: Language model pretraining on text-rich networks. *arXiv preprint arXiv:2305.12268*.
- Bowen Jin, Yu Zhang, Qi Zhu, and Jiawei Han. 2023c. Heterformer: Transformer-based deep node representation learning on heterogeneous text-rich networks. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1020–1031.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards understanding and mitigating social biases in language models. In *ICML*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Vipula Rawte, Amit Sheth, and Amitava Das. 2023. A survey of hallucination in large foundation models. *arXiv preprint arXiv:2309.05922*.
- Ali Sadeghian, Lakshman Sundaram, Daisy Zhe Wang, William F Hamilton, Karl Branting, and Craig Pfeifer. 2018. Automatic semantic edge labeling over legal citation graphs. *Artificial Intelligence and Law*, 26:127–144.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. 2008. Arnetminer: Extraction and mining of academic social networks. In *KDD'08*, pages 990–998.

- Yijun Tian, Huan Song, Zichen Wang, Haozhu Wang, Ziqing Hu, Fang Wang, Nitesh V Chawla, and Panpan Xu. 2023. Graph neural prompting with large language models. *arXiv preprint arXiv:2309.15427*.
- SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Mengting Wan and Julian McAuley. 2018. Item recommendation on monotonic behavior chains. In *Proceedings of the 12th ACM conference on recommender systems*, pages 86–94.
- Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul Kanakia. 2020. Microsoft academic graph: When experts are not enough. *Quantitative Science Studies*, 1(1):396–413.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Cheng Niu, Randy Zhong, Juntong Song, and Tong Zhang. 2023. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. *arXiv preprint arXiv:2401.00396*.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*.
- Shunxin Xiao, Shiping Wang, Yuanfei Dai, and Wenzhong Guo. 2022. Graph neural networks in node classification: survey and evaluation. *Machine Vision and Applications*, 33:1–19.
- Siheng Xiong, Ali Payani, Ramana Kompella, and Faramarz Fekri. 2024. Large language models can learn temporal reasoning. *arXiv preprint arXiv:2401.06853*.
- Junhan Yang, Zheng Liu, Shitao Xiao, Chaozhuo Li, Defu Lian, Sanjay Agrawal, Amit Singh, Guangzhong Sun, and Xing Xie. 2021. Graphformers: Gnn-nested transformers for representation learning on textual graph. *Advances in Neural Information Processing Systems*, 34:28798–28810.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.
- Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, and Yongfeng Zhang. 2023. Natural language is all a graph needs. *arXiv preprint arXiv:2308.07134*.
- Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. *Advances in neural information processing systems*, 31.
- Yu Zhang, Bowen Jin, Qi Zhu, Yu Meng, and Jiawei Han. 2023a. The effect of metadata on scientific literature tagging: A cross-field cross-model study. In *WWW’23*, pages 1626–1637.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023b. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- Jianan Zhao, Meng Qu, Chaozhuo Li, Hao Yan, Qian Liu, Rui Li, Xing Xie, and Jian Tang. 2022. Learning on large-scale text-attributed graphs via variational inference. *arXiv preprint arXiv:2210.14709*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

A Dataset

The detailed statistics of the graphs in GRBENCH are shown in Table 5. We will discuss the nodes' features and neighbors in each graph in the following paragraphs respectively.

Academic Graphs contain three types of nodes: paper, author, and venue. Here are examples to show the feature information and neighboring information for the three types of nodes respectively.

```
1 # paper node
2 {
3     'features': {
4         'title': ...,
5         'abstract': ...,
6         'keywords': [...],
7         'lang': ...,
8         'year': ...,
9     },
10    'neighbors': {
11        'author': [...],
12        'venue': [...],
13        'reference': [...],
14        'cited_by': [...],
15    }
16 }
17 # author node
18 {
19     'features': {
20         'name': ...,
21         'organization': ...,
22     },
23     'neighbors': {
24         'paper': [...],
25     }
26 }
27 # venue node
28 {
29     'features': {
30         'name': ...,
31     },
32     'neighbors': {
33         'paper': [...],
34     }
35 }
```

E-commerce Graph contains two types of nodes: item and brand. Here are examples to show the feature information and neighboring information for the two types of nodes respectively.

```
1 # item node
```

```
2 {
3     'features': {
4         'title': ...,
5         'description': ...,
6         'price': ...,
7         'category': [...],
8     },
9     'neighbors': {
10        'also_viewed_item': [...],
11        'buy_after_viewing_item':
12        [...],
13        'also_bought_item': [...],
14        'bought_together_item': [
15        ...],
16        'brand': [...],
17    }
18 }
19 # brand node
20 {
21     'features': {
22         'name': ...,
23     },
24     'neighbors': {
25         'item': [...],
26     }
27 }
```

Literature Graph contains four types of nodes: book, author, publisher, and series. Here are examples to show the feature information and neighboring information for the four types of nodes respectively.

```
1 # book node
2 {
3     'features': {
4         'country_code': ...,
5         'language_code': ...,
6         'is_ebook': ...,
7         'title': ...,
8         'description': ...,
9         'format': ...,
10        'num_pages': ...,
11        'publication_year': ...,
12        'popular_shelves': [...],
13        'genres': [...],
14    },
15    'neighbors': {
16        'author': [...],
17        'publisher': [...],
18    }
```



```

18     'series': [...],
19     'similar_books': [...],
20     }
21 }
22 # author node
23 {
24     'features': {
25         'name': ...,
26     },
27     'neighbors': {
28         'book': [...],
29     }
30 }
31 # publisher node
32 {
33     'features': {
34         'name': ...,
35     },
36     'neighbors': {
37         'book': [...],
38     }
39 }
40 # series node
41 {
42     'features': {
43         'title': ...,
44         'description': ...,
45     },
46     'neighbors': {
47         'book': [...],
48     }
49 }

```

Healthcare Graph contains eleven types of nodes: anatomy, biological process, cellular component, compound, disease, gene, molecular function, pathway, pharmacologic class, side effect, and symptom. Here are examples to show the feature information and neighboring information for the eleven types of nodes respectively.

```

1 # anatomy node
2 {
3     'features': {
4         'name': ...,
5     },
6     'neighbors': {
7         'Anatomy-expresses-Gene':
8         [...],
9     }
10 # biological process node

```

```

11 {
12     'features': {
13         'name': ...,
14     },
15     'neighbors': {
16         'Gene-participates-
17         Biological Process': [...],
18     }
19 # cellular component node
20 {
21     'features': {
22         'name': ...,
23     },
24     'neighbors': {
25         'Gene-participates-
26         Cellular Component': [...],
27     }
28 # compound node
29 {
30     'features': {
31         'name': ...,
32     },
33     'neighbors': {
34         'Compound-causes-Side
35         Effect': [...],
36         'Compound-upregulates-
37         Gene': [...],
38         'Compound-downregulates-
39         Gene': [...],
40     }
41 # disease node
42 {
43     'features': {
44         'name': ...,
45     },
46     'neighbors': {
47         'Disease-associates-Gene':
48         [...],
49         'Disease-localizes-
50         Anatomy': [...],
51         'Compound-treats-Disease':
52         [...],
53         'Disease-resembles-
54         Disease': [...],
55         'Disease-presents-Symptom':
56         [...],
57         'Disease-upregulates-Gene':
58         [...],
59     }

```

```

52 }
53 # gene node
54 {
55     'features': {
56         'name': ...,
57     },
58     'neighbors': {
59         'Gene-participates-
Biological Process': [...],
60         'Anatomy-upregulates-Gene
': [...],
61         'Anatomy-expresses-Gene':
[...],
62         'Anatomy-downregulates-
Gene': [...],
63         'Compound-upregulates-
Gene': [...],
64         'Gene-interacts-Gene': [
...],
65         'Gene-participates-
Molecular Function': [...],
66         'Gene-participates-
Cellular Component': [...],
67     }
68 }
69 # molecular function node
70 {
71     'features': {
72         'name': ...,
73     },
74     'neighbors': {
75         'Gene-participates-
Molecular Function': [...],
76     }
77 }
78 # pathway node
79 {
80     'features': {
81         'name': ...,
82     },
83     'neighbors': {
84         'Gene-participates-
Pathway': [...],
85     }
86 }
87 # pharmacologic node
88 {
89     'features': {
90         'name': ...,
91     },
92     'neighbors': {
93         'Pharmacologic Class-
includes-Compound': [...],
94     }
95 }
96 # side effect node
97 {
98     'features': {
99         'name': ...,
100     },
101     'neighbors': {
102         'Compound-causes-Side
Effect': [...],
103     }
104 }
105 # symptom node
106 {
107     'features': {
108         'name': ...,
109     },
110     'neighbors': {
111         'Disease-presents-Symptom
': [...],
112     }
113 }

```

Legal Graph contains four types of nodes: opinion, opinion cluster, docket, and court. Here are examples to show the feature information and neighboring information for the four types of nodes respectively.

```

1 # opinion node
2 {
3     'features': {
4         'plain_text': ...,
5     },
6     'neighbors': {
7         'opinion_cluster': [...],
8         'reference': [...],
9         'cited_by': [...],
10    }
11 }
12 # opinion cluster node
13 {
14     'features': {
15         'judges': ...,
16         'case_name': ...,
17         'attorneys': ...,
18         'syllabus': ...,
19     },
20     'neighbors': {
21         'opinion': [...],
22         'docket': [...],

```

```

23     }
24 }
25 # docket node
26 {
27     'features': {
28         'case_name': ...,
29         'pacer_case_id': ...,
30     },
31     'neighbors': {
32         'opinion_cluster': [...],
33         'court': [...],
34     }
35 }
36 # court node
37 {
38     'features': {
39         'citation_string': ...,
40         'full_name': ...,
41         'start_date': ...,
42         'end_date': ...,
43     },
44     'neighbors': {
45         'docket': [...],
46     }
47 }

```

- Who collaborate with author {author_name} in {org_name} to write paper "{paper_title}"?
- Who wrote both the paper "{paper1_title}" and paper "{paper2_title}"?
- Who is the closest collaborator with author {author_name} in {org_name}? Closeness is defined in terms of the number of collaborations together.
- How many collaborators does author {author_name} in {org_name} have in year?
- How many papers did {author_name1} in {org_name1} and {author_name2} in {org_name2} write together?
- Which venue did {author_name1} in {org_name1} and {author_name2} in {org_name2} collaborate most?
- How many people does author {author_name1} in {org_name1} need to know at least to know author {author_name2} in {org_name2}?"
- What is the research interests (top 3 keywords) of author {author_name} in {org_name}?

B Question Templates

In this section, we show the templates for easy, medium, and hard questions.

Academic Graphs.

- Easy

- Who are the authors of paper "{paper_title}"?
- What organization is researcher {author_name} affiliated with?
- Where is the paper "{paper_title}" published?
- How many papers cite the paper "{paper_title}"?
- How many papers does paper "{paper_title}" cite?
- Which is the most cited paper by author {author_name} in {org_name}?
- How many papers did author {author_name} in {org_name} write?

- Medium

- Hard

- Which paper should be recommended to the reader of paper {paper1_title}? Please select from the candidate list {paper2_title}, {paper3_title}, {paper4_title}, {paper5_title}, {paper6_title}, {paper7_title}, {paper8_title}, {paper9_title}, {paper10_title}, {paper11_title}. Please answer the paper title rather than ID."

E-commerce Graph.

- Easy

- What is the brand of item {item_title}?
- What is the category of item {item_title}?
- What is the price of item {item_title}?

- Medium

- How many co-viewed items does item {item_title} have?
- How many bought-together items does item {item_title} have?
- How many buy-after-viewing items does item {item_title} have?

Table 5: Detailed Dataset Statistics of GRBENCH.

Domain	Topic	Graph Statistics		Data	
		Nodes	Edges	# Templates	# Questions
Academic	CS	Paper (~5M) Author (~2M) Venue (~55K)	Written-by (~14M) Publish-in (~5M) Cited-by (~32M)	15	150
	Biology	Paper (~1M) Author (~2M) Venue (100)	written-by (~8M) publish-in (~1M) cited-by (~29M)	14	140
	Chemistry	Paper (~1M) Author (~2M) Venue (100)	written-by (~7M) publish-in (~1M) cited-by (~20M)	14	140
	Material Science	Paper (~1M) Author (~1M) Venue (99)	written-by (~6M) publish-in (~1M) cited-by (~14M)	14	140
	Medicine	Paper (~2M) Author (~4M) Venue (100)	written-by (~14M) publish-in (~2M) cited-by (~12M)	14	140
	Physics	Paper (~1M) Author (~1M) Venue (91)	written-by (~13M) publish-in (~1M) cited-by (~18M)	14	140
E-commerce	Amazon	Item (~9M) Brand (~110K)	also-viewed (~125M) buy-after-viewing (~9M) also-bought (~170M) bought-together (~6M) item-brand (~1M)	20	200
Literature	Goodreads	Book (~2M) Author (~829K) Publisher (~193K) Series (~400K)	written-by (~3M) published-in (~1M) book-series (~822K) similar-book (~16M)	24	240
Healthcare	Disease	11 nodes types See Sec A	24 edge types See Sec A	27	270
Legal	Freelaw	Opinion (~9M) Opinion-cluster (~8M) Docket (~66M) Court (~3K)	opinion-cluster (~9M) opinion-citation (~29M) cluster-docket (~8M) docket-court (~66M)	18	180
SUM	-	-	-	174	1740

- How many also-bought items does item {item_title} have?
- How many items are in brand {brand_name}?
- Find the items which are in the same brand and same category as item {item_title}.
- Which item shares over {num} co-viewed items with item {item_title}?
- Which item shares over {num} bought-together items with item {item_title}?
- How many items have the same bought-together items with item {item_title}?
- What is the average price of the bought-together/co-viewed items with {item_title}?

- What is the most popular category name of the bought-together/co-viewed items with {item_title}?

• Hard

- What next item should be recommended to the user based on his history: {item_titles}?
- What is the exact matched item given this query: {query_text}?
- What is the substitutive item given this query: {query_text}?
- What is the complementary item given this query: {query_text}?

Literature Graph.

• Easy

- Who are the authors of book {book title}?
- What is the publisher of book {book title}?
- Which shelves do we need to put book {book title} on?
- What genre does the book {book title} belong to?
- In which series is the book {book title} included?
- What is the publication year of book {book title}?
- How many pages does the book {book title} have?
- Is the book {book title} an eBook?
- What language is the book {book title} written in?
- How many books has author {author name} written?
- How many similar books does Book {book title} have?
- How many books does publisher {publisher name} publish?
- How many books are part of the series {series title}?

• **Medium**

- Find the book written by the same author and published by the same publisher as book {book title}.
- Find books by the same author and share similar genre with book {book title}.
- Find the earliest book written by the author of the book {book title}.
- Find the series in which the same author as the book {book title} has contributed, but the series is different from the book's series.
- How many authors have collaborated with the publisher {publisher name}?
- Which author has the most published books that have the same genre as the book {book title}?
- What is the most common publication format of books by author {author name}?
- What is the most frequent genre in the works of the author {author name}?
- Which publisher has released the majority of books in the genre {genre name}?

- What is the most common language among the books written by author {author name}?

• **Hard**

- What book should be recommended to the user based on his history: {book titles}?

Healthcare Graph.

• **Easy**

- What are the side effects of compound {compound name}?
- What are the symptoms of the disease {disease name}?
- What are the biological processes of gene {gene name}?
- What are the molecular functions of gene {gene name}?
- What anatomy can be downregulated by gene {gene name}?
- What anatomy can be expressed by gene {gene name}?
- What anatomy can be upregulated by gene {gene name}?
- How many resemble compounds do {compound name} have?
- How many resemble disease do {disease name} have?
- How many compounds can be used to treat {disease name}?

• **Medium**

- What compound can treat both {disease name1} and {disease name2}?
- What disease located in {anatomy name} can {compound name} palliate?
- What disease located in {anatomy name} can {compound name} treat?
- What disease is downregulated by {gene name} and located in {anatomy name}?
- What disease is associated by {gene name} and located in {anatomy name}?
- What disease is upregulated by {gene name} and located in {anatomy name}?
- Is there a correlation between {gene name} and {symptom name}? Please answer True or False

- Which pharmacologic class includes the most compounds that can palliate the disease with {symptom name}?
- Which pharmacologic class includes the most compounds that can treat the disease with {symptom name}?
- Which cellular component is participated by most genes that are upregulated in disease with {symptom name}?
- Which cellular component is participated by most genes that are associated in disease with {symptom name}?
- Which cellular component is participated by most genes that are downregulated in disease with {symptom name}?
- Which pathway is participated by most genes that are upregulated in disease with {symptom name}?
- Which pathway is participated by most genes that are associated in disease with {symptom name}?
- Which pathway is participated by most genes that are downregulated in disease with {symptom name}?
- How many genes participate the exact same biological processes with {gene name}?
- How many diseases present the exact same symptoms with {disease name}?
- Which members of the judiciary are responsible for the group of rulings that includes the following opinion: {opinion text}
- What docket includes this opinion: {opinion plain text}? Please answer with the pacer case ID.
- Which court is this opinion cluster syllabus published: {opinion cluster text}?
- How many times has the case {case name} been judged in different courts?
- How many opinions are contained in the opinion clusters about {case name}?
- How many opinions are contained in the opinion cluster with syllabus: {opinion cluster text}?
- How many opinions are contained in the opinion cluster with opinion {opinion text}?
- Which court is this opinion ({opinion text}) published?
- What is the preferred court to cite of judges in court {source court name}?

• **Hard**

- Is the given sentence supported by the given case? Sentence: text, case: {case name}.
- Find a case which can support this sentence: {text}.

Legal Graph.

• **Easy**

- what is the start date of court {court name}?
- what is the end date of court {court name}?
- what is the citation string of court {court name}?
- which court is handling the case listed under the PACER docket number {pacer id}?
- Who are the attorneys for the case corresponding to this opinion cluster: {opinion cluster text}?
- How many dockets have been processed in court {court name}?
- How many opinions are citing this opinion: {opinion text}?

• **Medium**

C Question Template Paraphrase Prompt

Paraphrase Prompt

Paraphrase the given template in four different ways. Keep the name in “ unchanged, don’t use ’ in question, and use the same format (‘question string’, ‘answer string’):

D Programmatic Automatic Answer Generation Examples

```

1 # Define graph walking functions
2 def one_hop(graph, center_node_type,
3             center_node_save_key,
4             neighbor_node_type,
5             neighbor_node_save_key, edge_type, k
6             ):
7     generated_data = []
8     cnt = 0
9     center_ids = list(graph[
10        center_node_type].keys())
11     random.shuffle(center_ids)

```

```

7   for center_id in center_ids:
8       center_name = graph[
center_node_type][center_id][
features']['name']
9       if edge_type not in graph[
center_node_type][center_id][
neighbors']:
10          continue
11          neighbor_ids = graph[
center_node_type][center_id][
neighbors][edge_type]
12          neighbor_names = [graph[
neighbor_node_type][neighbor_id][
features']['name'] for neighbor_id
in neighbor_ids]
13          if len(neighbor_names) > 5:
14              continue
15              generated_data.append({
center_node_save_key:center_name,
neighbor_node_save_key: ', '.join(
neighbor_names)})
16              cnt += 1
17              if cnt == k:
18                  break
19          return generated_data
20
21 # Generate examples
22 random.seed(2023)
23 question = "what are the side effects of
compound {compound_name}?"
24 answer = "{side_effects}"
25 generated_data = one_hop(graph, '
Compound_nodes', 'compound_name', '
Side_Effect_nodes', 'side_effects',
'Compound-causes-Side Effect', k)
26 assert len(generated_data) == k
27 all_generated_data[(question, answer)] =
generated_data

```

E Prompts in GRAPH-COT

The prompt to instruct LLMs for GRAPH-COT contains three parts: graph description, interaction function description, and demonstrations. The final prompt is shown below, where “graph definition”, “interaction function descriptions” and “examples” correspond to the three parts respectively:

GRAPH-COT prompt

Solve a question answering task with interleaving Thought, Interaction with Graph, Feedback from Graph steps.

In Thought step, you can think about what further information is needed, and In Interaction step, you can get feedback from graphs with four functions:

{interaction function descriptions}

You may take as many steps as necessary.

Here are some examples:

{examples}

(END OF EXAMPLES)

Definition of the graph: {graph definition}

Question: {question}

Please answer by providing node main feature (e.g., names) rather than node IDs.

E.1 Graph Description Prompts

MAG graph descriptions

There are three types of nodes in the graph: paper, author and venue. Paper nodes have features: title, abstract, year and label. Author nodes have features: name. Venue nodes have features: name. Paper nodes are linked to author nodes, venue nodes, reference nodes and cited by nodes. Author nodes are linked to paper nodes. Venue nodes are linked to paper nodes.

DBLP graph descriptions

There are three types of nodes in the graph: paper, author and venue. Paper nodes have features: title, abstract, keywords, lang, and year. Author nodes have features: name and organization. Venue nodes have features: name. Paper nodes are linked to their author nodes, venue nodes, reference nodes (the papers this paper cite) and cited by nodes (other papers which cite this paper). Author nodes are linked to their paper nodes. Venue nodes are linked to their paper nodes.

E-commerce graph descriptions

There are two types of nodes in the graph: item and brand. Item nodes have features: title, description, price, img, category. Brand nodes have features: name. Item nodes are linked to their brand nodes, also viewed item nodes, buy after viewing item nodes, also bought item nodes, bought together item nodes. Brand nodes are linked to their item nodes.

Healthcare graph descriptions

There are eleven types of nodes in the graph: Anatomy, Biological Process, Cellular Component, Compound, Disease, Gene, Molecular Function, Pathway, Pharmacologic Class, Side Effect, Symptom. Each node has name feature. There are these types of edges: Anatomy-downregulates-Gene, Anatomy-expresses-Gene, Anatomy-upregulates-Gene, Compound-binds-Gene, Compound-causes-Side Effect, Compound-downregulates-Gene, Compound-palliates-Disease, Compound-resembles-Compound, Compound-treats-Disease, Compound-upregulates-Gene, Disease-associates-Gene, Disease-downregulates-Gene, Disease-localizes-Anatomy, Disease-presents-Symptom, Disease-resembles-Disease, Disease-upregulates-Gene, Gene-covaries-Gene, Gene-interacts-Gene, Gene-participates-Biological Process, Gene-participates-Cellular Component, Gene-participates-Molecular Function, Gene-participates-Pathway, Gene-regulates-Gene, Pharmacologic Class-includes-Compound.

Literature graph descriptions

There are four types of nodes in the graph: book, author, publisher, and series. Book nodes have features: country code, language code, is ebook, title, description, format, num pages, publication year, url, popular shelves, and genres. Author nodes have features: name. Publisher nodes have features: name. Series nodes have features: title and description. Book nodes are linked to their author nodes, publisher nodes, series nodes and similar books nodes. Author nodes are linked to their book nodes. Publisher nodes are linked to their book nodes. Series nodes are linked to their book nodes.

Legal graph descriptions

There are four types of nodes in the graph: opinion, opinion cluster, docket, and court. Opinion nodes have features: plain text. Opinion cluster nodes have features: syllabus, judges, case name, attorneys. Docket nodes have features: pacer case id, case name. Court nodes have features: full name, start date, end date, citation string. Opinion nodes are linked to their reference nodes and cited by nodes, as well as their opinion cluster nodes. Opinion cluster nodes are linked to opinion nodes and docket nodes. Docket nodes are linked to opinion cluster nodes and court nodes. Court nodes are linked to docket nodes.

E.2 Interaction Function Description Prompts

Interaction function descriptions

- (1) RetrieveNode[keyword], which retrieves the related node from the graph according to the corresponding query.
- (2) NodeFeature[Node, feature], which returns the detailed attribute information of Node regarding the given "feature" key.
- (3) NodeDegree[Node, neighbor type], which calculates the number of "neighbor type" neighbors of the node Node in the graph.
- (4) NeighbourCheck[Node, neighbor type], which lists the "neighbor type" neighbours of the node Node in the graph and returns them.

E.3 Demonstrations

In GRAPH-CoT, we provide three demonstrations to teach LLMs how to utilize the four interaction functions. The demonstrations for academic domain graphs are shown in Figure 6. More detailed information about demonstrations for other domain graphs can be found at https://github.com/PeterGriffinJin/Graph-CoT/blob/main/Graph-CoT/code/graph_fewshots.py.

Academic graph demonstrations

Definition of the graph: {academic graph definition}

Question: When was the paper Strongly Interacting Higgs Sector in the Minimal Standard Model published?

Reasoning 1: The question is asking some basic information of a node (Strongly Interacting Higgs Sector in the Minimal Standard Model). We need to find the node in the graph.

Interaction 1: RetrieveNode[Strongly Interacting Higgs Sector in the Minimal Standard Model]

Execution 1: The ID of this node is 3101448248.

Reasoning 2: The question is asking the published date of a paper, we need to check the node feature (year) from the graph.

Interaction 2: NodeFeature[3101448248, year]

Execution 2: 1993

Reasoning 3: The published date of the paper is 1993.

Interaction 3: Finish[1993]

Definition of the graph: {academic graph definition}

Question: How many authors do the paper Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars have?

Reasoning 1: The question is asking information of a node (Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars). We need to find the node in the graph.

Interaction 1: RetrieveNode[Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars]

Execution 1: The ID of this node is 2090642949.

Reasoning 2: The question is asking the number of authors of a paper, we need to calculate the node's author neighbor degree from the graph.

Interaction 2: NodeDegree[2090642949, author]

Execution 2: 2

Reasoning 3: The number of the authors is 2

Interaction 3: Finish[2]

Definition of the graph: {academic graph definition}

Question: What was the publish venue of the paper Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars?

Reasoning 1: The question is asking information of a node (Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars). We need to find the node in the graph.

Interaction 1: RetrieveNode[Mass Accretion Rates in Self-Regulated Disks of T Tauri Stars]

Execution 1: The ID of this node is 2090642949.

Reasoning 2: The question is asking the published venue of a paper, we need to check the node's venue neighbor from the graph.

Interaction 2: NeighbourCheck[2090642949, venue]

Execution 2: ['1980519', '1053242']

Reasoning 3: The ID of the published venue are 1980519 and 1053242. We need to get their names.

Interaction 3: NodeFeature[1980519, name], NodeFeature[1053242, name]

Execution 3: the astrophysical journal, the atmosphere journal

Reasoning 4: The name of the published venues are the astrophysical journal and the atmosphere journal

Interaction 4: Finish[the astrophysical journal, the atmosphere journal]

Figure 6: Demonstrations for the academic domain graphs