

# A Survey of Graph Retrieval-Augmented Generation for Customized Large Language Models

Qinggong Zhang<sup>†\*</sup>, Shengyuan Chen<sup>†\*</sup>, Yuanchen Bei<sup>†\*</sup>, Zheng Yuan<sup>†</sup>  
Huachi Zhou<sup>†</sup>, Zijin Hong<sup>†</sup>, Junnan Dong<sup>†</sup>, Hao Chen<sup>†</sup>, Yi Chang<sup>‡</sup>, Xiao Huang<sup>†</sup>

<sup>†</sup>The Hong Kong Polytechnic University, Hong Kong SAR, China

<sup>‡</sup>Jilin University, Changchun, China

{qinggangg.zhang, yzheng.yuan, huachi.zhou, zijin.hong, hanson.dong}@connect.polyu.hk

{shengyuan-chen, xiao.huang}@polyu.edu.hk

{iyuanchenbei, sundaychenhao}@gmail.com, yichang@jlu.edu.cn

**Abstract**—Large language models (LLMs) have demonstrated remarkable capabilities in a wide range of tasks, yet their application to specialized domains remains challenging due to the need for deep expertise. Retrieval-Augmented generation (RAG) has emerged as a promising solution to customize LLMs for professional fields by seamlessly integrating external knowledge bases, enabling real-time access to domain-specific expertise during inference. Despite its potential, traditional RAG systems, based on flat text retrieval, face three critical challenges: (i) complex query understanding in professional contexts, (ii) difficulties in knowledge integration across distributed sources, and (iii) system efficiency bottlenecks at scale. This survey presents a systematic analysis of Graph-based Retrieval-Augmented Generation (GraphRAG), a new paradigm that revolutionizes domain-specific LLM applications. GraphRAG addresses traditional RAG limitations through three key innovations: (i) graph-structured knowledge representation that explicitly captures entity relationships and domain hierarchies, (ii) efficient graph-based retrieval techniques that enable context-preserving knowledge retrieval with multihop reasoning ability, and (iii) structure-aware knowledge integration algorithms that leverage retrieved knowledge for accurate and logical coherent generation of LLMs. In this survey, we systematically analyze the technical foundations of GraphRAG and examine current implementations across various professional domains, identifying key technical challenges and promising research directions. All the related resources of GraphRAG, including research papers, open-source data, and projects, are collected for the community in <https://github.com/DEEP-PolyU/Awesome-GraphRAG>.

**Index Terms**—Retrieval-Augmented Generation, Knowledge Graphs, Large Language Models, GraphRAG

## I. INTRODUCTION

Large language models (LLMs), like the GPT series [1], have surprised the world with their remarkable ability across a wide range of tasks, achieving breakthroughs in text comprehension [2], question answering [3], and content generation [4]–[6]. Despite their effectiveness, LLMs are always criticized for their limited ability to handle knowledge-intensive tasks, especially when faced with questions requiring domain expertise [7]. Specifically, the application of LLMs to specialized domains remains challenging for three fundamental reasons, including ❶ Knowledge limitations: LLMs’ pre-trained knowledge is broad but shallow in specialized fields. Their training data primarily

consists of general-domain content, leading to insufficient depth in professional domains and potential inconsistencies with current domain-specific standards and practices. ❷ Reasoning complexity: Specialized domains require precise, multi-step reasoning with domain-specific rules and constraints. LLMs often struggle to maintain logical consistency and professional accuracy throughout extended reasoning chains, particularly when dealing with technical constraints or domain-specific protocols. ❸ Context sensitivity: Professional fields often involve context-dependent interpretations where the same terms or concepts may have different meanings or implications based on specific circumstances. LLMs frequently fail to capture these nuanced contextual variations, leading to potential misinterpretations or inappropriate generalizations.

To adapt LLMs for specific or private domains, initial strategies involved **fine-tuning LLMs** with specialized datasets [8]. This method enhances performance by adding a limited number of parameters while fixing the parameters learned in the pre-training [9]. However, the significant distribution gap between the domain-specific dataset and the pre-training corpus makes it challenging for LLMs to integrate new knowledge without compromising their existing understanding [10]. A recent study by Google Research further highlighted the risks associated with using supervised fine-tuning to update knowledge, particularly in cases where new knowledge conflicts with pre-existing information; acquiring new knowledge through supervised fine-tuning can lead to the model generating new hallucinations and even experiencing severe catastrophic forgetting [11].

**Retrieval-Augmented generation (RAG)** offers a promising solution to customize LLMs for specific domains [12]. Rather than retraining LLMs to incorporate updates, RAG enhances these models by leveraging external knowledge from text corpora without modifying their architecture or parameters. This approach enables LLMs to generate responses by leveraging not only their pre-trained knowledge but also real-time retrieved domain-specific information, thereby providing more accurate and reliable answers. The naive RAG systems operate through three key steps: knowledge preparation, retrieval, and integration. During knowledge preparation, external sources such as documents, databases, or webpages are divided into manageable textual chunks and converted into vector representations for efficient indexing. In the retrieval stage, when a user submits a

\*Authors contributed equally to this research.




Fig. 1: The development trends in the field of GraphRAG with representative works.

query, the system searches for relevant chunks using keyword matching or vector similarity measures. The integration stage then combines these retrieved chunks with the original query to create an informed prompt for the LLM's response. Recently, some advanced RAG systems have evolved beyond simple text chunk retrieval to offer more sophisticated knowledge augmentation approaches. These include hierarchical RAG that preserves document structure through multi-level retrieval [13], [14], re-ranking systems that implement two-stage retrieval for higher recall and precision [15], [16], self-querying RAG that automatically decomposes complex queries [17], and adaptive RAG that dynamically adjusts retrieval strategies based on query types [18], [19]. These advanced RAG systems aim to address the limitations of naive RAG approaches by improving context awareness, retrieval accuracy, and handling complex queries more effectively.

The emergence of RAG has offered a promising approach for customizing LLMs with domain-specific knowledge. However, despite its potential, RAG faces several critical limitations that impact its effectiveness in practical applications. These limitations can be broadly categorized into four main challenges that significantly affect the performance and utility of RAG-enhanced LLMs. The primary challenge lies in **complex query understanding**. Specialized domains often involve intricate terminology and industry-specific jargon that requires precise interpretation [20]. User queries in these areas typically contain numerous technical terms and industry-specific expressions, with solutions often requiring reasoning across multiple related concepts. Traditional RAG approaches, which rely on simple keyword matching and vector similarity techniques, are inadequate for capturing the deep semantic nuances and multi-step reasoning processes necessary for accurate and comprehensive [21]. For instance, when queried about the connection between concept A and concept D, these systems

typically retrieve only directly related information, missing crucial intermediate concepts like B and C that could bridge the relationship. This narrow retrieval scope limits RAG's ability to a broad contextual understanding and complex reasoning.

Another key challenge involves **integrating domain knowledge from distributed sources**. Domain knowledge is usually collected from different sources, such as textbooks, research papers, industry reports, technical manuals, and maintenance logs. These textual documents may have varying levels of quality, accuracy, and completeness. The retrieved knowledge is often flattened, extensive, and intricate, while domain concepts are typically scattered across multiple documents without clear hierarchical relationships between different concepts [7], [22], [23]. Although RAG systems attempt to manage this complexity by dividing documents into smaller chunks for effective and efficient indexing, this approach inadvertently sacrifices crucial contextual information, significantly compromising retrieval accuracy and contextual comprehension. This limitation hampers the ability to establish robust connections between related knowledge points, leading to fragmented understanding and reduced efficacy in leveraging domain-specific expertise.

The third limitation stems from the **inherent constraints of LLMs**. While RAG systems can retrieve relevant information from vast knowledge bases, the LLM's ability to process this information is constrained by its fixed context window (typically 2K-32K tokens) [1], [24]. Long-range dependencies in complex documents cannot be fully captured, as the content exceeding the context window must be truncated or summarized, disrupting natural semantic units and logical flow. The challenge of maintaining coherence across extensive knowledge contexts becomes increasingly problematic in professional domains, as critical information may be lost during context window truncation. This fundamental limitation directly impacts the system's ability to process and synthesize comprehensive

information from large-scale knowledge bases.

The last challenge relates to system **efficiency and scalability**. The entire RAG pipeline - from initial corpus preprocessing and indexing to real-time retrieval and generation - faces significant efficiency bottlenecks [25], [26]. The external knowledge base contains a lot of domain-irrelevant information, while domain-specific terminologies are always sparsely distributed over these documents. RAG systems can be computationally expensive and time-consuming [25], especially when dealing with large-scale knowledge sources, as the model needs to search through vast amounts of unstructured text to find relevant information. Moreover, real-time retrieval and cross-document reasoning can introduce considerable latency, negatively impacting user experience. The scalability of RAG is further constrained by declining retrieval quality and accuracy as the size of the knowledge base grows [26], thereby limiting its practical deployment in extensive and dynamic professional environments.

To address these limitations, **graph retrieval-augmented generation (GraphRAG)** has recently emerged as a new paradigm to customize LLMs with well-organized background knowledge and improved contextual reasoning [25], [27]–[29]. Based on the utilization of graph structures, existing GraphRAG models can be categorized into three main categories: ① **Knowledge-based GraphRAG**, which uses graphs as knowledge carriers. ② **Index-based GraphRAG** uses graphs as index tools to retrieve relevant raw texts from the corpus, and ③ **Hybrid GraphRAG** which combines the strengths of both knowledge-based and index-based frameworks, providing more advanced solutions for complex reasoning tasks. Knowledge-based and Index-based GraphRAG represent two distinct paradigms for enhancing LLMs with structured knowledge. Knowledge-based GraphRAG focuses on transforming unstructured textual documents into explicit and structured KGs, where nodes represent domain concepts and edges capture semantic relationships between them, enabling better representation of hierarchical relationships and complex knowledge dependencies. In contrast, Index-based GraphRAG maintains the original textual form while utilizing graph structures primarily as an indexing mechanism to organize and retrieve relevant text chunks efficiently. By incorporating graph structures into text indexing, Index-based GraphRAG methods establish semantic connections between text chunks for efficient look-up operations and retrieval. While Knowledge-based GraphRAG emphasizes the explicit modeling of domain knowledge and semantic relationships through graph transformation, Index-based GraphRAG prioritizes efficient information retrieval and global navigation through the graph-based organization of the raw text. This fundamental difference in approach reflects their distinct purposes: Knowledge-based GraphRAG aims to create a structured knowledge representation for a better understanding of complex relationships with graph-based reasoning capability, whereas Index-based GraphRAG focuses on optimizing the retrieval and accessibility of relevant textual information through graph-based indexing strategies.

In this survey, we systematically analyze the technical foundations of GraphRAG and examine current implementations across various professional domains, identifying key technical

challenges and promising research directions. All the related resources of GraphRAG, including research papers, open-source data, and projects, are collected for the community in <https://github.com/DEEP-PolyU/Awesome-GraphRAG>.

## II. OVERVIEW

This survey provides a comprehensive analysis of GraphRAG, detailing its taxonomy, mechanisms, challenges, and future research directions, which is organized into seven main sections that progress from foundational concepts to practical implementations. Specifically, we begin in Section 2 II by establishing the foundational framework, tracing GraphRAG’s evolution from traditional RAG systems, examining RAG’s limitations in handling structured knowledge, and introducing GraphRAG’s core concepts and advantages in complex reasoning tasks. The following three sections systematically explore the key components of GraphRAG systems: two primary paradigms for structured knowledge organization, including knowledge carrier graphs and index graphs (Section 3 IV), retrieval techniques for extracting query-relevant factual information from structured knowledge bases (Section 4 V), and knowledge integration methods for effectively incorporating retrieved knowledge into LLMs (Section 5 VI). Moving toward practical applications, Section 6 VIII addresses implementation aspects by providing detailed guidelines, reviewing open-source projects, and presenting domain-specific case studies supported by comprehensive datasets and evaluation benchmarks. Finally, Section 7 VII concludes by identifying future research directions and discussing potential challenges in knowledge quality, retrieval efficiency, system generalization, and security, alongside practical guidance for building domain-specific GraphRAG systems. Throughout the survey, we balance fundamental concepts, the current state of the art, and practical implementations, making it valuable for researchers advancing the field and practitioners developing GraphRAG applications in real-world scenarios.

Our survey advances beyond the existing survey [28]–[30] through a more systematic and comprehensive approach to GraphRAG systems. While the previous survey offered a basic workflow description covering Graph-based Indexing, Graph-guided Retrieval, and Graph-enhanced Generation, we introduce a more sophisticated and comprehensive taxonomy that clearly categorizes GraphRAG approaches into three distinct categories (Knowledge-based, Index-based, and Hybrid GraphRAG), providing a more nuanced understanding of the field. Our survey features a more systematic six-section structure that progresses logically from theoretical foundations to practical implementations, offering a more detailed exploration of each component, including knowledge organization paradigms, retrieval techniques, and integration methods. Unlike the previous survey, we provide extensive practical guidance through detailed review of open-source projects, and domain-specific case studies supported by comprehensive datasets and evaluation benchmarks. We also offer a more thorough analysis of challenges and solutions across multiple dimensions, including knowledge quality, retrieval efficiency, system generalization, and security concerns. Finally, while the



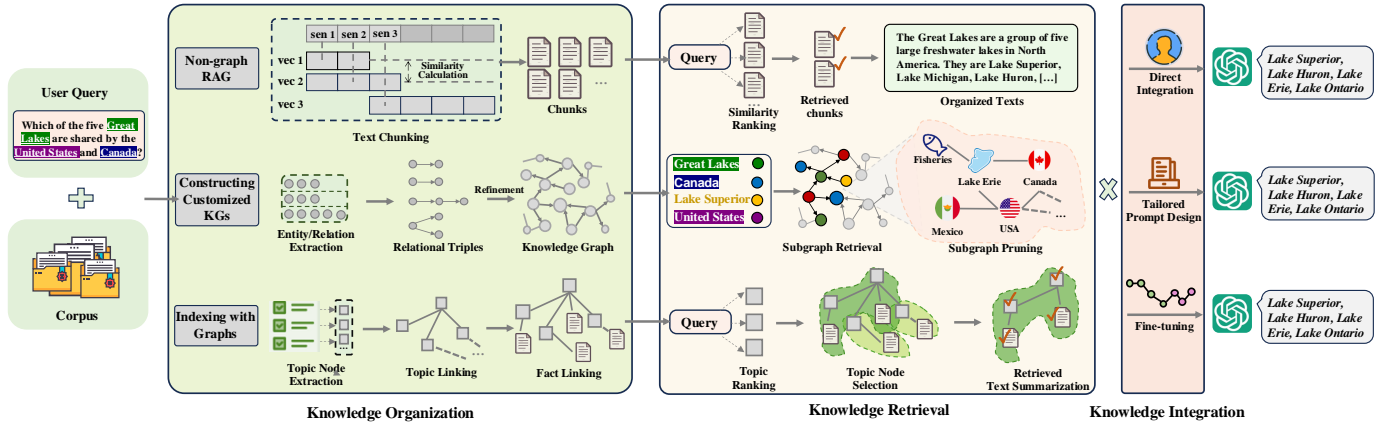


Fig. 2: A comprehensive overview of traditional RAG and two typical GraphRAG workflows. Non-graph RAG organizes the corpus into chunks, ranks them by similarity, and retrieves the most relevant text for generating responses. Knowledge-based GraphRAG extracts detailed knowledge graphs from the corpus using entity recognition and relation extraction, offering fine-grained, domain-specific information. Index-based GraphRAG summarizes the corpus into high-level topic nodes, which are linked to form an index graph, while fact linking maps topics to text. This two-layer structure combines efficient topic retrieval with detailed text knowledge, offering scalability and performance compared to Knowledge-based GraphRAG.

existing survey broadly discussed potential applications, we provide actionable insights supported by empirical evidence and implementation examples, making our survey a more valuable resource for practitioners deploying GraphRAG systems in production environments.

### III. WHAT IS GRAPH RAG

This section provides an overview of RAG with LLMs, including the traditional RAG pipeline, the GraphRAG, and the advantages GraphRAG offers over traditional RAG systems.

#### A. Traditional RAG Pipeline

A RAG framework begins by retrieving relevant information from pre-constructed external knowledge bases based on the query. This information is then used to prompt LLMs, guiding them in constructing credible reasoning chains. As a result, RAG enables LLMs to generate more substantiated and accurate content, effectively minimizing hallucinations and inconsistencies. The traditional RAG pipeline typically comprises three core components: knowledge organization, knowledge retrieval, and knowledge integration.

1) *Knowledge organization*: In traditional RAG systems, knowledge organization involves structuring and preparing external knowledge repositories to facilitate rapid and relevant retrieval when provided with a query. A common strategy is to split the large-scale text corpus into manageable chunks. These chunks are then transformed into embeddings using an embedding model, where the embeddings serve as keys of original text chunks in a vector database [81]–[83]. This setup enables efficient look-up operations and retrieval of relevant content via distance-based search in the semantic space.

As a crucial step in the pre-retrieval process, several methods [84], [85] have been proposed to optimize knowledge organization, focusing on two main aspects: granularity optimization and indexing optimization. Granularity optimization

aims to balance relevance and efficiency, as coarse-grained units provide richer context but risk redundancy and distraction, while fine-grained units may lack semantic integrity and increase the retrieval burden [86], [87]. To control granularity, chunking strategies are employed to split documents into chunks based on token limits. Methods such as recursive splits, sliding windows, and Small-to-Big [85], [88] strive to maintain semantic completeness while optimizing context length. Indexing optimization seeks to improve the structure and quality of content for retrieval. Metadata-addition techniques, which attach chunk text with metadata like titles, timestamps, categories, and keywords, enable filtering and re-ranking operations during the post-retrieval process [85]. Another type of technique is hierarchical indexing, which organizes files into parent-child relationships with summaries at each node, facilitating faster and more efficient data traversal while reducing retrieval errors [89]. Such tree-like indexing methods represent early attempts at structured knowledge organization and have inspired successors to harness the power of graph structures for knowledge organization, i.e., GraphRAG.

In summary, knowledge organization is foundational to the retrieval process. By carefully constructing the knowledge resources, RAG systems can ensure the fidelity and reliability of retrieved content.

2) *Knowledge retrieval*: The knowledge retrieval stage encompasses various methods and strategies designed to efficiently access and retrieve the necessary knowledge from pre-organized repositories, ensuring the selection of relevant information that can enhance the quality of generated outputs.

Current RAG works usually involve retrieval methods such as k-nearest neighbor retrieval (KNN), term frequency-inverse document frequency (TF-IDF), and best matching 25 (BM25) to retrieve the relevant content. RETRO [81] employs KNN to extract approximate relevant neighbors from the conducted key-value database by calculating the L2 distance. RETRO-prompt [90] extends this approach to construct a few-shot

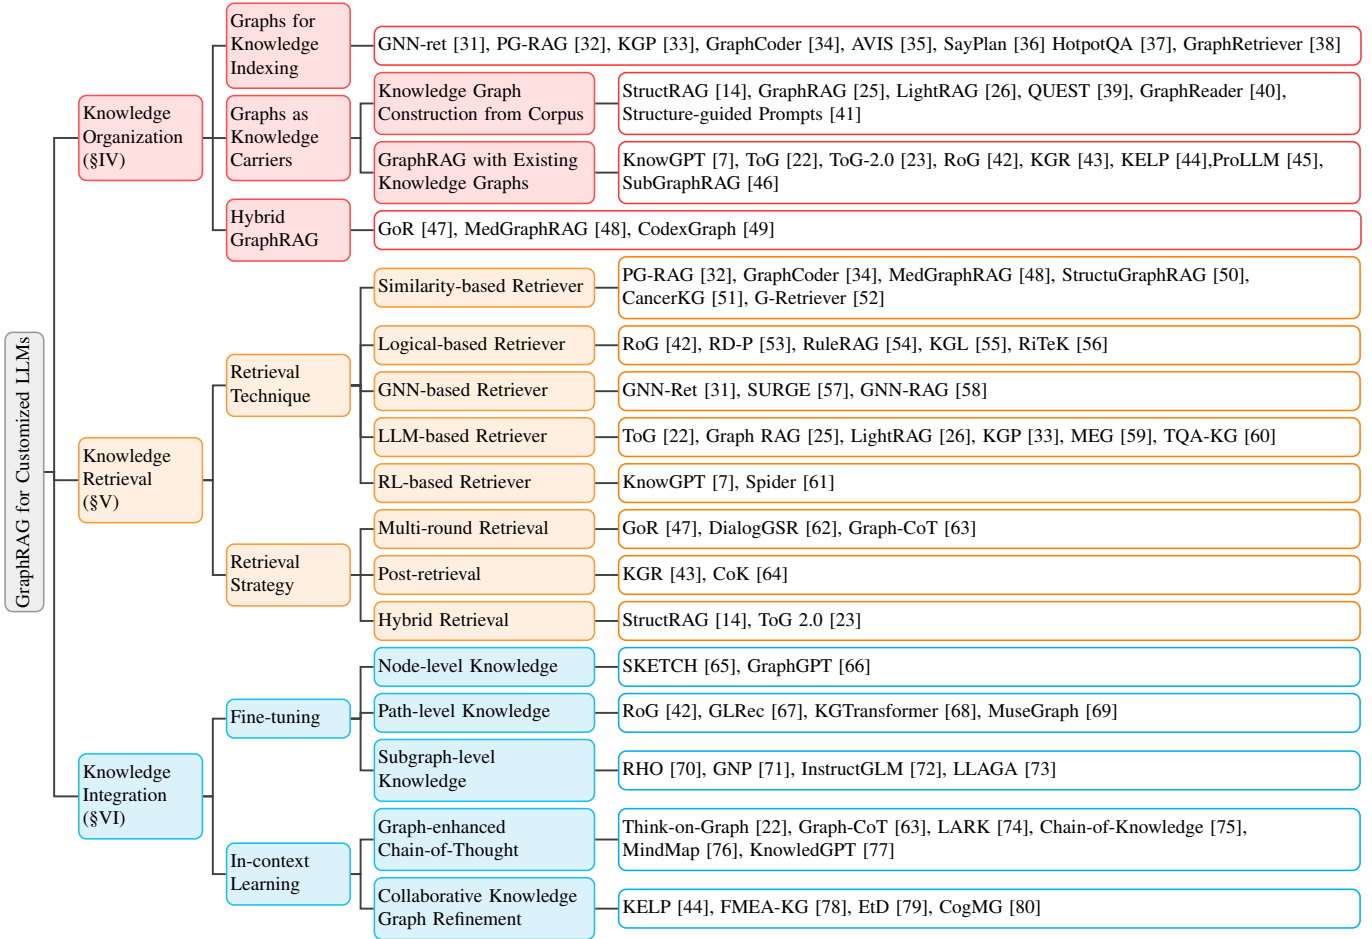


Fig. 3: The taxonomy for existing GraphRAG methods in the survey.

knowledge store, tailoring it specifically for prompt learning scenarios.

In addition, some specialized techniques are used prior to the retrieval method to improve the accuracy and efficiency of the retrieval. To capture multifaceted aspects of the query, GAR [91] introduces diverse context generation, enriching the initial query with additional contexts before applying BM25 retrieval. Enhancing this framework, EAR [92] implements a re-ranking process that selects the optimal candidate from multiple expanded queries to improve the retrieval accuracy. Furthermore, in tackling the computational challenges associated with exact retrieval methods like BM25, Doostmohammadi et al. [93] propose a hybrid approach to identify approximate neighbors using sentence transformers for representation and then apply BM25 for re-ranking, effectively balancing accuracy with computational efficiency on large-scale retrieval tasks.

Retrieval model training is another crucial aspect, where recent methodologies have explored the use of self-supervised techniques. For instance, REPLUG [94] demonstrates that labels generated by a frozen language model can be leveraged to directly supervise the training of retrieval models. This approach enhances retrieval quality without requiring additional manually annotated datasets. ATLAS [82] fine-tunes both the retriever and the LLMs in tandem. Furthermore, frameworks like FLARE [83] dynamically balance retrieval with generation

needs through an active RAG framework that intelligently determines retrieval timings based on the generation progression.

Moreover, the use of external APIs helps to broaden the scope of retrieval capabilities. By integrating external knowledge sources like Wikipedia and Google Search APIs, RAG systems can access expansive real-time databases, enriching their generated content. For example, Toolformer [95] taps into Wikipedia to access extensive knowledge bases, while the work by Lazaridou et al. [96] utilizes Google Search to ensure the most current and comprehensive information is incorporated, granting RAG systems a dynamic edge in content generation.

Overall, the knowledge retrieval process combines innovative methodologies with diverse data sourcing strategies, underpinning the generation of informed and contextually relevant outputs.

3) *Knowledge Integration*: The knowledge integration phase in RAG frameworks is crucial for synthesizing coherent and accurate responses based on both retrieved and inherent knowledge. At this stage, researchers utilize LLM to generate the output and employ several preprocessing and efficiency strategies to improve its quality and efficiency.

The quality of retrieved content can drastically affect generation, with irrelevant or misleading data potentially having deleterious effects. LeanContext [97] addresses this issue by utilizing reinforcement learning to selectively choose sentences

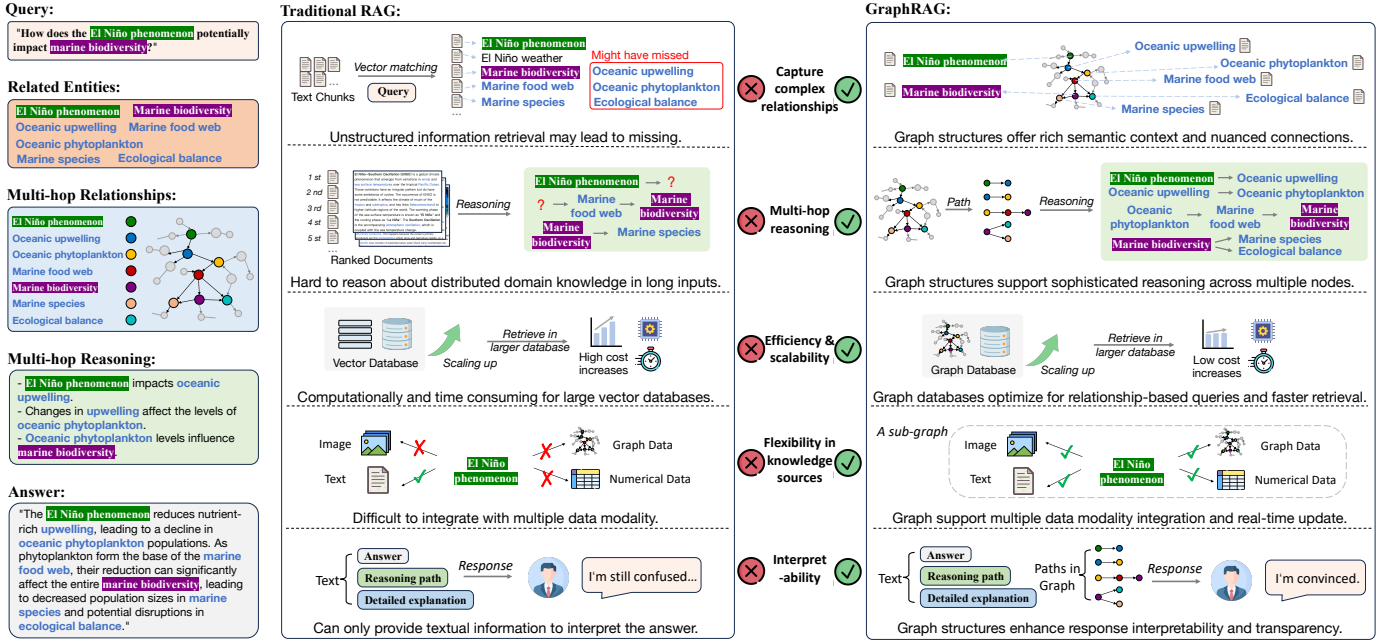


Fig. 4: The illustration of the comparison between traditional RAG and GraphRAG.

that are most pertinent to the query, effectively minimizing the context size and reducing computational costs. Similarly, SELF-RAG [17] introduces a self-reflection mechanism where the language model assesses its own generated and retrieved content, allowing for corrective measures to be implemented during the generation process. SKR [98] proposes a dynamic framework whereby LLMs can rely on their pre-trained knowledge for recognized queries, resorting to retrieval only when necessary. FILCO [99] enhances this by training a context filtering model that screens out irrelevant data, thereby mitigating hallucination risks. Recent research by Lyu et al. [100] introduces an evaluation method specifically designed to assess the importance of the retrieved content. Their findings suggest that by pruning or re-weighting parts of the retrieval corpus, RAG systems can enhance their performance without requiring additional training rounds. In scenarios where the retrieval corpus is limited, approaches such as Selfmem [101] construct a memory pool from LLM-generated results, using an iterative selection framework to enhance generative quality. SAIL [102] forms an instruction-tuning dataset based on retrieval outcomes, fine-tuning LLMs to ground responses in reliable content while excluding distracting elements.

Efficiency is another crucial concern, since encoding a large volume of retrieved passages is resource-intensive, leading to significant computation and memory overheads. To circumvent this, [103] adopts a pre-encoding strategy that stores dense representations of the text corpus. LUMEN [104] takes it a step further by conditioning the encoder on current inputs, and fine-tuning it specifically for this task to ensure high-quality pre-encoded memory representations. Innovative parallelism solutions, like PipeRAG [105], advance efficiency by enabling concurrent retrieval and generation, thus introducing alignment between retrieved content and real-time generation states. RAGCache [106] further optimizes resource usage by caching

intermediate states within a GPU-resident knowledge tree, employing a prefix-aware replacement policy to maintain the most critical key-value tensors. To efficiently compress retrieved contents, Xu et al. [16] introduce an extractive compressor to identify useful sentences and an abstract compressor to distill summaries; this dual compression system refines textual summaries before they are integrated into the context. Similarly, TCRA-LLM [107] utilizes a token compression strategy, optimizing the semantic relevance of tokens to conserve computational resources.

In conclusion, the generation stage within RAG focuses on bolstering the quality and efficiency of LLM-driven content generation. These methodologies continually evolve to maximize the potential of LLMs in generating context-aware, accurate, and efficient responses for RAG.

### B. Limitations of traditional RAG

Although researchers have extensively explored traditional RAG, there are still some unresolved limitations due to the constraints of the data structure itself.

1) *Complex query understanding*: Traditional RAG faces significant challenges in precisely answering complex queries, mainly due to the intrinsic limitation of its knowledge organization (vector database). Given a query, these RAG methods can only retrieve information from chunks containing anchor entities and are incapable of multi-hop reasoning. This limitation becomes more pronounced as granularity decreases, making it challenging to handle domain knowledge, as domain knowledge and complex logic often require multi-hop reasoning for effective understanding. Some traditional RAG methods have attempted to improve the complex query understanding, such as enhancing queries before retrieval [91], enhancing multiple query candidates for re-ranking [92], and using related metadata to provide richer information [85]. However,

these methods are still limited by the chunking style of knowledge organization and are challenging to capture multi-hop information in complex queries effectively.

2) *Distributed domain knowledge*: Domain-specific queries often involve jargon that requires contextual comprehension. However, domain knowledge is typically sparsely distributed across various documents and data sources. Although RAG uses chunking to divide documents into smaller pieces to manage this complexity and improve indexing efficiency, it sacrifices critical contextual information, significantly reducing retrieval accuracy and contextual understanding. Additionally, vector databases store the text chunks without a hierarchical organization of vague or abstract concepts, so it becomes difficult to resolve such queries, reducing the efficiency of utilizing expertise in specific domains. Some current traditional RAG works continuously align retrieved content with real-time generation states [105], and introduce external APIs to provide dynamic and rich auxiliary information [96]. While these approaches enhance RAG’s understanding of dispersed domain knowledge, they still rely on the traditional chunking method for knowledge organization, which cannot fully address distributed domain knowledge. This limitation restricts their ability to construct robust connections between pieces of knowledge.

3) *Inherent constraints of LLMs*: Traditional RAG often uses the vector similarity-based retrieval module, which usually lacks effective filtering of the retrieved contents from vast knowledge bases, providing excessive but maybe unnecessary information. Considering the inherent constraints of LLMs, such as the fixed context window (typically 2K-32K tokens) [1], [24] and the challenge of fully capturing long-range dependencies in complex documents, it is hard for LLMs in traditional RAG to capture the necessary information from the excessive retrieved contents. While scaling chunk granularity could alleviate these challenges, this approach significantly increases computational costs and response latency. Furthermore, indexing-based methods fail to effectively prune irrelevant information during retrieval, leading to the target knowledge being overwhelmed. These limitations present practical challenges for deploying RAG in resource-constrained environments where LLMs with smaller input context windows are preferred. Existing traditional RAG approaches have proposed methods such as recursive splits, sliding windows, and Small-to-Big [85], [88] to optimize context length while preserving semantic integrity. They also assess and filter retrieved information to extract key insights [17], [97], [99], [100], and fine-tune LLMs to enhance their capability to capture crucial information from a vast amount of retrieved data [102]. These approaches effectively reduce the volume of retrieved content, aiding the handling of LLMs’ limited context windows, but they still fall short of capturing long-distance dependencies. This shortcoming arises because vector similarity-based retrieval makes it challenging to establish sufficiently clear connections between different pieces of retrieved content, limiting LLMs’ understanding of long-distance dependencies.

4) *Efficiency and scalability*: Since large-scale knowledge sources often contain a significant amount of non-domain-specific information and domain-specific terms are usually sparsely distributed across diverse knowledge carriers, the

retrieval module of RAG systems often needs to search through a vast amount of unstructured text to find relevant information, requiring considerable computational resources and time [25]. Although there are existing traditional RAG methods that propose preprocessing before employing retrieval methods to reduce costs [93], [103], [104], or using innovative parallelism solutions to extract and maintain key information in intermediate processes to optimize resource use [16], [106], [107], they inevitably need to retrieve information from a large volume of unstructured text. This makes resource consumption a persisting issue for traditional RAG, affecting its scalability.

### C. GraphRAG

To address the limitations of traditional RAG, a novel paradigm known as Graph Retrieval-Augmented Generation (GraphRAG) has been introduced. Leveraging structured knowledge, GraphRAG provides an efficient and accurate solution for organizing and retrieving information with structural databases, enhancing the performance and reliability of RAG systems.

1) *Definition of GraphRAG*: GraphRAG can be formally defined as a subclass of RAG framework that **leverage graph structure to organize and retrieve knowledge**. Unlike traditional RAG methods, which rely on vector databases for knowledge organization, GraphRAG employs structural databases where graphs are used to model dependencies among knowledge pieces. This approach enhances the accuracy and efficiency of information retrieval, enabling more effective augmented generation for LLMs.

Specifically, GraphRAG systems can utilize graphs as either the carrier of knowledge or the indexing tools for efficient retrieval from chunked textual data. By modeling dependencies between nodes, GraphRAG enables the discovery of related knowledge pieces centered around a topic or anchor entity, ensuring comprehensive knowledge retrieval. Moreover, these connections support efficient search by navigating through relevant pathways and meanwhile pruning irrelevant information during the retrieval process.

2) *Workflow of GraphRAG*: Like traditional RAG, the workflow of GraphRAG can be divided into three key stages: knowledge organization, knowledge retrieval, and knowledge integration. However, due to the structured nature of GraphRAG, they have some special strategies compared to traditional RAG:

- **Knowledge Organization** The knowledge organization stage structures external information using graph-based methods, either through explicit knowledge representation (graphs as knowledge carriers) or indexing mechanisms (graphs for knowledge indexing). These approaches enable efficient, context-aware information retrieval.
- **Knowledge Retrieval**. GraphRAG models employ graph-based planners (either learnable planners or graph algorithm-based planners) to retrieve relevant information based on the input query. These retrieval techniques not only consider the semantic similarity between the query and each text chunk, but also the logical coherence between the query type and the retrieved subgraph.
- **Knowledge Integration**. Once the relevant knowledge is retrieved, the GraphRAG model integrates it with

the input query to generate the final output. The goal of the integration process is to seamlessly incorporate the retrieved knowledge into the generated text, thereby enhancing its quality and informativeness. A key design consideration is how to preserve the richness of the retrieved subgraph information in the final text-based prompt without introducing redundancy or misplacing emphasis on less critical aspects of the textual descriptions.

#### D. Traditional RAG vs GraphRAG

GraphRAG provides several key advantages over traditional RAG systems, enhancing the capabilities of AI-driven information retrieval and generation. Below, we discuss these benefits and illustrate the comparison between traditional RAG and GraphRAG in Figure 4.

1) *Enhanced Knowledge Representation*: GraphRAG utilizes graph structures to represent knowledge, capturing complex relationships between entities and concepts. This approach allows for a more nuanced and contextual understanding of information compared to the flat document-based representation used in traditional RAG. The graph structure in GraphRAG can represent hierarchies, associations, and multi-hop relationships, providing a richer semantic context for queries and revealing non-obvious connections between different pieces of information. This capability can lead to new insights and discoveries, making GraphRAG particularly valuable in research and analysis applications. By representing multiple possible interpretations or relationships in the graph, GraphRAG can better handle ambiguous queries. It can explore different semantic paths and provide responses that account for various possibilities, offering a more nuanced understanding of complex topics.

2) *Flexibility in Knowledge Sources*: GraphRAG systems can adapt to and integrate various knowledge sources, including structured databases, semi-structured data (like JSON or XML), and unstructured text. This versatility allows organizations to leverage their existing data infrastructure while benefiting from the advanced capabilities of GraphRAG. The system can connect data types, providing a unified view of an organization’s knowledge. GraphRAG can incorporate different types of data (text, images, numerical data) into a single graph structure. This capability allows for more comprehensive knowledge representation and the ability to answer queries that span multiple data modalities.

3) *Efficiency and Scalability*: GraphRAG systems built on fast graph databases can handle large-scale datasets efficiently. Graph databases are optimized for relationship-based queries, allowing for quick traversal of complex data structures. This efficiency translates to faster response times, especially for queries that require exploring multiple relationships. Research has shown that GraphRAG systems can generate LLM responses using 26% to 97% fewer tokens compared to traditional methods, indicating significant improvements in both speed and resource utilization. GraphRAG systems can more easily accommodate updates to the knowledge base. New information can be added as nodes or edges in the graph without requiring a complete reindexing of the entire knowledge base. This feature

allows for real-time updates and ensures that the system always has access to the most current information. The graph structure allows GraphRAG to consider the context of a query more effectively. Instead of relying solely on keyword matching or vector similarity, it can leverage the semantic relationships in the graph to retrieve more relevant information. This context awareness leads to more accurate and pertinent responses to user queries.

4) *Interpretability*: The graph structure in GraphRAG allows for better visualization and understanding of how the system arrives at its responses. Users can trace the path of reasoning through the knowledge graph, seeing which entities and relationships were considered in formulating the answer. This transparency is crucial for building trust in AI systems and is especially valuable in fields like healthcare, finance, or legal applications where decision-making processes need to be auditable. By integrating LLMs with graph-based knowledge representation, GraphRAG can perform more sophisticated reasoning tasks. It can follow chains of logic across multiple nodes in the graph, combining information from various sources to answer complex queries. This capability allows for multi-hop question answering and the ability to draw insights that may not be explicitly stated in any single document.

In conclusion, GraphRAG represents a significant advancement over traditional RAG systems, offering improved knowledge representation, reasoning capabilities, and efficiency. Its ability to capture and utilize complex relationships in data makes it a powerful tool for next-generation AI applications that require deep understanding and contextual awareness.

## IV. KNOWLEDGE ORGANIZATION

A major distinction of GraphRAG techniques compared to traditional RAG methods lies in their ability to leverage graph structures for efficient knowledge organization, enhancing the effectiveness of query responses. In this setup, the LLM functions as an intelligent agent, while the graph structures support its ability to organize and integrate information more comprehensively and precisely. In real-world applications, external knowledge sources may include extensive text corpora, document collections, search results, historical user data, or interaction logs. Properly organizing these sources is essential, as direct retrieval can be prone to irrelevant results and missing contexts, leading to information overload, knowledge conflicts, and compromised comprehension.

To address this, GraphRAG methods employ a two-step process: first constructing a graph structure to organize knowledge, then retrieving and integrating information relevant to the query. The organization of knowledge varies by task and source type, with three primary paradigms: (1) *Graphs for Knowledge Indexing* for **index-based** graphRAG: Text chunks are organized by symbolizing each chunk as a node within a graph, within which edges between nodes signify query-oriented relationships, making it easier to localize relevant knowledge. Considering knowledge hierarchy, further work builds multi-level index graphs, involving bottom-up knowledge summarization and top-down knowledge localization to enable coarse-to-fine knowledge indexing. (2) *Graphs as Knowledge*



TABLE I: Comparison of three knowledge organization paradigms: (i) vector databases in traditional RAG systems (“vector database”), (ii) graphs for knowledge indexing (“index graphs”), and (iii) graphs as knowledge carriers (“knowledge graphs”).

RAG Technique		Traditional RAG		GraphRAG	
Knowledge Organization Paradigm		Vector Database	Index Graphs	Knowledge Graphs	
Advantages	Explicit Reasoning on Graphs	✗	✗	✓	✓
	Index Original Content	✓	✓	✗	✗
	Logic-guided Retrieval	✗	✗	✓	✓
	Community Detection	✗	✓	✓	✓
	Prune Query-irrelevant Facts	✗	✗	✓	✓
	Missing Fact Completion	✗	✗	✓	✓
	Multi-hop Knowledge Edit	✗	✗	✓	✓
Disadvantages	Sensitive to Chunk Granularity	✓	✓	✗	✗
	Information Loss	✗	✗	✓	✓

**Carrier for knowledge-based graphRAG:** Here, summarized knowledge (e.g. atomic facts, community summarizations) is extracted from text chunks and integrated into a unified knowledge base, serving as a carrier of essential knowledge condensed from the raw corpora. (3) **Hybrid GraphRAG:** This approach combines both index graph and knowledge carrier functions, benefiting from the logical reasoning of knowledge graphs while preserving the detailed information in raw text chunks through indexing. The following sections detail each of these paradigms, explaining how each type of graph structure facilitates efficient and effective retrieval in diverse scenarios.

#### A. Graphs for Knowledge Indexing

Index-based GraphRAG methods utilize graph structures to index and retrieve relevant raw text chunks, which are then fed into LLMs for knowledge injection and contextual comprehension. These index graphs apply principles such as semantic similarities or domain-specific relations to effectively bridge connections among separate text passages. Compared to using graphs solely as knowledge carriers, this technique provides more informative answers by directly summarizing information from query-related raw text chunks.

The inherent challenges associated with the construction and maintenance of these graphs raises several research questions, giving rise to several research questions that shape current research directions: (1) **Conciseness and Relevance:** Ensuring the constructed graph captures only relevant relationships without overloading it with unnecessary connections is a significant challenge, thus facilitating effective recalling of relevant text chunks without redundancies, (2) **Consistency and Conflict Resolution:** Different data chunks may introduce conflicting information. It is crucial to resolve these conflicts and ensure that the graph remains consistent, reliable, and well-structured. Researchers are actively exploring tailored index graph structures to address these challenges effectively.

For *general-purpose* passages indexing, GNN-ret [31] creates a passage-level graph, connecting passages through structural and keyword similarities. PG-RAG [32] treats LLMs as learners, forming a pseudo-graph by linking concise reading material summaries through common topics or complementary facts, generating an interconnected pseudo-graph database. KGP [33] examines the logical associations of three major types of questions and develops three methods for constructing

similarity-based index graphs. The first method uses TF-IDF to create lexical-similarity-based indexing graphs. The second employs KNN-ST/MDR [37] to build semantic-similarity-based indexing graphs. The third utilizes TAGME [38] to construct indexing graphs based on shared entities.

A number of methods focus on *domain-specific* index graph construction. In the code completion task, GraphCoder [34] utilizes a code context graph (CCG) that captures control flow and data dependencies between code statements, aiding in retrieving related code snippets for completion. AVIS [35] generates a transition graph based on user decisions, mapping states, and permissible actions. For robot task planning, SayPlan [36] employs 3D scene graph representations, enabling LLMs to perform semantic searches for planning purposes. These specialized methods tailor index graphs to meet the unique demands of various domain-specific tasks.

In summary, index graph-based methods have evolved from general-purpose passage and entity retrieval strategies to sophisticated pseudo-graphs for coherent information organization and, finally, to highly specialized graphs for domain-specific applications. Each development has addressed unique challenges, from organizing large text corpora and contextual summaries to leveraging inductive bias for enhanced domain-specific tasks like code completion, decision-making, and robotic planning. This evolution underscores the adaptability of graph-based structures in meeting the diverse and growing demands of LLM-driven applications across multiple domains.

#### B. Graphs as Knowledge Carriers

The paradigm of representing knowledge explicitly through graph structures is increasingly recognized for its effectiveness, offering several notable advantages: 1) **Efficient retrieval of query-related knowledge facts.** The explicit structure of knowledge graphs facilitates logic-guided chain retrieval, efficiently identifying missing facts while pruning the search space through reasoning paths; 2) **Coherent multi-step reasoning over long spans.** Previous studies have revealed the importance of the reasoning steps, i.e., lengthening the reasoning steps in prompts brought significant reasoning performance enhancement [108]. Knowledge graphs excel in maintaining logical coherence for multi-step reasoning, where natural language often struggles. By leveraging planning algorithms, they identify optimal reasoning paths, enhancing both

interpretability and clarity in handling complex information; **3) Effective handling of linguistic diversity.** Natural language often presents challenges with varying terminologies for the same entities or relationships. Knowledge graphs address this by enabling efficient entity and relation resolution through off-the-shelf graph algorithms, streamlining the integration of diverse data sources.

Research in this area can be categorized into two main directions based on the source of the knowledge graphs: *GraphRAG with KGs constructed from corpus*, which transforms unstructured text into structured graph representations, and *GraphRAG with existing KGs*, which focus on how to effectively leverage well-established KGs. Below we present the details of each direction and discuss the limitations of them.

1) *Knowledge Graph Construction from Corpus*: This research direction aims to transform unstructured text corpora into structured KGs, enabling efficient and precise information retrieval. It heavily relies on *Open Information Extraction* techniques, which automate the domain-independent extraction of relational facts from large corpora. These techniques include both traditional OIE approaches [109]–[112] and advanced OIE methods based on LLMs [113]–[116]. However, constructing effective KGs from domain-specific corpora remains challenging due to variations in ontology across domains. A key objective in this area is to incorporate inductive biases to design graph structures that capture essential domain-specific information while minimizing unnecessary complexity.

In scenarios where the query prompt consists of short passages and no external corpus is provided, Structure-guided Prompts [41] generate logically coherent responses by categorizing reasoning tasks into six types based on logical dependencies and organizing the passage’s information using tailored graph structures specific to each type. Other studies have focused on scaling language models to handle multi-document inputs through dynamically constructed knowledge graphs. For instance, [117] transformed sentences into relational triples via Coreference Resolution and Open Information Extraction, subsequently reducing graph redundancy by merging nodes and edges based on TF-IDF overlap. Similarly, AutoKG [118] constructs a weighted knowledge graph where nodes represent keywords extracted from the original text, and edges, evaluated through graph Laplacian learning, represent relationships with assigned weights. To offer more informative knowledge augmentation, QUEST [39], GraphRAG [25], and GraphReader [40] construct attributed KGs where nodes are enriched with side information. Specifically, QUEST emphasizes integrating entities, relations, types, and semantic alignments, while GraphRAG leverages an LLM to generate community summaries as enriched knowledge for related entities, subsequently using these summaries to produce partial answers for the entities. Similarly, GraphReader employs an LLM to summarize text chunks into atomic facts and extract key elements (nouns, verbs, and adjectives). Each key element, combined with its relevant atomic facts, forms a node, with links established between nodes sharing elements in the same atom facts. To leverage LLMs for customer service, [119] propose using RAG to retrieve and integrate historical issues. They create an attributed knowledge graph where nodes

represent past issues and edges denote their connections. Each node includes a tree structure of the issue’s information, capturing both inter-issue and intra-issue relationships to improve question answering. Recent studies have extended graph structures to more sophisticated forms. For example, StructRAG [14] considers five candidate structure types tailored to different knowledge-intensive tasks: tables for statistical tasks, graphs for long-chain tasks, algorithms for planning tasks, catalogs for summarizing tasks, and chunks for simple single-hop tasks. A DPO-based router is trained to determine the optimal structure for representing knowledge in each case.

2) *GraphRAG with existing KGs*: This line of research focuses on utilizing well-established knowledge graphs, which can be either domain-specific—such as Lynx [120] for multilingual compliance services in the legal domain, AceKG [121] for academic applications, SPOKE [122] for biomedical applications, STRING [123] for protein-protein interaction prediction, — or general-purpose, like DBpedia [124] and YAGO [125]. The central challenge is to develop planning algorithms that dynamically retrieve reasoning paths or subgraphs from these KGs, tailored to the specific characteristics of the queries and the underlying graph structures. Researchers are investigating techniques to efficiently navigate these large-scale graphs, aiming to extract meaningful insights while ensuring the retrieved information is both relevant and concise.

The retrieval of reasoning paths typically involves generating relation paths grounded in KGs as plans, which are then used to extract factual knowledge from the KGs in the form of reasoning chains (relation paths with anchored entities). ProLLM [45] applies graph-based RAG to predict protein-protein interactions, modeling signaling pathways as reasoning processes where biological signals pass from upstream proteins through intermediates to downstream proteins via the shortest path. Approaches such as RoG [42] and KGR [43] directly utilize LLMs as planning agents to identify optimal plans. To address the low recall of reasoning paths caused by the incompleteness of KGs, ToG [22] employs a beam search algorithm to dynamically explore the most probable relational triples, forming coherent reasoning chains. Its successor, ToG-2 [23], enables joint retrieval from a KG and textual documents, through alternating between logic chain extensions on KGs that explore neighboring entities of current ones, and contextual knowledge expansion from retrieved relevant documents. In the biomedical domain, KG-RAG [126] focuses on retrieving accurate and trustworthy biomedical contexts. It optimizes the retrieval process by using techniques such as disease entity recognition, context pruning, and leveraging a biomedical KG [122]. To optimize the planning process, KnowGPT [7] formulates the exploration of plans as a planning problem and solves it using deep reinforcement learning. Meanwhile, KELP [44] adopts a simpler semantic-matching strategy, encoding both the question and candidate knowledge paths in an encoder fine-tuned on a latent semantic space, and selecting the knowledge paths with the smallest semantic distance. To balance precision and recall in reasoning chain retrieval, subgraphRAG [46] introduces a lightweight multilayer perceptron to extract subgraphs, discovering key facts within the KGs. This method allows the subgraph size to be adjusted,

achieving an efficiency-effectiveness trade-off. More details about the retrieval techniques are presented in Section V.

3) *Limitations*: Although using graphs as knowledge carriers is efficient and effective, these methods are limited in several aspects: (i) **Lack of high-quality KGs**. For directly using KGs as external knowledge bases, this line of research is constrained by the availability of high-quality KGs. Constructing KGs is resource-intensive, and although powerful KG refinement techniques [127]–[132] have been developed to improve KG completeness, most publicly available KGs remain far from comprehensive. Furthermore, the absence of a unified ontology poses challenges to designing transferable planning algorithms. These limitations have motivated recent research to focus on constructing KGs from corpora. However, this line of work introduces additional challenges: (ii) **Trade-off between efficiency and effectiveness**. When constructing KGs from text corpora, the granularity of the extracted knowledge plays a crucial role in balancing efficiency and effectiveness. Preserving fine-grained information results in larger, more detailed KGs, which can hinder computational efficiency. Conversely, compact KGs may sacrifice important details, leading to potential information loss. This trade-off complicates the knowledge summarization process. Moreover, approaches like GraphRAG [25] and GraphReader [40], which leverage LLMs for knowledge summarization, face additional challenges due to the high computational costs associated with LLMs.

### C. Hybrid GraphRAG

This paradigm utilizes graph structures both as carriers of knowledge and as indexing tools. A common approach involves constructing a graph that encapsulates key information from the original text, with each node linked to corresponding text chunks. These text chunks act as a complementary knowledge source, providing detailed contextual information.

To enhance long-context summarization, GoR [47] constructs a graph by creating nodes for text chunks and their LLM-generated responses to simulated queries, with edges linking the queries, the retrieved text chunks, and the responses. This structure captures complex correlations among the elements through query-specified indexing, offering a comprehensive representation for effective long-context global summarization. In the medical domain, MedGraphRAG [48] introduces a unique triple graph construction. This approach creates a triple-linked structure that connects user documents to credible medical sources and controlled vocabularies, utilizing medical domain relations. This facilitates evidence-based medical responses with enhanced safety and reliability when handling private medical data. To enhance repo-level code tasks, CodexGraph [49] uses static analysis to construct a code graph that links source code symbols (e.g., MODULE, CLASS, FUNCTION) through edges denoting relationships (e.g., CONTAINS, INHERITS, USES). This graph acts as a knowledge carrier by implying the workflow or data flow of the code and serves as an index for the metadata (source code pieces) associated with each node.

## V. KNOWLEDGE RETRIEVAL

Based on the well-built knowledge base with graph-based organization strategies, we need to conduct retrieval from the knowledge base. Specifically, given a user query and a graph knowledge base with dense information, retrieving factual information relevant to the given query from the knowledge base is very important in developing effective and efficient GraphRAG systems. This retrieval process forms the cornerstone of a robust GraphRAG architecture, directly impacting its overall performance and utility. In this section, we will introduce current knowledge retrieval methods for graph-based retrieval-augmented generation in detail.

### A. Overall Pipeline of Knowledge Retriever

The knowledge retrieval process in GraphRAG focuses on extracting relevant background knowledge from a graph database in response to a given query. This process follows three distinct and sequential steps that transform raw graph data into usable, contextual knowledge. The overall process of knowledge retrieval is shown in Figure 5.

1) *Query/Graph Preprocessing*: The preprocessing stage operates simultaneously on both the query and graph databases to prepare them for efficient retrieval. For query preprocessing, the system transforms the input question into a structured representation through vectorization or key term extraction [34], [50], [51]. These representations serve as search indices for subsequent retrieval operations. On the graph side, the graph database undergoes more comprehensive processing where pre-trained language models transform graph elements (entities, relations, and triples) into dense vector representations that serve as retrieval anchors [23], [33], [52], [60]. Additionally, some advanced retrieval models apply graph neural networks (GNNs) on the graph database to extract high-level structural features, while a few methods even adopt rule mining algorithms to generate rule banks as rich, searchable indexes of the graph knowledge [22], [31], [54], [57].

2) *Matching*: The matching stage establishes connections between the preprocessed query and the indexed graph database. This process compares the query representations against the graph indices to identify relevant knowledge fragments. The matching algorithm considers both semantic similarity and structural relationships within the graph. Based on the matching score, the system retrieves connected components and subgraphs that demonstrate high relevance to the query, creating an initial set of candidate knowledge.

3) *Knowledge Pruning*: The pruning stage refines the initially retrieved knowledge to improve its quality and relevance. This refinement process addresses the common challenge of retrieving excessive or irrelevant information, particularly when dealing with complex queries or large graph databases. The pruning algorithm applies a series of refinement operations to consolidate and summarize the retrieved knowledge [7], [32], [52]. Specifically, the system first removes clearly irrelevant or noisy information. It then consolidates related knowledge fragments and generates concise summaries of complex graph knowledge. The purpose of knowledge summarization is to facilitate comprehensive knowledge synthesis, which is vital

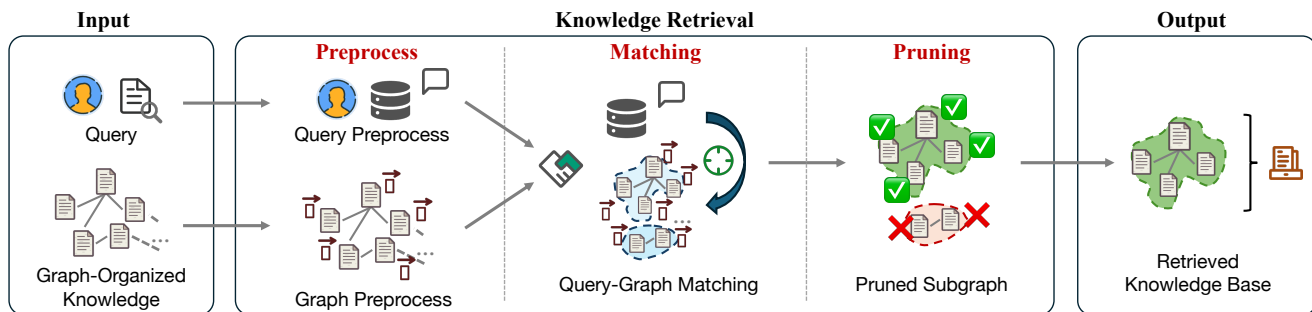


Fig. 5: Illustration of the overall pipeline of knowledge retrieval.

for the generation process of LLMs. By providing a distilled and focused summary, we enable LLMs to better understand the context and nuances of the information, leading to more accurate and meaningful responses.

### B. Retrieval Techniques

In this subsection, we provide a comprehensive overview of the various retrieval techniques employed in the GraphRAG knowledge base. These techniques can be broadly categorized into six main classes, including string matching, semantic similarity, logical reasoning, LLM-based approaches, RL-based methods, and GNN-based models, as shown in Table II.

1) *Semantics Similarity-based Retriever*: The semantic similarity-based retriever primarily conducts appropriate answer retrieval by measuring the similarity between queries in discrete linguistic space or continuous vector space and the knowledge base [32], [34], [50]–[52]. (i) **Discrete space modeling**. Discrete space modeling methods primarily leverage linguistic discrete statistical knowledge to directly model text strings, which is a kind of simple and straightforward retrieval technique that relies on the exact matching of query terms with entities and relations in the graph base. It employs algorithms like substring matching, regular expressions, and exact phrase matching to identify relevant nodes and edges [23], [51], [60], [84], [133]. Although straightforward to implement, string matching is limited by its inability to handle variations in terminology, synonyms, and contextual nuances. Despite these limitations, it serves as a foundational approach, especially useful in domains with well-defined vocabularies and minimal ambiguity. (ii) **Embedding space modeling**. Embedding space modeling techniques are employed to evaluate the contextual and conceptual connections between queries and the elements of a graph base to transcend the constraints of string matching [32], [34], [50], [52]. Utilizing approaches like pre-train language models and word embeddings, such as TF-IDF, Word2Vec, and GloVe [134], [135], these methods capture the semantic ties that bind different components. By mapping both queries and graph entities into a continuous vector space, these sophisticated techniques enable the retrieval of semantically pertinent information, bridging gaps even in the absence of precise term matches. However, the limitation of semantics similarity-based models is that they neglect the full exploitation of graph structures resulting in a significant underutilization of the inherent advantages of graph bases.

2) *Logical Reasoning-based Retriever*: Logical rule-based retrieval employs symbolic reasoning to deduce and extract pertinent information from graph knowledge bases. This methodology encompasses the creation of logical rules and constraints that articulate the relationships and hierarchies intrinsic to the knowledge bases. Utilizing techniques such as rule mining [54], inductive logic programming [136], and constraint satisfaction [42], this approach uncovers insights that may not be explicitly present in the data. By harnessing these logical inferences, GraphRAG systems are capable of retrieving information that is congruent with the underlying structure and semantics of the graph bases.

3) *GNN-based Retriever*: Graph neural networks (GNNs) have become a mainstream tool for graph modeling and mining, achieving significant results in tasks such as node classification [137]–[140] and link prediction [141]–[143]. The message-passing mechanism in GNNs operates by iteratively aggregating information from a node’s neighbors to update its own representation, enabling the network to learn from the structure and features of the graph [137], [144]. GNN-based knowledge retriever primarily utilizes graph neural networks to encode nodes within the constructed graph bases [31], [57], [57], [58], [62]. The localizing and positioning of knowledge mainly rely on the encoding similarity of the node representations that contain both sentiment meanings and structural relationship understandings. GNN-based retrievers require the training of a GNN encoder. Additionally, due to the lack of explicitly labeled data, the focus of training is on designing an appropriate loss function that enables the GNN to learn to accurately locate the target knowledge through representation encoding.

4) *LLM-based Retriever*: With the advent of powerful LLMs, retrieval techniques have increasingly incorporated these models to enhance the semantic understanding and contextual relevance of the retrieved information. LLM-based retrieval leverages the deep contextual embeddings and generative capabilities of LLMs to interpret queries, generate relevant queries for the KG, and even synthesize information from multiple sources within the knowledge bases [22], [23], [26]. These models can perform tasks such as query expansion, contextual disambiguation, and the generation of tailored retrieval paths, thereby improving the overall effectiveness of the retrieval process. Regarding the constructed graph base, the LLM-based knowledge retriever primarily focuses on leveraging the LLM to comprehend the graph and identify key subgraphs. Compared to methods



TABLE II: Representative knowledge retrieval techniques and strategies used in different GraphRAG Systems.

Category	GraphRAG Model	Input		Implementation Details			Output
		Query Side	Graph Side	Query/Graph Preprocess	Matching	Pruning method	
Similarity-based	StruGraphRAG [50]	query embedding	entity embedding	BERT/BERT	similarity calculation	×	literal context
	CancerKG [51]	keywords	entity	NA/NA	TF-IDF	×	tabular results
	G-Retriever [52]	query embedding	entity/relation embedding	SentenceBert/SentenceBert	k-nearest neighbors	PCST	subgraph
	PG-RAG [32]	query embedding	Pseudo-Graph	SentenceBert/SentenceBert	depth-first search	PGR	literal context, path
	GraphCoder [34]	query slice	CCG node	CodeBERT/CodeBERT	Jaccard index	subgraph edit distance	code snippet
	MedGraphRAG [48]	query	entity description	SentenceBert/SentenceBert	top-down search	bottom-up refine	triple
Logical-based	RoG [42]	query embedding	relation path	Llama2/path generation	breadth-first search	LLM agent	reasoning path
	RD-P [53]	query embedding	topic entity	RoBERTa/RoBERTa	path expansion	path discriminator	reasoning path
	RuleRAG [54]	query embedding	rule bank	Llama2/rule mining	similarity calculation	×	literal context, path
	KGL [55]	query embedding	clarification path	Roberta/path generation	similarity calculation	×	literal context
	RiTeK [56]	query embedding	triple graph	SentenceBERT/SentenceBERT	MCTS/R-MCTS	LLM agent	reasoning path
GNN-based	GNN-Ret [31]	subquestions	graph feature	SentenceBERT/RGNN	DPR	×	subgraph
	SURGE [57]	query embedding	triple embedding	T5-small/GNN	similarity calculation	DHT	triple
	GNN-RAG [58]	query embedding	graph feature	NA/GNN	similarity calculation	×	literal context
LLM-based	KGP [33]	query embedding	passage (node description)	T5/T5	TF-IDF	LLM agent	literal context
	Graph RAG [25]	query	community summary	×/community detection	LLM agent	×	literal context
	ToG [22]	query embedding	relation path	Llama2/relation exploration	beam search	LLM agent	reasoning path
	LightRAG [26]	keywords	entity, global keys	GPT-4o/GPT-4o	keyword-search	global keyword match	literal context
	MEG [59]	query embedding	graph feature	SapBERT/SapBERT	token generation	disambiguation	literal context
	TQA-KG [60]	keyword	triple	entity extract/×	keyword-search	LLM agent	subgraph
RL-based	KnowGPT [7]	query embedding	entity/relation embedding	BERT	similarity calculation	RL agent	reasoning path
	Spider [61]	query embedding	entity embedding	NA	similarity calculation	RL agent	subgraph
Multi-round	DialogGSR [62]	query embedding	linearized graph	T5/T5	subgraph generation	×	linearized subgraph
	Graph-CoT [63]	query embedding	entity embedding	Llama2/Llama2	similarity calculation	LLM agent	subgraph
	GoR [47]	query embedding	entity embedding	Mixtral-7B/GAT	similarity calculation	×	literal context
Post-retrieval	CoK [64]	pseudo evidence	triple	NA/NA	F2-Verification	×	final answer
	KGR [43]	claim	entity, triple	claim extraction/NA	claim verification	×	final answer
Hybrid Retrieval	StructRAG [14]	query embedding	entity embedding	Qwen2/Qwen2	similarity calculation	LLM agent	literal context
	ToG 2.0 [23]	keyword	entity, relation	keyword extract/NA	similarity calculation	entity/relation prune	literal context, path

based on semantic similarity, it provides more powerful text representations. Specifically, the LLM-based knowledge retriever capitalizes on the formidable comprehension and generation prowess of LLMs to interpret the query and extract or synthesize pertinent candidate instances. Utilizing similarity-based retrieval strategies, it calculates the affinity between these candidates and the nodes within the graph, thereby identifying the most pivotal subgraph segments. For instance, PG-RAG [32] initially employs LLMs to generate key points corresponding to the query. Subsequently, it identifies seed nodes within the graph by computing a similarity matrix. Starting from these seed nodes, it searches for candidate nodes and ascertains their contribution to the seed nodes through a contribution matrix. Ultimately, the path (fact trajectory) with the highest cumulative contribution is selected as the retrieved collection of node information. On the constructed hierarchical tree, Thought Graph [145] employs LLMs to recursively expand and vote on the generated biological process terms with the prompt, selecting the term set that most accurately describes the gene set. HOLMES [146] begins the process by leveraging the named entities extracted from the question as the foundation for a breadth-first traversal of the entity-document graph. By harnessing the capabilities of an LLM, it extracts KG triples from the document nodes and elevates them to a hyper-relations KG. Subsequently, the hyper-relations KG undergoes pruning through the computation of cosine similarity between relational embeddings, thereby ensuring that only the information most pertinent to the question is preserved. In general, the LLM-based retriever can guide the model through the use of prompts without model training. The need for retriever training arises when there is a requirement to utilize the retrieval results to guide the LLM in reflecting on or fine-tuning its existing knowledge to enhance its capability [42]. Recent studies also reveal the capability of LLMs to achieve few-shot or zero-shot transferable performance in graph-related tasks [147]–[149],

showcasing their potential for resource-constrained scenarios involving large-scale and diverse graphs.

5) *Reinforcement Learning-based Retriever*: Reinforcement learning (RL) provides an adaptive and dynamic strategy for retrieval within GraphRAG systems. By framing the retrieval process as a sequential decision-making challenge, RL-based methods enable an agent to learn and traverse the graph base in search of the most pertinent information, guided by environmental feedback. Advanced techniques like Deep Q-Networks [150], Policy Gradients [151], and Actor-Critic [152] methods are deployed to refine retrieval strategies progressively. This methodology endows the system with the capacity to enhance its retrieval performance continuously through active interaction and accrued experience. This process can be described as follows: The relevant reasoning background lies in a question-specific subgraph  $\mathcal{G}_{\text{sub}}$  that contains all the *source* entities  $Q_s$ , *target* entities  $Q_t$ , and their neighbors. An ideal subgraph  $\mathcal{G}_{\text{sub}}$  is expected to have the following properties: (i)  $\mathcal{G}_{\text{sub}}$  encompasses as many source and target entities as possible; (ii) The entities, and relations within  $\mathcal{G}_{\text{sub}}$  exhibit a strong relevance to question context; (iii)  $\mathcal{G}_{\text{sub}}$  is concise with little redundant information such that it can be fed into LLMs with limited lengths. However, it is challenging to find such a  $\mathcal{G}_{\text{sub}}$  since extracting a subgraph is NP-hard. To effectively and efficiently find a satisfactory  $\mathcal{G}_{\text{sub}}$ , some researchers develop tailored knowledge extraction methods that employ deep RL to sample reasoning chains in a trial-and-error fashion [7], [61].

### C. Retrieval Enhancement Strategies

Based on the retrieval technologies, some strategies are designed to enhance the retrieved resources, aiming to improve the relevance of target queries. We will introduce insights of these strategies in this subsection.

1) *Multi-round Retrieval*: Beyond static retrieval techniques, some GraphRAG systems incorporate contextual information

TABLE III: Representative fine-tuning knowledge integration techniques used in GraphRAG systems.

Category	Model	Input Format	Preprocess	Applied (L)LMs
Node-level	SKETCH [65]	Node attribute	Concatenate Node Attribute	LLAMA3-8B
	GraphGPT [66]	Node attribute & Graph Embeddings	Learning Node Embeddings	Baichuan-7B
Path-level	GLRec [67]	User-item description	Path Weight Calculation	LLAMA-7B
	KGTransformer [68]	Sampled Knowledge Graph Sequence	None	Transformers
	MuseGraph [69]	Node Path & Node attribute	COT	LLAMA-7B
	RoG [42]	Knowledge Graph Path	Instruction Generation	LLaMA2-Chat-7B
Subgraph-level	RHO [70]	KG Embeddings & KG SubGraph	None	BART
	GNP [71]	Graph & Query Embeddings	None	FLAN-T5-xxlarge
	InstructGLM [72]	Node attribute & Graph Embeddings	Flatten subgraph to sequence	LLAMA-7B
	LLAGA [73]	Node attribute & Graph Embeddings	Flatten subgraph to sequence	Vicuna-7B

and user feedback to refine and adapt the retrieval process dynamically, which we call multi-round retrieval. These approaches aim to gradually align the retrieval mechanism more closely with specific use cases and evolving user intents [47], [62], [63].

2) *Post-retrieval*: Traditionally, GraphRAG mainly adopted prior retrieval, which means the retrieval process is conducted before generation. Recently, a new type of post-retrieval strategy has received more and more attention [43], [64]. This kind of method conducts retrieval after the generation process. In this way, the retrieved results can be adopted to evaluate whether the generated answers are faithful and accurate. The evaluation can be used as the guidance for LLM rethinking and correcting.

3) *Hybrid Retrieval*: Combining different forms of data for retrieval can integrate their strengths and complement their weaknesses. In GraphRAG systems, there are also some recent works that adopt both graphs and other data architectures to collaboratively finish the RAG process. Specifically in recent years, to enhance the scope of the retrieval candidates, *Graph+Vector RAG* [119], [153] combine the knowledge graph and the vector database as the retrieval sources and integrate the retrieved results for the answer generation. For example, ToG-2 [23] performs joint retrieval on a knowledge graph and text documents. It begins with entities identified from the query and iteratively retrieves relevant knowledge. Due to the static nature of knowledge graphs, to further ensure the timeliness and accuracy of the retrieved information, some efforts have begun to combine knowledge graphs with online web resources for joint retrieval, which can be summarized as *Graph+Online Web Resource RAG* [154]–[159]. Further, K2 [160] explores the LLM tuning with multiple retrieval sources, such as the textual corpus, self-instructed LLM knowledge, and online web resources. In future works, hybrid RAG with multiple retrieval resources simultaneously can be a potential research direction for complex LLM system building under a wide range of knowledge.

#### D. Discussions

The true magic of graph-based knowledge resources lies in their supernatural ability to connect data while being flexible and scalable during retrieval. Graphs can seamlessly integrate various data sources, including structured and unstructured data. This fusion creates a unified view of information, often revealing hidden connections and patterns. Advanced graph knowledge bases incorporate ontologies and semantic schemas,

providing formal definitions for concepts and their interrelationships within a domain. While traditional retrieval methods are like tourists asking random pedestrians for directions, GraphRAG is the mayor who knows every nook and cranny of the city. It doesn't just look at individual buildings (or pieces of information); it sees the entire cityscape, considering the intricate network of streets, neighborhoods, and communities. This bird's-eye view allows GraphRAG to navigate these streets with the efficiency of a seasoned cab driver who knows all the shortcuts. Here is a vivid example: Let's consider GraphRAG to be a sophisticated urban planning and information retrieval system for the knowledge base aka city. This system operates in two main phases: planning and retrieval.

In the planning phase, the city planners (LLMs) first divide the urban landscape (source documents) into manageable districts (concepts). They then survey each district, identifying key landmarks and connections (element instances) - think of this as mapping out important buildings, parks, and the roads connecting them. These individual elements are then summarized into concise descriptions (element summaries), like creating brief profiles of each neighborhood landmark. Next, using advanced urban analysis tools (community detection algorithms like Leiden), the planners group these elements into natural communities, much like identifying distinct neighborhoods or boroughs within the city. Finally, they craft detailed "neighborhood reports" (community summaries) for each of these communities, providing a rich, multi-layered guidebook to the entire city. This indexing phase creates a comprehensive, hierarchical understanding of the urban landscape, from individual buildings to entire districts.

After the planning, we obtain a high-quality graph knowledge base. Then, when it comes to retrieval, GraphRAG shines in handling citywide inquiries. Instead of dispatching surveyors to random locations, it consults the pre-prepared neighborhood reports. Each community contributes its local perspective to the broader question, like local town halls holding simultaneous meetings (community summaries to community answers). These local insights are then synthesized into a comprehensive city report (global answer), much like an urban planner combining feedback from various boroughs to understand citywide trends. This approach proved particularly effective for "global" questions about the entire urban landscape, outperforming traditional methods in both the breadth and diversity of insights offered.

## VI. KNOWLEDGE INTEGRATION

The knowledge retrieval phase is crafted to gather pertinent documents from various external sources in alignment with a specified query. Following this, the integration phase focuses on seamlessly synthesizing documents obtained from knowledge retrieval into a cohesive prompt, simultaneously setting appropriate training goals for the purpose of optimization. Recognizing that LLMs form a sturdy and foundational framework, the exploration during the integration phase predominantly steers clear of modifying the internal structure of individual layers or necessitating a complete reinitialization of LLMs for training. In this section, we delve into the comprehensive pipeline, relevant technologies, and strategies for enhancing knowledge integration.

### A. Overall Pipeline of Knowledge Integration

Integrating graph-retrieved knowledge into LLMs mainly includes two main ways: fine-tuning and in-context learning. The overall pipeline of these integration methods can be summarized as follows.

*Fine-tuning.* To directly leverage information retrieved from graph searches to enhance open-source LLMs, fine-tuning offers a straightforward solution for the integration, such as LoRA-based tuning [8], [161] and other data-efficient fine-tuning strategies [162], [163]. It injects the retrieved knowledge directly into the LLMs, focusing on graph-retrieved information at three knowledge levels: node-level knowledge, path-level knowledge, and subgraph-level knowledge for model tuning. In this way, graph information from different levels enhances the different capabilities of LLMs.

*In-context Learning.* While numerous open-source LLMs have been released to date, many state-of-the-art LLMs remain closed-source in practice. The integration of closed-source LLMs is constrained since it is not feasible to jointly train or fine-tune closed-source LLMs in an end-to-end manner. Thus, in-context learning provides an indirect strategy for knowledge integration, which can be roughly decomposed into two steps: prompt format choice and LLM response optimization. (i) Prompt format choice: The choice of prompt format is crucial for knowledge integration. This is because LLMs are highly sensitive to the prompt format. For example, the order of examples in in-context learning can lead to different responses [164], [165]. (ii) LLM response optimization: Then the appropriate prompt format joins the retrieved content and the questions as input to prompt the LLMs for response generation or further optimization.

### B. Integration Techniques

1) *Fine-tuning Techniques:* The fine-tuning process, leveraging various graph information, can be delineated into three distinct categories based on the granularity of the input target: (i) **Node-level Knowledge:** Focusing on individual nodes within the graph. (ii) **Path-level Knowledge:** Concentrating on the connections and sequences between nodes. (iii) **Subgraph-level Knowledge:** Considering larger structures composed of multiple nodes and their interconnections. We will explore each of these aspects in detail.

**Fine-tuning with Node-level Knowledge.** In many graph-based RAG systems, each node is linked to a document, such as an abstract in a citation network [168]. Since domain-specific data is seldom present in pre-training corpora [160], some studies employ instruction tuning to bolster domain-specific knowledge comprehension before proceeding with downstream task fine-tuning [169], [170]. A straightforward fine-tuning approach involves feeding node and neighboring text as contextual information into LLMs to aid in predictions [65], [157], [171]. Given that retrieved documents can be extensive, researchers can leverage LLMs to distill these texts into a single embedding [12]. Despite the absence of pre-training data for out-of-vocabulary tokens, LLMs are capable of discerning the information within these embeddings in practice [172]. To augment the information within these embeddings, alignment techniques are integrated to fuse multi-modal data, such as images [66].

**Fine-tuning with Path-level Knowledge.** Linguistic tasks often involve intricate reasoning and require a clear understanding of factual relationships [173]. Utilizing knowledge graph paths, LLMs are guided through the transitory relationships and entities, thereby enhancing their reasoning capabilities with evidence-based support [174]. These paths can either be the most direct routes from question entities to answer entities or be mined using graph retrieval models [167] or heuristic methods [69]. They can function as both input and output, but when multiple paths exist between two nodes, it is crucial to filter out noisy paths while preserving the relationships within the knowledge graph [67]. To maintain the integrity of entity representations and their relationships along the paths, certain approaches focus on using these paths as training objectives, predicting nodes and relations along a path between two nodes [68] or even across multiple paths [42]. This enables LLMs to engage in edge-level reasoning and produce reliable outputs.

**Fine-tuning with Subgraph-level Knowledge.** A subgraph represents a segment of a graph, encompassing a subset of nodes and edges from the original structure. Unlike the linear topology of path data, subgraph data exhibits a more complex, irregular topology [66]. This complexity arises because subgraphs can contain a multitude of connections beyond those found in paths, allowing them to capture intricate relationships between nodes and posing a greater challenge for LLMs to learn [73], [175]. One straightforward approach is to employ a graph encoder to condense subgraph-level information into a readout embedding [52], [71]. Alternatively, some researchers transform graph data into sequences, leveraging LLMs' inherent strength in processing sequential data, such as executable graph database queries [176]. However, these methods often overlook the rich textual content within subgraphs and fail to make LLMs cognizant of the underlying graph structure. To address this, efforts are divided: some focus on adapting the transformer architecture to better handle structured data [70], [177], [178], while others incorporate descriptions of nodes and edges directly into the prompt [72], [179]. However, there still remain challenges for existing methods. The former risks losing knowledge acquired during pre-training due to architectural alterations, while the latter may struggle with dense graphs featuring a large number of nodes and edges.

TABLE IV: Representative in-context learning knowledge integration techniques used in GraphRAG systems.

Methods	Applied LLMs	Zero-shot	Few-shot	CoT	Graph Query	Release Time
LARK [74]	LLaMA-2	✗	✗	✓	✓	May-2023
Chain-of-Knowledge [75]	GPT-3.5	✗	✓	✓	✗	May-2023
Think-on-Graph [22]	GPT-4, GPT-3.5, LLaMA-2	✗	✓	✓	✗	Jul-2023
MindMap [76]	GPT-3.5	✓	✓	✓	✗	Aug-2023
KnowledGPT [77]	GPT-4	✓	✗	✓	✓	Aug-2023
Align-NL2GQL [166]	Qwen, Baichuan2, ChatGLM	✗	✗	✗	✓	Feb-2024
Graph-CoT [63]	LLaMA-2, Mixtral, GPT-3.5	✗	✓	✓	✗	Apr-2024
GNN-RAG [167]	LLaMA-2, GPT-3.5	✓	✗	✗	✗	May-2024
FMEA-KG [78]	GPT-4	✗	✗	✗	✓	Jun-2024
KELP [44]	GPT-3.5	✗	✓	✗	✗	Jun-2024
CogMG [80]	Qwen	✗	✓	✗	✗	Jun-2024
EtD [79]	LLaMA-2	✓	✗	✗	✗	Jun-2024

2) *In-context Learning Techniques*: Recent research on knowledge integration through in-context learning (ICL) has introduced sophisticated frameworks that can be broadly divided into two stages: (i) **Graph-enhanced Chain-of-Thought**: This approach leverages graph structures to enhance the reasoning process within LLMs. (ii) **Collaborative Knowledge Graph Refinement**: This method highlights the distinct techniques used in optimizing LLM responses through collaborative refinement of well-built knowledge graphs.

**Graph-enhanced Chain-of-Thought**. Since Chain-of-Thought (CoT) has demonstrated effective performance improvement and scalability [180], CoT-based prompting techniques on knowledge integration also elicit strong potential in complex multi-hop reasoning. The Graph-enhanced CoT, such as Reasoning on Graphs [42], emphasizes chain reasoning based on the interaction between graphs and LLMs. Think-on-Graph [22], [23] proposes a reasoning framework based on the KG-LLMs interaction; LLMs are required to retrieve the relevant sub-graph of the proposed keyword first, then deduce the final result through chain reasoning. Similarly, in Graph CoT [63], given a question, LLM reasons the key component in the external graph that is required for the final answer, then uses a multi-round graph execution and interaction to finish the chain reasoning that concludes the final answer. Chain-of-Knowledge [75], firstly constructs exemplars that answer the given question based on evidence triples retrieved from the knowledge base, then utilizes faithfulness and factuality verification to check the reliability of the chain prediction. Program-of-Thought (PoT) [181], [182] is also a well-designed prompting technique that extends CoT prompting into coding scenarios. Following this, the KnowledGPT [77] framework uses PoT to generate the research language of knowledge bases to retrieve the answer step by step. The CoT-based methods achieve promising progress in knowledge integration for GraphRAG, significantly improving the reasoning capability of complex questions for LLMs. However, error accumulation is still a severe problem in this ICL paradigm [183], [184]. To address this, research is being conducted to diversify the reasoning paths that LLMs can take to reach the same conclusion [185], [186], which can help mitigate the effects of error propagation. Additionally, the integration of visual in-context learning for Large Vision-Language Models is being explored to enhance cross-modal interactions and representation alignment [187], [188], which may provide further insights

into managing error accumulation in reasoning tasks.

**Collaborative Knowledge Graph Refinement**. Instead of a well-designed prompting technique utilized to generate a better prompt or elicit LLMs for better reasoning, refining the LLM’s original response based on the factual knowledge in knowledge graphs is also an effective method to prevent LLMs from hallucination scenarios. The timeliness and accuracy of knowledge graphs are crucial, as they greatly influence the quality of augmented generation [63], [189], [190]. Maintaining and refining the content in knowledge graphs through feedback in LLM generation is an important strategy in GraphRAG. Representatively, KG-based Retrofitting [43] is a framework that combines LLMs with KGs to mitigate hallucination during the reasoning process. It retrofits the initial draft responses of LLMs based on the factual knowledge stored in KGs with an autonomous knowledge verifying and refining procedure. Similarly, KELP [44] refines a trainable encoder for path selection based on the LLM’s response. The Explore-then-Determine framework [79] uses the integration of LLMs and KGs to determine the final answer. CogMG [80] has designed a framework for collaborative augmentation between LLMs and KGs, leveraging KGs to augment LLMs in generations, explicitly targeting incomplete knowledge. LLMs are then required to identify and decompose required knowledge triples that are not present in the KG, enriching them and aligning updates with real-world demands. The collaboration between KGs and LLMs aligns factual knowledge to ensure the accuracy of generation and the up-to-date quality of KGs. However, ensuring the correctness of the intermediate process is challenging, and the accuracy of the refining process should also be evaluated with tailored metrics [191]–[193].

### C. Integration Enhancement Strategies.

1) *Training with other domain-specific models*: Training with domain-specific models is a critical approach to enhance the capabilities of LLMs in handling multi-modal information representations and to improve their performance on specific tasks. (i) **Enhancing Multi-modal Information Processing**: Integrating multi-modal information representations, such as image representations learned by vision-language models [194], into LLMs, is a significant advancement. These representations not only reduce input lengths but also encapsulate rich information from specific inputs. However, they also pose challenges for LLMs, as they may be treated as out-of-vocabulary tokens,



hindering the full exploitation of the original multi-modal information. Furthermore, recent efforts demonstrate the large number of additional tokens especially in the extremely sparse case where each additional token is only paired with a few training samples can cause LLMs to perform even worse over the small LMs [195]. To address this, recent efforts have focused on fine-tuning LLMs with domain-specific models to increase their generalization capabilities across different domains [196]. (ii) *Fine-tuning LLM with Domain-specific Task*: Fine-tuning an LLM involves using a pre-trained model and refining its weights by training it on a small set of annotated data with a slow learning rate. This principle allows the language model to adopt new knowledge from the data while retaining its initial learnings [197]. Transfer learning, a technique that allows a pre-trained model to apply its knowledge to a new task, is particularly useful when sufficient datasets for fine-tuning are not available [198]. For instance, MedPaLM [199] is a domain-specific model built upon PaLM, which demonstrated exceptional performance in complex tasks by using prompting strategies with annotated medical questions and answers.

2) *Multi-round Integration*: The knowledge integration of LLMs through a single round of interaction often falls short of providing satisfactory answers. To address this issue, researchers have been exploring multi-round integration techniques that enhance the quality of responses from LLMs for the following reasons: (i) *Enhancing Retrieved Content Quality*: One of the primary reasons LLMs may fail to answer questions effectively in a single round is the inconsistency in the quality of retrieved content. To mitigate this, multi-round integration allows for iterative refinement of the retrieved information. For instance, the IM-RAG approach [200] integrates retrieval systems with LLMs to support multi-round RAG through learning Inner Monologues, which are akin to the human inner voice that narrates one’s thoughts. (ii) *Handling Multi-hop Complex Reasoning*: Complex questions often require multi-hop reasoning, which can be challenging for LLMs due to the complexity involved. To address this, multi-round integration enables the generation of intermediate reasoning steps, leading to more accurate reasoning results. The study [201] comprehensively analyzed the latent multi-hop reasoning capabilities of LLMs. They found that while LLMs can exhibit multi-hop reasoning, their performance is significantly influenced by the structure of the prompt and the relational information within. (iii) *Aligning Output with Target Labels*: LLMs’ output may not always meet specific requirements, such as aligning with target labels. To overcome this, multi-round integration can help in aligning the model’s output with the desired outcomes. The X-InSTA method [202] proposes a cross-lingual in-context source-target alignment strategy that aligns prompt examples in a cross-lingual scenario. (iv) *Multi-round Interaction for Quality Assurance*: Multi-round integration can better ensure the stability and quality of the answers. For example, AgentPS [203] is another framework that integrates agentic process supervision into LLMs via multi-round question answering during fine-tuning, highlighting the effectiveness of integrating process supervision and structured sequential reasoning for multimodal content quality assurance.

#### D. Discussion.

The knowledge integration allows the users to inject domain-specific knowledge into LLMs. When integrating, the following aspects need to be considered for the design:

1) *Computational Overhead*: Incorporating retrieved text into the input consumes a portion of the available input length. The computational overhead may increase further due to the complexity of handling graph-structured data from graph-retrieved information. GraphRAG combines well-built graphs with LLMs, which requires additional processing to retrieve and manage graph elements such as nodes, triples, and paths. This added complexity can lead to much higher costs. Given that prompt processing time scales quadratically with the length of the prompt, large retrieved texts can significantly amplify the computational overhead for LLMs, especially in downstream applications like billion-scale recommender systems [204], [205]. To mitigate this, domain-specific knowledge should be permanently integrated into LLMs through fine-tuning, enhancing their responses, and tailoring them to specific applications. This approach could lead to specialized LLMs that reduce reliance on extensive text retrieval from external sources, thereby improving response speed and reducing both time and financial costs [206], [207]. For instance, deploying GPT-4 for pilot-scale customer service could incur costs in the thousands of dollars, while fine-tuning a large language model might require hundreds of gigabytes of memory. The practical deployment of LLMs in real-world scenarios is thus heavily influenced by these cost considerations.

2) *Graph-structured Input*: LLMs are not inherently equipped to process graph-structured data. The common practice is to convert graph data into natural language descriptions, such as explaining the connections between nodes [208], [209]. However, these descriptions alone are insufficient for LLMs to grasp the complex geometric properties of graphs. Graph-structured data is a fundamental input format for numerous tasks, including node classification and link prediction in text-attributed networks. The success of LLMs in graph-related tasks depends on their capacity to understand and process graph-structured inputs. As graph reasoning tasks grow in complexity, such as calculating the minimum cut problem for a randomly generated graph, the importance of accurate input modeling becomes even more critical.

3) *Large-scale Graph*: The scalability of models is crucial for dealing with real-world graphs [210], [211]. Handling graph-structured input becomes particularly challenging with large graphs, where the number of edges and nodes may surpass the input length limitations of LLMs. Therefore, efficiently managing large-scale graph problems within the constraints of LLMs is a significant area of research [212]. Advances in this field could enable LLMs to handle more complex and larger graphs, which is essential for real-world applications in fields such as social network analysis, bioinformatics, and knowledge graph construction.

## VII. LIMITATIONS AND FUTURE OPPORTUNITIES

In this section, we systematically analyze several critical limitations of existing GraphRAG systems regarding knowledge quality, knowledge conflict, data privacy, and efficiency

TABLE V: Open-source Projects and Applications.

Model	Domain	Task	Dataset	Background Knowledge
Graph RAG [25]	General Domain	Graph Construction, QA	-	Podcast Transcripts, News Articles
ToG [22]	General Domain	QA	WebQSP, CWQ	Freebase
ToG 2.0 [23]	General Domain	QA	WebQSP, CWQ	Freebase
RoG [42]	General Domain	QA	WebQSP, CWQ	Freebase
KnowGPT [7]	General Domain	QA	CommenconseQA, OpenbookQA	ConceptNet
SubgraphRAG [46]	General Domain	QA	WebQSP, CWQ	Freebase
NQ-RAG [58]	General Domain	Graph Construction, QA	NQ	Wikipedia
LBR-GNN [213]	General Domain	QA	CommonsenseQA, OpenbookQA	Wikipedia, ConceptNet
SG-RAG [214]	General Domain	QA	MetaQA, LC-QuAD, ComplexWebQuestions	Freebase
StructRAG [14]	General Domain	QA	Loong, Podcast Transcripts	-
R-DP [53]	General Domain	QA	WebQSP, CWQ	Freebase
GAIL [215]	General Domain	QA	WebQSP, CWQ, GrailQA	Freebase
KG-RAG [126]	Biomedical	QA	MCQ	SPOKE
MEG [59]	Medical	QA	MedQA, PubMedQA, MedMCQA	UMLS
DialogGSR [62]	Medical	Dialog Generation	OpenDialogKG, KOMODIS	Freebase, IMDB
ReTek [56]	Medical	Graph Construction, QA	ReTek	PharmKG, ADInt
KG4Diagnosis [216]	Medical	Graph Construction, Dialog Generation	-	SNOMED-CT, UMLS
CancerKG [51]	Medical	QA	CancerKG	PrimeKG, PubMed
MedGraphRAG [48]	Healthcare, Medical	QA	MultiMedQA, DiverseHealth	MedC-K, UMLS
AGENTiGraph [217]	Legislation, Healthcare	Graph Construction, QA	-	UK Legislation, MMedC
KGL [55]	Water Conservancy	Graph Construction, QA	KGQS	Special Report
SURGE [57]	Movie	Dialog Generation	OpendialKG, KOMODIS	Freebase, IMDB
TQA-KG [60]	Education	Graph Construction, QA	TQA-KG	CK12-QA, AI2D
GraphFusion [218]	Scientific Research	Graph Construction, QA	TutorQA	NLP KG, TutorialBank, NLP-Papers
StructuGraphRAG [50]	Scientific Research	Graph Construction, QA	NSDUH	Codebook
Soccer-GraphRAG [219]	Soccer	Graph Construction, QA	Soccer KG	SoccerNet-Echoes
LightRAG [26]	Cross Domain	Graph Construction, QA	-	UltraDomain
PG-RAG [32]	Cross Domain	Graph Construction, QA	Mindmap	CRUD
DiaKoP [136]	Cross Domain	Dialog Generation	ConvQuestions	Wikidata
Graph-CoT [63]	Cross Domain	QA	-	-
G-Retriever [52]	Mutiple Domain	Graph Construction, QA	GraphQA	ExplaGraphs, SceneGraphs, WebQSP

issues, and propose targeted research directions for practical advancements.

1) *Knowledge Quality*: The effectiveness of GraphRAG models fundamentally depends on the quality of the external knowledge, necessitating the development of sophisticated mechanisms for knowledge engineering. This encompasses advanced techniques for (i) systematic knowledge organization, (ii) automated quality refinement, and (iii) intelligent knowledge base expansion. First, knowledge organization demands more expressive graph structures that capture complex semantic relationships, temporal dynamics, and hierarchical dependencies through hybrid neuro-symbolic approaches and advanced embedding techniques. Second, knowledge refinement requires automated quality assurance frameworks that leverage cross-validation, statistical analysis, and machine learning to identify inconsistencies, remove redundancies, and validate factual accuracy. Third, knowledge expansion is critical in enhancing GraphRAG systems’ practical effectiveness through continuous knowledge enrichment and adaptation. The practical implementation of knowledge expansion typically combines multiple approaches: automated web crawling for public information updates, API integrations with authoritative databases, expert feedback loops for validation, and machine learning models for relationship inference. This multi-faceted approach ensures robust and reliable knowledge growth while maintaining data quality and relevance. The seamless integration of these components, supported by a scalable infrastructure for dynamic updates and version control, will be crucial for developing robust GraphRAG systems across various applications.

Beyond traditional text-based knowledge, integrating multi-modal information, including images and videos, offers promising opportunities to enrich the external knowledge databases with more comprehensive domain expertise and conceptual understanding. Equally crucial is developing robust knowledge

validation frameworks to ensure reliability and maintain data integrity across different modalities. These frameworks should incorporate advanced techniques from anomaly detection, knowledge base completion, and alignment to systematically identify and correct errors, address information gaps, and maintain consistency across the integrated knowledge base. Such comprehensive approaches to knowledge quality assurance will be fundamental in enhancing the reliability and effectiveness of GraphRAG systems across diverse applications.

2) *Knowledge Conflict*: Integrating multiple knowledge sources in GraphRAG models introduces significant challenges in managing conflicting information and maintaining knowledge consistency. As such, another critical research priority is developing sophisticated techniques for knowledge reconciliation, conflict resolution, and truth discovery. These approaches should incorporate advanced algorithms capable of identifying contradictory statements, evaluating source reliability, and determining the most probable accurate information based on available evidence and contextual facts. Beyond conflict resolution, ensuring seamless alignment between external knowledge and LLM-generated outputs presents another fundamental challenge. This alignment requires innovative methods for knowledge distillation, fine-tuning, and cross-modal integration to harmonize structured knowledge with the LLM’s learned representations. Incorporating uncertainty modeling and probabilistic reasoning frameworks offers a promising direction for handling ambiguous or conflicting information more effectively. By associating knowledge with confidence scores and probability distributions, GraphRAG systems can make more nuanced decisions and generate outputs that accurately reflect the inherent uncertainty in the integrated knowledge. This comprehensive approach to knowledge consistency, combining conflict resolution, alignment techniques, and uncertainty modeling, will be essential for developing more

robust and reliable GraphRAG systems capable of handling real-world information complexity.

3) *Data Privacy*: The integration of external knowledge in GraphRAG systems raises critical privacy concerns that demand sophisticated technical solutions and robust governance frameworks. Privacy-preserving knowledge integration and retrieval represent critical challenges requiring advanced cryptographic approaches, including secure multi-party computation, homomorphic encryption, and differential privacy mechanisms. These techniques enable GraphRAG systems to leverage sensitive information from knowledge graphs while maintaining strict privacy guarantees for individuals and organizations. Beyond technical solutions, developing comprehensive data governance frameworks becomes essential for responsible deployment. Such frameworks should establish clear standards for access control, data handling protocols, and ethical guidelines that align with evolving privacy regulations while maintaining system utility. This dual focus on technical privacy preservation and governance frameworks will be crucial in building trustworthy GraphRAG systems that balance the competing demands of knowledge accessibility and privacy protection, ultimately enabling broader adoption across sensitive domains such as healthcare, finance, and personal information management.

4) *Efficiency*: The practical deployment of GraphRAG systems in real-world applications hinges critically on model efficiency. Future work should investigate techniques for optimizing the knowledge retrieval and integration processes to reduce computational overhead and memory requirements. This may involve exploring knowledge distillation, pruning, and compression methods to create more compact and efficient representations of domain knowledge. Developing scalable and efficient algorithms for subgraph matching, graph traversal, and reasoning is another important direction for future work. By leveraging techniques from graph theory, database systems, and parallel computing, researchers can enable faster inference times and real-time generation of responses in GraphRAG models. Furthermore, investigating hardware acceleration techniques, such as utilizing GPUs and TPUs, can significantly boost the performance of GraphRAG models. Exploiting the parallelism and computational power of specialized hardware can enable the efficient processing of large-scale knowledge graphs and support the integration of extensive knowledge into language models. Through advancements in these research directions, future research can significantly improve the practical applicability of GraphRAG systems by reducing computational overhead, memory requirements, and response latency while maintaining or enhancing system performance.

## VIII. OPEN-SOURCE PROJECTS AND APPLICATIONS

The advent of GraphRAG has opened new avenues for enhancing information retrieval and generation processes through the integration of graph-based methodologies. As researchers and developers explore the potential of GraphRAG, a variety of open-source projects and applications have emerged, demonstrating its versatility across different domains. This section delves into the real-world implementation of GraphRAG systems, including benchmark datasets, that serve as evaluation

standards and downstream tasks that gauge their effectiveness in real-world scenarios. Additionally, it highlights notable open-source projects that exemplify the implementation of GraphRAG principles, paving the way for innovative applications in areas such as chatbots, knowledge management, and research development.

### A. Benchmark Datasets

We categorize representative datasets based on their complexity and specific characteristics:

#### 1) *Simple Question Answering*:

- **SimpleQuestion** [220] consists of 100k questions constructed from the Freebase knowledge graph. The dataset emphasizes basic evidence retrieval without requiring complex reasoning chains. Its straightforward nature makes it ideal for evaluating baseline performance and validating fundamental retrieval mechanisms in KGQA systems. The questions are designed to test direct fact extraction, making it a valuable resource for assessing retrieval accuracy.
- **WebQ** [221] comprises 4,737 questions, split into 3,098 training and 1,639 testing examples. The questions were collected through Google Suggest API, capturing natural user queries. Each question is annotated with SPARQL queries, enabling systematic evaluation of semantic parsing capabilities. The dataset leverages the Freebase knowledge graph, providing a realistic test bed for question-answering systems dealing with real-world information needs.

#### 2) *Multi-hop Reasoning*:

- **CWQ** [222] extends WebQSP to 34,689 questions requiring sophisticated reasoning patterns. The dataset is carefully balanced across four question types: composition (45%) involving multiple-step reasoning chains, conjunction (45%) combining multiple constraints, comparative (5%) handling numerical or temporal comparisons, and superlative (5%) dealing with maximum or minimum value queries. Questions require up to 4-hop reasoning paths, making it a comprehensive benchmark for evaluating complex reasoning capabilities over the Freebase knowledge graph.
- **MetaQA** [223] specializes in movie domain knowledge, containing over 400k questions structured around a knowledge graph with 135k triples, 43k entities, and 9 relations. The dataset is organized into progressive difficulty levels based on reasoning hop requirements. Each question includes detailed annotations of head entities, answers, and the entities involved in the reasoning path, providing a controlled environment for evaluating multi-hop reasoning capabilities.
- **MetaQA-3** [223] focuses specifically on 3-hop questions, containing over 100k examples built on the WikiMovies dataset. The knowledge graph encompasses 43k entities and 135k triples. The training set typically uses 1,000 questions (1% of the total), making it particularly suitable for evaluating models' ability to generalize from limited training data in complex reasoning scenarios.

### 3) Large-scale Complex QA:

- **LC-QuAD** [224] contains 5,000 question-SPARQL pairs based on DBpedia 2016. The dataset supports both public endpoint access and local query execution, making it versatile for different evaluation setups. It focuses on testing systems' ability to generate and execute complex SPARQL queries accurately.
- **LC-QuAD v2** [225] expands to 30,000 question-query pairs, supporting both Wikidata and DBpedia 2018. The dataset provides 22,792 entity QIDs and 3,627 predicates, offering broader coverage and increased complexity compared to its predecessor. Its dual compatibility enables comprehensive evaluation across different knowledge graphs.
- **KQAPro** [226] represents the largest KGQA dataset with 94,376 training and 11,797 validation/test examples. Built on a dense subset of Wikidata, it features multiple inference types and logical operations including unions and intersections. Each example provides both SPARQL queries and KoPL logical forms, with entities and predicates represented in their natural form rather than ID-based notation.

### 4) Domain-specific QA:

- **Mintaka** [227] contains 20k questions across 8 different languages, focusing on logical refinement and answer revision capabilities. The dataset emphasizes cross-lingual reasoning and the ability to systematically improve answers based on available evidence.
- **FACTKG** [228] encompasses 108,000 claims requiring validation against DBpedia. The binary classification task tests systems' ability to verify factual accuracy using knowledge graph information, providing a unique perspective on knowledge validation capabilities.
- **WebQSP** [229] features 4,737 natural language questions operating over a vast knowledge graph containing 164.6M facts and 24.9M entities. The questions are distributed across different reasoning types: 30% require two-hop aggregation, 7% involve constraint reasoning, and 63% need single-fact retrieval.
- **GraiiQA** [230] focuses on complex question answering over knowledge bases, incorporating advanced reasoning patterns. The dataset challenges systems to demonstrate sophisticated query understanding and execution capabilities, pushing the boundaries of KGQA system development.
- **TutorQA** [218] is a dataset specifically designed for research in educational question answering and tutoring systems. It consists of questions and answers extracted from real-world tutoring interactions.
- **CRUD** [231] is constructed by collecting recent high-quality news articles from major Chinese news websites, which provides a fresh and challenging benchmark for RAG systems.
- **UltraDomain** [232] is a benchmark for evaluating RAG systems on complex queries. It includes 428 college textbooks across 18 domains, such as agriculture, humanities, and computer science, covering topics like machine learning, and big data.

### B. GraphRAG Applications

GraphRAG systems have emerged as a powerful framework for real-world applications, offering robust solutions to challenges across diverse domains by leveraging graph-based representations of knowledge. These systems excel in tasks like QA, dialog generation and other natural language processing tasks, where they integrate structured and unstructured data to enable advanced reasoning and contextual understanding. Their versatility stems from their ability to utilize a wide range of datasets and background knowledge, adapting to the specific needs of various fields.

In the **general domain**, GraphRAG models such as GraphRAG [25], SubgraphRAG [46], and StructRAG [14] are designed to handle tasks that require broad contextual understanding. These models utilize extensive knowledge sources, including Freebase, Wikipedia, ConceptNet, and even unstructured data such as podcast transcripts and news articles. By constructing graphs that represent relationships between entities, these systems are capable of answering complex queries that go beyond simple fact retrieval. For example, GraphRAG models can synthesize information from different domains to address multifaceted questions, making them highly effective for general-purpose knowledge exploration and reasoning.

In the **biomedical** and **medical domains**, GraphRAG systems have proven particularly impactful. Models like KG-RAG [126], MEG [59], and MedGraphRAG [48] are tailored to handle the intricate and high-stakes nature of medical data. They rely on specialized datasets such as PubMedQA, MedQA, and UMLS, which provide rich sources of structured medical knowledge. By constructing graphs that capture relationships between medical concepts, these models can assist in a range of tasks, including medical QA, diagnostic support, and the integration of healthcare data. For instance, KG-RAG leverages knowledge graphs to link symptoms, diseases, and treatments, enabling healthcare professionals to access accurate and interpretable insights. Similarly, MedGraphRAG addresses the need for personalized healthcare solutions by integrating multimodal data from sources like MultiMedQA and DiverseHealth.

In the **legislation area**, GraphRAG systems like AGENTiGraph [217] bridge the gap between regulatory requirements and practical applications. AGENTiGraph focuses on integrating legislative and healthcare-related knowledge graphs to address domain-specific QA tasks. By utilizing datasets such as UK legislation and MMedC (medical corpora), this model enables the navigation of complex regulatory environments, such as understanding compliance requirements in healthcare or analyzing the implications of new policies. This integration of structured legislative data and healthcare knowledge allows AGENTiGraph to support decision-making processes in areas where legal and medical considerations intersect.

In the realm of **education** and **scientific research**, GraphRAG systems like TQA-KG [60] and GraphFusion [218] play a critical role in facilitating knowledge dissemination and academic inquiry. These models are designed to process datasets such as CK12-QA, AI2D, and TutorQA, which provide educational content in structured formats. By constructing



knowledge graphs that represent interconnected concepts, these systems enable students, educators, and researchers to access relevant information efficiently. For example, TQA-KG can assist in answering educational queries by linking concepts across subjects, while GraphFusion integrates scientific literature and tutorial resources to support researchers in navigating complex academic topics. For scientific research, StructGraphRAG [50] organizes and processes research papers, integrating data from scientific corpora to create graphs that highlight relationships between key findings, methodologies, and datasets. This capability aids researchers in identifying trends and synthesizing new insights.

Beyond these domains, GraphRAG systems demonstrate their adaptability to specialized fields. In **water conservancy**, models like KGL [55] utilize datasets such as KGQS to create domain-specific knowledge graphs that facilitate the management and analysis of water resources. These graphs enable stakeholders to identify relationships between hydrological data, infrastructure, and policies, improving decision-making processes. In the domain of **sports analytics**, Soccer-GraphRAG [219] applies graph-based reasoning to soccer data, utilizing datasets like SoccerNet-Echoes to analyze player performance, match statistics, and team dynamics. This capability supports applications such as player scouting and strategy development. In the **entertainment industry**, systems like SURGE [57] are used to construct knowledge graphs centered on movie-related data. These graphs enhance applications such as content recommendation, audience analysis, and storytelling, enabling deeper engagement with viewers.

Overall, GraphRAG systems demonstrate unparalleled versatility by adapting their graph-based frameworks to meet the unique demands of various real-world domains. Their ability to integrate structured and unstructured data into cohesive knowledge representations enables them to address complex problems across general and highly specialized fields. By constructing graph bases tailored to specific applications, these systems contribute to advancements in healthcare, education, legislation, scientific research, sports analytics, and other specific domains. Their impact lies in their ability to provide interpretable, context-aware insights that empower decision-making and drive innovation.

### C. Open-source Project

GraphRAG has inspired a range of open-source projects that explore its principles and adapt its framework to various use cases. These implementations highlight the versatility and growing interest in combining knowledge graphs with RAG techniques to enhance LLMs.

**Microsoft GraphRAG**<sup>1</sup> serves as the foundational framework that combines knowledge graphs with RAG to enhance the performance of LLMs. It constructs a knowledge graph from private datasets, leveraging graph machine learning to enrich the query process through prompt augmentation at runtime. By doing so, Microsoft GraphRAG achieves remarkable improvements in answering complex queries that require detailed reasoning or domain-specific mastery. Its ability

to outperform traditional methods demonstrates its effectiveness in handling private datasets where accuracy and contextual understanding are critical.

Building upon this foundation, a variety of open-source projects have emerged to explore and adapt GraphRAG’s principles for diverse needs. One such project is **Nano-GraphRAG**<sup>2</sup>, a lightweight and highly customizable implementation designed with simplicity in mind. Featuring a concise codebase of approximately 1,100 lines, Nano-GraphRAG is accessible to developers aiming to prototype or deploy graph-enhanced LLM systems without the overhead of a full-scale solution. It supports multiple graph storage systems, such as Neo4j, and integrates seamlessly with vector databases like Milvus and Faiss, offering flexibility for a range of applications.

**Azure GraphRAG**<sup>3</sup> extends the concept by offering a solution accelerator built on the graphrag Python package. This project provides API endpoints hosted on Azure, allowing users to trigger indexing pipelines and query knowledge graphs seamlessly. Azure GraphRAG focuses on demonstrating how knowledge graph memory structures can enhance LLM outputs in a hosted service environment. While not an official Microsoft product, it serves as a practical toolkit for developers and researchers to experiment with and deploy GraphRAG-based solutions, bridging the gap between research concepts and real-world applications.

**Fast GraphRAG**<sup>4</sup> takes a performance-oriented approach to the GraphRAG concept. By leveraging asynchronous operations and parallelized graph querying, it prioritizes speed and efficiency, making it well-suited for real-time applications that demand rapid data retrieval and reasoning. This implementation highlights the importance of optimization for time-sensitive domains, where traditional graph traversal methods might fall short.

In contrast, **LightRAG**<sup>5</sup> focuses on reducing computational complexity and resource requirements, providing a streamlined solution for environments with limited infrastructure. Its lightweight architecture makes it ideal for deployment on edge devices or within constrained systems, such as IoT networks or small-scale reasoning platforms.

Another notable adaptation is **Medical GraphRAG**<sup>6</sup>, a domain-specific variant tailored to the healthcare industry. This implementation integrates medical knowledge graphs, such as SNOMED and ICD, to provide contextually accurate and clinically relevant outputs. By bridging the gap between medical terminology and generative AI capabilities, Medical GraphRAG facilitates applications like clinical decision support and patient education, where precision and reliability are paramount.

These open-source projects, along with the original Microsoft GraphRAG, collectively demonstrate the adaptability and growing importance of integrating structured knowledge graphs with LLMs. They highlight the potential of this approach to create intelligent, context-aware systems that address complex challenges across a wide array of industries and domains.

<sup>2</sup><https://github.com/gusye1234/nano-graphrag.git>

<sup>3</sup><https://github.com/Azure-Samples/graphrag-accelerator.git>

<sup>4</sup><https://github.com/circlemind-ai/fast-graphrag.git>

<sup>5</sup><https://github.com/HKUDS/LightRAG.git>

<sup>6</sup><https://github.com/SuperMedIntel/Medical-Graph-RAG.git>

<sup>1</sup><https://github.com/microsoft/graphrag.git>

## IX. CONCLUSION

LLMs have demonstrated impressive capabilities in natural language processing, achieving remarkable performance across a wide range of tasks. However, their effectiveness diminishes significantly in specialized domains, especially when faced with knowledge-intensive tasks requiring domain expertise that have almost never appeared in the pre-training corpus. RAG has emerged as a promising solution that enhances LLMs with external knowledge bases to improve their domain-specific capabilities. However, traditional RAG systems encounter several critical challenges in customizing LLMs for specific domains, including (i) efficient processing of extensive document collections, (ii) effective integration of knowledge dispersed across multiple sources, and (iii) maintenance of contextual coherence during information retrieval and generation. These limitations become particularly critical in applications requiring complex reasoning, where successful inference depends on both comprehensive knowledge integration and deep contextual understanding. To address these challenges, GraphRAG emerges as a pioneering approach that enriches LLMs with well-organized background knowledge and clear reasoning chains. This survey provides a comprehensive analysis of GraphRAG, detailing its taxonomy, mechanisms, challenges, and future research directions. We also provide guidelines for building domain-specific GraphRAG systems, assisting in framework selection for various use cases in practical scenarios.

## REFERENCES

- [1] OpenAI, “Gpt-4 technical report,” 2023.
- [2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [3] D. Khashabi, S. Min, T. Khot, A. Sabharwal, O. Tafjord, P. Clark, and H. Hajishirzi, “Unifiedqa: Crossing format boundaries with a single qa system,” *arXiv preprint arXiv:2005.00700*, 2020.
- [4] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, “Palm: Scaling language modeling with pathways,” *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.
- [5] Q. Zhang, J. Dong, H. Chen, W. Li, F. Huang, and X. Huang, “Structure guided large language model for sql generation,” *arXiv preprint arXiv:2402.13284*, 2024.
- [6] Z. Hong, Z. Yuan, Q. Zhang, H. Chen, J. Dong, F. Huang, and X. Huang, “Next-generation database interfaces: A survey of llm-based text-to-sql,” *arXiv preprint arXiv:2406.08426*, 2024.
- [7] Q. Zhang, J. Dong, H. Chen, D. Zha, Z. Yu, and X. Huang, “Knowgpt: Knowledge graph based prompting for large language models,” in *Conference on Neural Information Processing Systems*, 2024.
- [8] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations*, 2022.
- [9] J. He, C. Zhou, X. Ma, T. Berg-Kirkpatrick, and G. Neubig, “Towards a unified view of parameter-efficient transfer learning,” in *International Conference on Learning Representations*, 2022.
- [10] Z. Gekhman, G. Yona, R. Aharoni, M. Eyal, A. Feder, R. Reichart, and J. Herzig, “Does fine-tuning llms on new knowledge encourage hallucinations?” *arXiv preprint arXiv:2405.05904*, 2024.
- [11] Y. Zhai, S. Tong, X. Li, M. Cai, Q. Qu, Y. J. Lee, and Y. Ma, “Investigating the catastrophic forgetting in multimodal large language model fine-tuning,” in *Conference on Parsimony and Learning*, 2024, pp. 202–227.
- [12] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, “Retrieval-augmented generation for knowledge-intensive nlp tasks,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [13] R. Chen, X. Zhang, J. Wu, W. Fan, X.-Y. Wei, and Q. Li, “Multi-level querying using a knowledge pyramid,” *arXiv preprint arXiv:2407.21276*, 2024.
- [14] Z. Li, X. Chen, H. Yu, H. Lin, Y. Lu, Q. Tang, F. Huang, X. Han, L. Sun, and Y. Li, “Structrag: Boosting knowledge intensive reasoning of llms via inference-time hybrid information structurization,” *arXiv preprint arXiv:2410.08815*, 2024.
- [15] M. Glass, G. Rossiello, M. F. M. Chowdhury, A. R. Naik, P. Cai, and A. Gliozzo, “Re2g: Retrieve, rerank, generate,” *arXiv preprint arXiv:2207.06300*, 2022.
- [16] F. Xu, W. Shi, and E. Choi, “Recomp: Improving retrieval-augmented llms with compression and selective augmentation,” *arXiv preprint arXiv:2310.04408*, 2023.
- [17] A. Asai, Z. Wu, Y. Wang, A. Sil, and H. Hajishirzi, “Self-rag: Learning to retrieve, generate, and critique through self-reflection,” *arXiv preprint arXiv:2310.11511*, 2023.
- [18] X. Tang, Q. Gao, J. Li, N. Du, Q. Li, and S. Xie, “Mba-rag: a bandit approach for adaptive retrieval-augmented generation through question complexity,” *arXiv preprint arXiv:2412.01572*, 2024.
- [19] P. Sarthi, S. Abdullah, A. Tuli, S. Khanna, A. Goldie, and C. D. Manning, “Raptor: Recursive abstractive processing for tree-organized retrieval,” *arXiv preprint arXiv:2401.18059*, 2024.
- [20] C. Ling, X. Zhao, J. Lu, C. Deng, C. Zheng, J. Wang, T. Chowdhury, Y. Li, H. Cui, X. Zhang *et al.*, “Domain specialization as the key to make large language models disruptive: A comprehensive survey,” *arXiv preprint arXiv:2305.18703*, 2023.
- [21] T. Bruckhaus, “Rag does not work for enterprises,” *arXiv preprint arXiv:2406.04369*, 2024.
- [22] J. Sun, C. Xu, L. Tang, S. Wang, C. Lin, Y. Gong, L. Ni, H.-Y. Shum, and J. Guo, “Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph,” in *International Conference on Learning Representations*, 2024.
- [23] S. Ma, C. Xu, X. Jiang, M. Li, H. Qu, and J. Guo, “Think-on-graph 2.0: Deep and interpretable large language model reasoning with knowledge graph-guided retrieval,” *arXiv preprint arXiv:2407.10805*, 2024.
- [24] A. Anthropic, “The claude 3 model family: Opus, sonnet, haiku,” *Claude-3 Model Card*, 2024.
- [25] D. Edge, H. Trinh, N. Cheng, J. Bradley, A. Chao, A. Mody, S. Truitt, and J. Larson, “From local to global: A graph rag approach to query-focused summarization,” *arXiv preprint arXiv:2404.16130*, 2024.
- [26] Z. Guo, L. Xia, Y. Yu, T. Ao, and C. Huang, “Lightrag: Simple and fast retrieval-augmented generation,” *arXiv preprint arXiv:2410.05779*, 2024.
- [27] L. Liang, M. Sun, Z. Gui, Z. Zhu, Z. Jiang, L. Zhong, Y. Qu, P. Zhao, Z. Bo, J. Yang *et al.*, “Kag: Boosting llms in professional domains via knowledge augmented generation,” *arXiv preprint arXiv:2409.13731*, 2024.
- [28] B. Peng, Y. Zhu, Y. Liu, X. Bo, H. Shi, C. Hong, Y. Zhang, and S. Tang, “Graph retrieval-augmented generation: A survey,” *arXiv preprint arXiv:2408.08921*, 2024.
- [29] T. T. Procko and O. Ochoa, “Graph retrieval-augmented generation for large language models: A survey,” in *2024 Conference on AI, Science, Engineering, and Technology (AIxSET)*. IEEE, 2024, pp. 166–169.
- [30] H. Han, Y. Wang, H. Shomer, K. Guo, J. Ding, Y. Lei, M. Halappanavar, R. A. Rossi, S. Mukherjee, X. Tang *et al.*, “Retrieval-augmented generation with graphs (graphrag),” *arXiv preprint arXiv:2501.00309*, 2024.
- [31] Z. Li, Q. Guo, J. Shao, L. Song, J. Bian, J. Zhang, and R. Wang, “Graph neural network enhanced retrieval for question answering of llms,” *arXiv preprint arXiv:2406.06572*, 2024.
- [32] X. Liang, S. Niu, S. Zhang, S. Song, H. Wang, J. Yang, F. Xiong, B. Tang, C. Xi *et al.*, “Empowering large language models to set up a knowledge retrieval indexer via self-learning,” *arXiv preprint arXiv:2405.16933*, 2024.
- [33] Y. Wang, N. Lipka, R. A. Rossi, A. Siu, R. Zhang, and T. Derr, “Knowledge graph prompting for multi-document question answering,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, 2024, pp. 19 206–19 214.
- [34] W. Liu, A. Yu, D. Zan, B. Shen, W. Zhang, H. Zhao, Z. Jin, and Q. Wang, “Graphcoder: Enhancing repository-level code completion via code context graph-based retrieval and language model,” *arXiv preprint arXiv:2406.07003*, 2024.
- [35] Z. Hu, A. Iscen, C. Sun, K.-W. Chang, Y. Sun, D. Ross, C. Schmid, and A. Fathi, “Avis: Autonomous visual information seeking with large language model agent,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.

- [36] K. Rana, J. Haviland, S. Garg, J. Abou-Chakra, I. Reid, and N. Suen-derhauf, "Sayplan: Grounding large language models using 3d scene graphs for scalable robot task planning," in *7th Annual Conference on Robot Learning*, 2023.
- [37] W. Xiong, X. L. Li, S. Iyer, J. Du, P. Lewis, W. Y. Wang, Y. Mehdad, W.-t. Yih, S. Riedel, D. Kiela *et al.*, "Answering complex open-domain questions with multi-hop dense retrieval," *arXiv preprint arXiv:2009.12756*, 2020.
- [38] S. Min, D. Chen, L. Zettlemoyer, and H. Hajishirzi, "Knowledge guided text retrieval and reading for open domain question answering," *arXiv preprint arXiv:1911.03868*, 2019.
- [39] X. Lu, S. Pramanik, R. Saha Roy, A. Abujabal, Y. Wang, and G. Weikum, "Answering complex questions by joining multi-document evidence with quasi knowledge graphs," in *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, 2019, pp. 105–114.
- [40] S. Li, Y. He, H. Guo, X. Bu, G. Bai, J. Liu, J. Liu, X. Qu, Y. Li, W. Ouyang, W. Su, and B. Zheng, "GraphReader: Building graph-based agent to enhance long-context abilities of large language models," in *Findings of the Association for Computational Linguistics: EMNLP 2024*. Association for Computational Linguistics, 2024.
- [41] K. Cheng, N. K. Ahmed, T. Willke, and Y. Sun, "Structure guided prompt: Instructing large language model in multi-step reasoning by exploring graph structure of the text," *arXiv preprint arXiv:2402.13415*, 2024.
- [42] L. LUO, Y.-F. Li, R. Haf, and S. Pan, "Reasoning on graphs: Faithful and interpretable large language model reasoning," in *The Twelfth International Conference on Learning Representations*, 2024.
- [43] X. Guan, Y. Liu, H. Lin, Y. Lu, B. He, X. Han, and L. Sun, "Mitigating large language model hallucinations via autonomous knowledge graph-based retrofitting," in *AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 18 126–18 134.
- [44] H. Liu, S. Wang, Y. Zhu, Y. Dong, and J. Li, "Knowledge graph-enhanced large language models via path selection," in *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics, Aug. 2024.
- [45] M. Jin, H. Xue, Z. Wang, B. Kang, R. Ye, K. Zhou, M. Du, and Y. Zhang, "Prollm: Protein chain-of-thoughts enhanced llm for protein-protein interaction prediction," *bioRxiv*, pp. 2024–04, 2024.
- [46] M. Li, S. Miao, and P. Li, "Simple is effective: The roles of graphs and large language models in knowledge-graph-based retrieval-augmented generation," *arXiv preprint arXiv:2410.20724*, 2024.
- [47] H. Zhang, T. Feng, and J. You, "Graph of records: Boosting retrieval augmented generation for long-context summarization with graphs," *arXiv preprint arXiv:2410.11001*, 2024.
- [48] J. Wu, J. Zhu, and Y. Qi, "Medical graph rag: Towards safe medical large language model via graph retrieval-augmented generation," *arXiv preprint arXiv:2408.04187*, 2024.
- [49] X. Liu, B. Lan, Z. Hu, Y. Liu, Z. Zhang, W. Zhou, F. Wang, and M. Shieh, "Codexgraph: Bridging large language models and code repositories via code graph databases," *arXiv preprint arXiv:2408.03910*, 2024.
- [50] X. Zhu, X. Guo, S. Cao, S. Li, and J. Gong, "Structgraphrag: Structured document-informed knowledge graphs for retrieval-augmented generation," in *Proceedings of the AAAI Symposium Series*, vol. 4, no. 1, 2024, pp. 242–251.
- [51] M. Gubanov, A. Pyayt, and A. Karolak, "Cancerkg.org-a web-scale, interactive, verifiable knowledge graph-llm hybrid for assisting with optimal cancer treatment and care," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 4497–4505.
- [52] X. He, Y. Tian, Y. Sun, N. V. Chawla, T. Laurent, Y. LeCun, X. Bresson, and B. Hooi, "G-retriever: Retrieval-augmented generation for textual graph understanding and question answering," *arXiv preprint arXiv:2402.07630*, 2024.
- [53] Y. Huang and G. Zeng, "Rd-p: A trustworthy retrieval-augmented prompter with knowledge graphs for llms," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 942–952.
- [54] Z. Chen, C. Xu, D. Wang, Z. Huang, Y. Dou, and J. Guo, "Rulerag: Rule-guided retrieval-augmented generation with language models for question answering," *arXiv preprint arXiv:2410.22353*, 2024.
- [55] Y. Yang, S. Chen, Y. Zhu, X. Liu, S. Pan, and X. Wang, "Intelligent question answering for water conservancy project inspection driven by knowledge graph and large language model collaboration," *LHB*, vol. 110, no. 1, p. 2397337, 2024.
- [56] J. Huang, M. Li, Z. Yao, Z. Yang, Y. Xiao, F. Ouyang, X. Li, S. Han, and H. Yu, "Ritek: A dataset for large language models complex reasoning over textual knowledge graphs," *arXiv preprint arXiv:2410.13987*, 2024.
- [57] M. Kang, J. M. Kwak, J. Baek, and S. J. Hwang, "Knowledge graph-augmented language models for knowledge-grounded dialogue generation," *arXiv preprint arXiv:2305.18846*, 2023.
- [58] Y. Dong, S. Wang, H. Zheng, J. Chen, Z. Zhang, and C. Wang, "Advanced rag models with graph structures: Optimizing complex knowledge reasoning and text generation," in *2024 5th International Symposium on Computer Engineering and Intelligent Communications (ISCEIC)*. IEEE, 2024, pp. 626–630.
- [59] L. Cabello, C. Martin-Turrero, U. Akujuobi, A. Sogaard, and C. Bobed, "Meg: Medical knowledge-augmented large language models for question answering," *arXiv preprint arXiv:2411.03883*, 2024.
- [60] M. He, A. Zhou, and X. Shi, "Enhancing textbook question answering with knowledge graph-augmented large language models," in *The 16th Asian Conference on Machine Learning (Conference Track)*, 2024.
- [61] T. Shen, E. Cambria, J. Wang, Y. Cai, and X. Zhang, "Insight at the right spot: Provide decisive subgraph information to graph llm with reinforcement learning," *Information Fusion*, vol. 117, p. 102860, 2025.
- [62] J. Park, M. Joo, J.-K. Kim, and H. J. Kim, "Generative subgraph retrieval for knowledge graph-grounded dialog generation," *arXiv preprint arXiv:2410.09350*, 2024.
- [63] B. Jin, C. Xie, J. Zhang, K. K. Roy, Y. Zhang, Z. Li, R. Li, X. Tang, S. Wang, Y. Meng *et al.*, "Graph chain-of-thought: Augmenting large language models by reasoning on graphs," *arXiv preprint arXiv:2404.07103*, 2024.
- [64] J. Wang, Q. Sun, N. Chen, X. Li, and M. Gao, "Boosting language models reasoning with chain-of-knowledge prompting," *arXiv preprint arXiv:2306.06427*, 2023.
- [65] Anonymous, "Large language models based graph convolution for text-attributed networks," in *Submitted to The Thirteenth International Conference on Learning Representations, 2024*, under review. [Online]. Available: <https://openreview.net/forum?id=x5FfUvslIE>
- [66] J. Tang, Y. Yang, W. Wei, L. Shi, L. Su, S. Cheng, D. Yin, and C. Huang, "Graphgpt: Graph instruction tuning for large language models," *arXiv preprint arXiv:2310.13023*, 2023.
- [67] L. Wu, Z. Qiu, Z. Zheng, H. Zhu, and E. Chen, "Exploring large language model for graph data understanding in online job recommendations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 8, 2024, pp. 9178–9186.
- [68] W. Zhang, Y. Zhu, M. Chen, Y. Geng, Y. Huang, Y. Xu, W. Song, and H. Chen, "Structure pretraining and prompt tuning for knowledge graph transfer," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 2581–2590.
- [69] Y. Tan, H. Lv, X. Huang, J. Zhang, S. Wang, and C. Yang, "Musegraph: Graph-oriented instruction tuning of large language models for generic graph mining," *arXiv preprint arXiv:2403.04780*, 2024.
- [70] Z. Ji, Z. Liu, N. Lee, T. Yu, B. Wilie, M. Zeng, and P. Fung, "Rho: Reducing hallucination in open-domain dialogues with knowledge grounding," in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, pp. 4504–4522.
- [71] Y. Tian, H. Song, Z. Wang, H. Wang, Z. Hu, F. Wang, N. V. Chawla, and P. Xu, "Graph neural prompting with large language models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, 2024, pp. 19 080–19 088.
- [72] R. Ye, C. Zhang, R. Wang, S. Xu, and Y. Zhang, "Language is all a graph needs," in *Findings of the Association for Computational Linguistics: EACL 2024*, 2024, pp. 1955–1973.
- [73] R. Chen, T. Zhao, A. K. JAISWAL, N. Shah, and Z. Wang, "LLaGA: Large language and graph assistant," in *Forty-first International Conference on Machine Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=B48Pzc4oKi>
- [74] N. Choudhary and C. K. Reddy, "Complex logical reasoning over knowledge graphs using large language models," *arXiv preprint arXiv:2305.01157*, 2023.
- [75] X. Li, R. Zhao, Y. K. Chia, B. Ding, S. Joty, S. Poria, and L. Bing, "Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources," in *The Twelfth International Conference on Learning Representations*, 2024.
- [76] Y. Wen, Z. Wang, and J. Sun, "Mindmap: Knowledge graph prompting sparks graph of thoughts in large language models," *arXiv preprint arXiv:2308.09729*, 2023.
- [77] X. Wang, Q. Yang, Y. Qiu, J. Liang, Q. He, Z. Gu, Y. Xiao, and W. Wang, "Knowledgegpt: Enhancing large language models with retrieval and storage access on knowledge bases," *arXiv preprint arXiv:2308.11761*, 2023.

- [78] L. Bahr, C. Wehner, J. Wewerka, J. Bittencourt, U. Schmid, and R. Daub, "Knowledge graph enhanced retrieval-augmented generation for failure mode and effects analysis," *arXiv preprint arXiv:2406.18114*, 2024.
- [79] G. Liu, Y. Zhang, Y. Li, and Q. Yao, "Explore then determine: A gnn-llm synergy framework for reasoning over knowledge graph," *arXiv preprint arXiv:2406.01145*, 2024.
- [80] T. Zhou, Y. Chen, K. Liu, and J. Zhao, "Cogmg: Collaborative augmentation between large language model and knowledge graph," *arXiv preprint arXiv:2406.17231*, 2024.
- [81] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. B. Van Den Driessche, J.-B. Lespiau, B. Damoc, A. Clark *et al.*, "Improving language models by retrieving from trillions of tokens," in *International conference on machine learning*. PMLR, 2022, pp. 2206–2240.
- [82] G. Izacard, P. Lewis, M. Lomeli, L. Hosseini, F. Petroni, T. Schick, J. Dwivedi-Yu, A. Joulin, S. Riedel, and E. Grave, "Atlas: Few-shot learning with retrieval augmented language models," *Journal of Machine Learning Research*, vol. 24, no. 251, pp. 1–43, 2023.
- [83] Z. Jiang, F. F. Xu, L. Gao, Z. Sun, Q. Liu, J. Dwivedi-Yu, Y. Yang, J. Callan, and G. Neubig, "Active retrieval augmented generation," pp. 7969–7992, 2023.
- [84] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, and H. Wang, "Retrieval-augmented generation for large language models: A survey," *arXiv preprint arXiv:2312.10997*, 2023.
- [85] X. Wang, Z. Wang, X. Gao, F. Zhang, Y. Wu, Z. Xu, T. Shi, Z. Wang, S. Li, Q. Qian *et al.*, "Searching for best practices in retrieval-augmented generation," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024, pp. 17 716–17 736.
- [86] W. Yu, H. Zhang, X. Pan, K. Ma, H. Wang, and D. Yu, "Chain-of-note: Enhancing robustness in retrieval-augmented language models," *arXiv preprint arXiv:2311.09210*, 2023.
- [87] Z. Zhong, H. Liu, X. Cui, X. Zhang, and Z. Qin, "Mix-of-granularity: Optimize the chunking granularity for retrieval-augmented generation," *arXiv preprint arXiv:2406.00456*, 2024.
- [88] LlamaIndex, "Llamaindex - build knowledge assistants over your enterprise data," 2024, accessed: 2025-01-08. [Online]. Available: <https://www.llamaindex.ai/>
- [89] Z. Wang, H. Yuan, W. Dong, G. Cong, and F. Li, "Corag: A cost-constrained retrieval optimization system for retrieval-augmented generation," *arXiv preprint arXiv:2411.00744*, 2024.
- [90] X. Chen, L. Li, N. Zhang, X. Liang, S. Deng, C. Tan, F. Huang, L. Si, and H. Chen, "Decoupling knowledge from memorization: Retrieval-augmented prompt learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 23 908–23 922, 2022.
- [91] Y. Mao, P. He, X. Liu, Y. Shen, J. Gao, J. Han, and W. Chen, "Generation-augmented retrieval for open-domain question answering," *arXiv preprint arXiv:2009.08553*, 2020.
- [92] Y.-S. Chuang, W. Fang, S.-W. Li, W.-t. Yih, and J. Glass, "Expand, rerank, and retrieve: Query reranking for open-domain question answering," *arXiv preprint arXiv:2305.17080*, 2023.
- [93] E. Doostmohammadi, T. Norlund, M. Kuhlmann, and R. Johansson, "Surface-based retrieval reduces perplexity of retrieval-augmented language models," pp. 521–529, 2023.
- [94] W. Shi, S. Min, M. Yasunaga, M. Seo, R. James, M. Lewis, L. Zettlemoyer, and W.-t. Yih, "Replug: Retrieval-augmented black-box language models," *arXiv preprint arXiv:2301.12652*, 2023.
- [95] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [96] A. Lazaridou, E. Gribovskaya, W. Stokowiec, and N. Grigorev, "Internet-augmented language models through few-shot prompting for open-domain question answering," *arXiv preprint arXiv:2203.05115*, 2022.
- [97] M. A. Arefeen, B. Debnath, and S. Chakradhar, "Leancontext: Cost-efficient domain-specific question answering using llms," *Natural Language Processing Journal*, vol. 7, p. 100065, 2024.
- [98] Y. Wang, P. Li, M. Sun, and Y. Liu, "Self-knowledge guided retrieval augmentation for large language models," *arXiv preprint arXiv:2310.05002*, 2023.
- [99] Z. Wang, J. Araki, Z. Jiang, M. R. Parvez, and G. Neubig, "Learning to filter context for retrieval-augmented generation," *arXiv preprint arXiv:2311.08377*, 2023.
- [100] X. Lyu, S. Grafberger, S. Biegel, S. Wei, M. Cao, S. Schelker, and C. Zhang, "Improving retrieval-augmented large language models via data importance learning," *arXiv preprint arXiv:2307.03027*, 2023.
- [101] X. Cheng, D. Luo, X. Chen, L. Liu, D. Zhao, and R. Yan, "Lift yourself up: Retrieval-augmented text generation with self-memory," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [102] H. Luo, Y.-S. Chuang, Y. Gong, T. Zhang, Y. Kim, X. Wu, D. Fox, H. Meng, and J. Glass, "Sail: Search-augmented instruction learning," *arXiv preprint arXiv:2305.15225*, 2023.
- [103] Z. Li, R. Guo, and S. Kumar, "Decoupled context processing for context augmented language modeling," *Advances in Neural Information Processing Systems*, vol. 35, pp. 21 698–21 710, 2022.
- [104] M. De Jong, Y. Zemlyanskiy, N. FitzGerald, J. Ainslie, S. Sanghai, F. Sha, and W. W. Cohen, "Pre-computed memory or on-the-fly encoding? a hybrid approach to retrieval augmentation makes the most of your compute," in *International Conference on Machine Learning*. PMLR, 2023, pp. 7329–7342.
- [105] W. Jiang, S. Zhang, B. Han, J. Wang, B. Wang, and T. Kraska, "Piperag: Fast retrieval-augmented generation via algorithm-system co-design," *arXiv preprint arXiv:2403.05676*, 2024.
- [106] C. Jin, Z. Zhang, X. Jiang, F. Liu, X. Liu, X. Liu, and X. Jin, "Ragcache: Efficient knowledge caching for retrieval-augmented generation," *arXiv preprint arXiv:2404.12457*, 2024.
- [107] J. Liu, L. Li, T. Xiang, B. Wang, and Y. Qian, "Tera-llm: Token compression retrieval augmented large language model for inference cost reduction," *arXiv preprint arXiv:2310.15556*, 2023.
- [108] M. Jin, Q. Yu, D. Shu, H. Zhao, W. Hua, Y. Meng, Y. Zhang, and M. Du, "The impact of reasoning step length on large language models," in *Findings of the Association for Computational Linguistics: ACL 2024*, 2024.
- [109] A. Yates, M. Banko, M. Broadhead, M. J. Cafarella, O. Etzioni, and S. Soderland, "Textrunner: open information extraction on the web," in *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 2007, pp. 25–26.
- [110] G. Angeli, M. J. J. Premkumar, and C. D. Manning, "Leveraging linguistic structure for open domain information extraction," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015, pp. 344–354.
- [111] W. Li, H. Zhou, J. Dong, Q. Zhang, Q. Li, G. Baciú, J. Cao, and X. Huang, "Constructing low-redundant and high-accuracy knowledge graphs for education," in *International Conference on Web-Based Learning*. Springer, 2022, pp. 148–160.
- [112] K. Kolluru, V. Adlakha, S. Aggarwal, S. Chakrabarti *et al.*, "Openie6: Iterative grid labeling and coordination analysis for open information extraction," *arXiv preprint arXiv:2010.03147*, 2020.
- [113] C. Wang, X. Liu, and D. Song, "Language models are open knowledge graphs," *arXiv preprint arXiv:2010.11967*, 2020.
- [114] C. Wang, X. Liu, Z. Chen, H. Hong, J. Tang, and D. Song, "Zero-shot information extraction as a unified text-to-triple translation," *arXiv preprint arXiv:2109.11171*, 2021.
- [115] R. Cohen, M. Geva, J. Berant, and A. Globerson, "Crawling the internal knowledge-base of language models," *arXiv preprint arXiv:2301.12810*, 2023.
- [116] H. Chen, X. Shen, Q. Lv, J. Wang, X. Ni, and J. Ye, "Sac-kg: Exploiting large language models as skilled automatic constructors for domain knowledge graphs," *arXiv preprint arXiv:2410.02811*, 2024.
- [117] A. Fan, C. Gardent, C. Braud, and A. Bordes, "Using local knowledge graph construction to scale Seq2Seq models to multi-document inputs," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Nov. 2019.
- [118] B. Chen and A. L. Bertozzi, "Autokg: Efficient automated knowledge graph generation for language models," in *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 2023, pp. 3117–3126.
- [119] Z. Xu, M. J. Cruz, M. Guevara, T. Wang, M. Deshpande, X. Wang, and Z. Li, "Retrieval-augmented generation with knowledge graphs for customer service question answering," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024, pp. 2905–2909.
- [120] E. Montiel-Ponsoda and V. Rodriguez-Doncel, "Lynx: building the legal knowledge graph for smart compliance services in multilingual europe," in *Proceedings of the 1st workshop on LREC (Language Resources and Technologies for the Legal Knowledge Graph) workshop*, 2018, pp. 19–22.
- [121] R. Wang, Y. Yan, J. Wang, Y. Jia, Y. Zhang, W. Zhang, and X. Wang, "Acekg: A large-scale knowledge graph for academic data mining," in

- Proceedings of the 27th ACM international conference on information and knowledge management*, 2018, pp. 1487–1490.
- [122] J. H. Morris, K. Soman, R. E. Akbas, X. Zhou, B. Smith, E. C. Meng, C. C. Huang, G. Ceroni, G. Schenk, A. Rizk-Jackson *et al.*, “The scalable precision medicine open knowledge engine (spoke): a massive knowledge graph of biomedical information,” *Bioinformatics*, vol. 39, no. 2, p. btad080, 2023.
- [123] D. Szklarczyk, A. L. Gable, D. Lyon, A. Junge, S. Wyder, J. Huerta-Cepas, M. Simonovic, N. T. Doncheva, J. H. Morris, P. Bork *et al.*, “String v11: protein–protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets,” *Nucleic acids research*, vol. 47, no. D1, pp. D607–D613, 2019.
- [124] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, “DBpedia: A nucleus for a web of open data,” in *The semantic web*. Springer, 2007, pp. 722–735.
- [125] F. M. Suchanek, M. Alam, T. Bonald, L. Chen, P.-H. Paris, and J. Soria, “Yago 4.5: A large and clean knowledge base with a rich taxonomy,” in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR ’24, 2024, p. 131–140.
- [126] K. Soman, P. W. Rose, J. H. Morris, R. E. Akbas, B. Smith, B. Peetoom, C. Villouta-Reyes, G. Ceroni, Y. Shi, A. Rizk-Jackson, S. Israni, C. A. Nelson, S. Huang, and S. E. Baranzini, “Biomedical knowledge graph-optimized prompt generation for large language models,” *Bioinformatics*, vol. 40, no. 9, p. btac560, 2024.
- [127] C. Shengyuan, Y. Cai, H. Fang, X. Huang, and M. Sun, “Differentiable neuro-symbolic reasoning on large-scale knowledge graphs,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [128] S. Chen, Q. Zhang, J. Dong, W. Hua, Q. Li, and X. Huang, “Entity alignment with noisy annotations from large language models,” *arXiv preprint arXiv:2405.16806*, 2024.
- [129] S. Chen, Q. Zhang, J. Dong, W. Hua, J. Cao, and X. Huang, “Neuro-symbolic entity alignment via variational inference,” *arXiv preprint arXiv:2410.04153*, 2024.
- [130] Q. Zhang, J. Dong, K. Duan, X. Huang, Y. Liu, and L. Xu, “Contrastive knowledge graph error detection,” in *ACM International Conference on Information and Knowledge Management*, 2022.
- [131] J. Dong, Q. Zhang, X. Huang, Q. Tan, D. Zha, and Z. Zihao, “Active ensemble learning for knowledge graph error detection,” in *Proceedings of the sixteenth ACM international conference on web search and data mining*, 2023, pp. 877–885.
- [132] Q. Zhang, J. Dong, Q. Tan, and X. Huang, “Integrating entity attributes for error-aware knowledge graph embedding,” *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [133] Y. Zhu, H. Yuan, S. Wang, J. Liu, W. Liu, C. Deng, H. Chen, Z. Liu, Z. Dou, and J.-R. Wen, “Large language models for information retrieval: A survey,” *arXiv preprint arXiv:2308.07107*, 2023.
- [134] S. Selva Birunda and R. Kanniga Devi, “A review on word embedding techniques for text classification,” *Innovative Data Communication Technologies and Application: Proceedings of ICIDCA 2020*, pp. 267–281, 2021.
- [135] T. Kenter and M. De Rijke, “Short text similarity with word embeddings,” in *Proceedings of the 24th ACM international conference on information and knowledge management*, 2015, pp. 1411–1420.
- [136] Z. Lee, Z. Huang, Z. Yao, J. Liu, A. Xin, L. Hou, and J. Li, “Diakop: Dialogue-based knowledge-oriented programming for neural-symbolic knowledge base question answering,” in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 5234–5238.
- [137] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [138] Y. Bei, S. Zhou, Q. Tan, H. Xu, H. Chen, Z. Li, and J. Bu, “Reinforcement neighborhood selection for unsupervised graph anomaly detection,” in *2023 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2023, pp. 11–20.
- [139] Z. Li, X. Sun, Y. Luo, Y. Zhu, D. Chen, Y. Luo, X. Zhou, Q. Liu, S. Wu, L. Wang *et al.*, “Gslb: the graph structure learning benchmark,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [140] Z. Liu, C. Shengyuan, K. Zhou, D. Zha, X. Huang, and X. Hu, “Rsc: accelerate graph neural networks training via randomized sparse computations,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 21 951–21 968.
- [141] M. Zhang and Y. Chen, “Link prediction based on graph neural networks,” *Advances in neural information processing systems*, vol. 31, 2018.
- [142] H. Chen, Y. Bei, W. Huang, S. Chen, F. Huang, and X. Huang, “Graph cross-correlated network for recommendation,” *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [143] Y. Bei, H. Chen, S. Chen, X. Huang, S. Zhou, and F. Huang, “Non-recursive cluster-scale graph interacted model for click-through rate prediction,” in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023, pp. 3748–3752.
- [144] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *International Conference on Learning Representations*, 2017.
- [145] C.-Y. Hsu, K. Cox, J. Xu, Z. Tan, T. Zhai, M. Hu, D. Pratt, T. Chen, Z. Hu, and Y. Ding, “Thought graph: Generating thought process for biological reasoning,” in *Companion Proceedings of the ACM on Web Conference 2024*, 2024, pp. 537–540.
- [146] P. Panda, A. Agarwal, C. Devaguptapu, M. Kaul *et al.*, “Holmes: Hyper-relational knowledge graphs for multi-hop question answering using llms,” *arXiv preprint arXiv:2406.06027*, 2024.
- [147] Y. Li, Z. Li, P. Wang, J. Li, X. Sun, H. Cheng, and J. X. Yu, “A survey of graph meets large language model: Progress and future directions,” *arXiv preprint arXiv:2311.12399*, 2023.
- [148] Y. Li, P. Wang, Z. Li, J. X. Yu, and J. Li, “Zero: Investigating cross-dataset zero-shot transferability in graphs,” in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 1725–1735.
- [149] J. Li, X. Sun, Y. Li, Z. Li, H. Cheng, and J. X. Yu, “Graph intelligence with large language models and prompt learning,” in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 6545–6554.
- [150] Z. Wang, S. X. Teo, J. Ouyang, Y. Xu, and W. Shi, “M-rag: Reinforcing large language model performance through retrieval-augmented generation with multiple partitions,” *arXiv preprint arXiv:2405.16420*, 2024.
- [151] M. Kulkarni, P. Tangarajan, K. Kim, and A. Trivedi, “Reinforcement learning for optimizing rag for domain chatbots,” *arXiv preprint arXiv:2401.06800*, 2024.
- [152] N. Gupta, S. Kirtania, P. Gupta, K. Kariya, S. Gulwani, A. Iyer, S. Parthasarathy, A. Radhakrishna, S. K. Rajamani, and G. Soares, “Stackfeed: Structured textual actor-critic knowledge base editing with feedback,” *arXiv preprint arXiv:2410.10584*, 2024.
- [153] B. Sarmah, B. Hall, R. Rao, S. Patel, S. Pasquali, and D. Mehta, “Hybridrag: Integrating knowledge graphs and vector retrieval augmented generation for efficient information extraction,” *arXiv preprint arXiv:2408.04948*, 2024.
- [154] W. Xie, X. Liang, Y. Liu, K. Ni, H. Cheng, and Z. Hu, “Weknow-rag: An adaptive approach for retrieval-augmented generation integrating web search and knowledge graphs,” *arXiv preprint arXiv:2408.07611*, 2024.
- [155] R. Mendes, D. Oliveira, and V. Garcia, “Application of generative ai as an enterprise wikibase knowledge graph q&a system,” in *Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models (KaLLM 2024)*, 2024, pp. 35–42.
- [156] M. Alhanahnah, Y. Boshmaf, and B. Baudry, “Depesrag: Towards managing software dependencies using large language models,” *arXiv preprint arXiv:2405.20455*, 2024.
- [157] M. Dehghan, M. A. Alomrani, S. Bagga, D. Alfonso-Hermelo, K. Bibi, A. Ghaddar, Y. Zhang, X. Li, J. Hao, Q. Liu *et al.*, “Ewek-qa: Enhanced web and efficient knowledge graph retrieval for citation-based question answering systems,” *arXiv preprint arXiv:2406.10393*, 2024.
- [158] Y. Zhao, Z. Jin, F. Zhang, X. Xu, and D. Liu, “Bridging the gap: Advancing commonsense question answering with integrated multimodal knowledge,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 46, 2024.
- [159] F. F. Bayat, K. Qian, B. Han, Y. Sang, A. Belyy, S. Khorshidi, F. Wu, I. Ilyas, and Y. Li, “Fleek: Factual error detection and correction with evidence retrieved from external knowledge,” in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2023, pp. 124–130.
- [160] C. Deng, T. Zhang, Z. He, Q. Chen, Y. Shi, Y. Xu, L. Fu, W. Zhang, X. Wang, C. Zhou *et al.*, “K2: A foundation language model for geoscience knowledge understanding and utilization,” in *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 2024, pp. 161–170.
- [161] Y. Li, Y. Yu, C. Liang, N. Karampatziakis, P. He, W. Chen, and T. Zhao, “Loftq: LoRA-fine-tuning-aware quantization for large language models,” in *The Twelfth International Conference on Learning Representations*, 2024.

- [162] Z. Hu, L. Wang, Y. Lan, W. Xu, E.-P. Lim, L. Bing, X. Xu, S. Poria, and R. Lee, "Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023, pp. 5254–5276.
- [163] X. Lin, W. Wang, Y. Li, S. Yang, F. Feng, Y. Wei, and T.-S. Chua, "Data-efficient fine-tuning for llm-based recommendation," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024, pp. 365–374.
- [164] J. He, M. Rungta, D. Koleczek, A. Sekhon, F. X. Wang, and S. Hasan, "Does prompt formatting have any impact on llm performance?" *arXiv preprint arXiv:2411.10541*, 2024.
- [165] I. Arawjo, C. Swoopes, P. Vaithilingam, M. Wattenberg, and E. L. Glassman, "Chainforge: A visual toolkit for prompt engineering and llm hypothesis testing," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1–18.
- [166] Y. Liang, K. Tan, T. Xie, W. Tao, S. Wang, Y. Lan, and W. Qian, "Aligning large language models to a domain-specific graph database," *arXiv preprint arXiv:2402.16567*, 2024.
- [167] C. Mavromatis and G. Karypis, "Gnn-rag: Graph neural retrieval for large language model reasoning," *arXiv preprint arXiv:2405.20139*, 2024.
- [168] X. Huang, K. Han, Y. Yang, D. Bao, Q. Tao, Z. Chai, and Q. Zhu, "Can gnn be good adapter for llms?" in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 893–904.
- [169] Y. Li, Z. Li, K. Zhang, R. Dan, S. Jiang, and Y. Zhang, "Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge," *Cureus*, vol. 15, no. 6, 2023.
- [170] C. Wu, W. Lin, X. Zhang, Y. Zhang, W. Xie, and Y. Wang, "Pmc-llama: toward building open-source language models for medicine," *Journal of the American Medical Informatics Association*, p. ocae045, 2024.
- [171] M. Fatehikia, J. K. Lucas, and S. Chawla, "T-rag: lessons from the llm trenches," *arXiv preprint arXiv:2402.07483*, 2024.
- [172] Y. Zhang, K. Bao, M. Yan, W. Wang, F. Feng, and X. He, "Text-like encoding of collaborative information in large language models for recommendation," *arXiv preprint arXiv:2406.03210*, 2024.
- [173] F. Bianchini, M. Calamo, F. De Luzi, M. Macri, and M. Mecella, "Enhancing complex linguistic tasks resolution through fine-tuning llms, rag and knowledge graphs (short paper)," in *International Conference on Advanced Information Systems Engineering*. Springer, 2024, pp. 147–155.
- [174] D. Shu, T. Chen, M. Jin, Y. Zhang, M. Du, and Y. Zhang, "Knowledge graph large language model (kg-llm) for link prediction," *arXiv preprint arXiv:2403.07311*, 2024.
- [175] N. Chen, Y. Li, J. Tang, and J. Li, "Graphwiz: An instruction-following language model for graph problems," *arXiv preprint arXiv:2402.16029*, 2024.
- [176] H. Luo, Z. Tang, S. Peng, Y. Guo, W. Zhang, C. Ma, G. Dong, M. Song, W. Lin *et al.*, "Chatkbqa: A generate-then-retrieve framework for knowledge base question answering with fine-tuned large language models," *arXiv preprint arXiv:2310.08975*, 2023.
- [177] Y. Yan, P. Zheng, and Y. Wang, "Enhancing large language model capabilities for rumor detection with knowledge-powered prompting," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108259, 2024.
- [178] S. Yuan, E. Nie, M. Färber, H. Schmid, and H. Schütze, "Gnavi: Navigating the information flow in large language models by graph neural network," *arXiv preprint arXiv:2402.11709*, 2024.
- [179] R. Chen, T. Zhao, A. K. JAISWAL, N. Shah, and Z. Wang, "Llaga: Large language and graph assistant," in *Forty-first International Conference on Machine Learning*, 2024.
- [180] Z. Zhang, A. Zhang, M. Li, and A. Smola, "Automatic chain of thought prompting in large language models," in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=5NTt8GFjUHkr>
- [181] L. Wang, W. Xu, Y. Lan, Z. Hu, Y. Lan, R. K.-W. Lee, and E.-P. Lim, "Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 2609–2634.
- [182] Z. Bi, N. Zhang, Y. Jiang, S. Deng, G. Zheng, and H. Chen, "When do program-of-thought works for reasoning?" in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 17 691–17 699.
- [183] L. Ma and J. Sanders, "Markov chains and hitting times for error accumulation in quantum circuits," in *Performance Evaluation Methodologies and Tools: 14th EAI International Conference, VALUETOOLS 2021, Virtual Event, October 30–31, 2021, Proceedings*. Springer, 2021, pp. 36–55.
- [184] Y. Bai, F. Chen, H. Wang, C. Xiong, and S. Mei, "Transformers as statisticians: Provable in-context learning with in-context algorithm selection," *Advances in neural information processing systems*, vol. 36, 2024.
- [185] R. Li, Y. Wang, Z. Wen, M. Cui, and Q. Miao, "Different paths to the same destination: Diversifying llms generation for multi-hop open-domain question answering," *Knowledge-Based Systems*, p. 112789, 2024.
- [186] S. M. Xie, A. Raghunathan, P. Liang, and T. Ma, "An explanation of in-context learning as implicit bayesian inference," in *International Conference on Learning Representations*, 2022.
- [187] Y. Zhou, X. Li, Q. Wang, and J. Shen, "Visual in-context learning for large vision-language models," in *Findings of the Association for Computational Linguistics: ACL 2024*, Aug. 2024, pp. 15 890–15 902.
- [188] Y. Zhang, K. Zhou, and Z. Liu, "What makes good examples for visual in-context learning?" *Advances in Neural Information Processing Systems*, vol. 36, pp. 17 773–17 794, 2023.
- [189] J. Dong, Q. Zhang, H. Zhou, D. Zha, P. Zheng, and X. Huang, "Modality-aware integration with large language models for knowledge-based visual question answering," *arXiv preprint arXiv:2402.12728*, 2024.
- [190] Z. Hong, Z. Yuan, H. Chen, Q. Zhang, F. Huang, and X. Huang, "Knowledge-to-SQL: Enhancing SQL generation with data expert LLM," in *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics, Aug. 2024.
- [191] S. Xu, L. Pang, H. Shen, X. Cheng, and T.-S. Chua, "Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks," in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 1362–1373.
- [192] K. Shridhar, K. Sinha, A. Cohen, T. Wang, P. Yu, R. Pasunuru, M. Sachan, J. Weston, and A. Celikyilmaz, "The art of llm refinement: Ask, refine, and trust," in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024, pp. 5872–5883.
- [193] Z. Hong and J. Liu, "Towards better question generation in QA-based event extraction," in *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics, Aug. 2024.
- [194] S. Geng, Z. Fu, Y. Ge, L. Li, G. De Melo, and Y. Zhang, "Improving personalized explanation generation through visualization," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 244–255.
- [195] H. Zhou, S. Zhou, H. Chen, N. Liu, F. Yang, and X. Huang, "Enhancing explainable rating prediction through annotated macro concepts," in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024, pp. 11 736–11 748.
- [196] A. Tariq, M. Luo, A. Urooj, A. Das, J. Jeong, S. Trivedi, B. Patel, and I. Banerjee, "Domain-specific llm development and evaluation—a case-study for prostate cancer," *medRxiv*, pp. 2024–03, 2024.
- [197] J. Shen, N. Tenenholz, J. B. Hall, D. Alvarez-Melis, and N. Fusi, "Tag-llm: Repurposing general-purpose llms for specialized domains," *arXiv preprint arXiv:2402.05140*, 2024.
- [198] C. Jeong, "Fine-tuning and utilization methods of domain-specific llms," *arXiv preprint arXiv:2401.02981*, 2024.
- [199] K. Singhal, S. Azizi, T. Tu, S. S. Mahdavi, J. Wei, H. W. Chung, N. Scales, A. Tanwani, H. Cole-Lewis, S. Pfohl *et al.*, "Large language models encode clinical knowledge," *Nature*, vol. 620, no. 7972, pp. 172–180, 2023.
- [200] D. Yang, J. Rao, K. Chen, X. Guo, Y. Zhang, J. Yang, and Y. Zhang, "Im-rag: Multi-round retrieval-augmented generation through learning inner monologues," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024, pp. 730–740.
- [201] S. Yang, E. Gribovskaya, N. Kassner, M. Geva, and S. Riedel, "Do large language models latently perform multi-hop reasoning?" *arXiv preprint arXiv:2402.16837*, 2024.
- [202] E. Tanwar, S. Dutta, M. Borthakur, and T. Chakraborty, "Multilingual llms are better cross-lingual in-context learners with alignment," in *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- [203] G. Liu, Y. Sun, R. Sun, X. Dong, and H. Xiong, "Agentps: Agentic process supervision for multi-modal content quality assurance through multi-round qa," *arXiv preprint arXiv:2412.15251*, 2024.



- [204] J. Wang, H. Lu, Y. Liu, H. Ma, Y. Wang, Y. Gu, S. Zhang, N. Han, S. Bi, L. Baugher *et al.*, “Llms for user interest exploration in large-scale recommendation systems,” in *Proceedings of the 18th ACM Conference on Recommender Systems*, 2024, pp. 872–877.
- [205] S. Kim, H. Kang, S. Choi, D. Kim, M. Yang, and C. Park, “Large language models meet collaborative filtering: An efficient all-round llm-based recommender system,” in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 1395–1406.
- [206] Y. Xia, J. Kim, Y. Chen, H. Ye, S. Kundu, C. C. Hao, and N. Talati, “Understanding the performance and estimating the cost of llm fine-tuning,” in *2024 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2024, pp. 210–223.
- [207] W. Sun, J. Wang, Q. Guo, Z. Li, W. Wang, and R. Hai, “Cebench: A benchmarking toolkit for the cost-effectiveness of llm pipelines,” *arXiv preprint arXiv:2407.12797*, 2024.
- [208] Z. Liu, X. He, Y. Tian, and N. V. Chawla, “Can we soft prompt llms for graph learning tasks?” in *Companion Proceedings of the ACM on Web Conference 2024*, 2024, pp. 481–484.
- [209] J. Huang, X. Zhang, Q. Mei, and J. Ma, “Can llms effectively leverage graph structural information through prompts, and why?” *Transactions on Machine Learning Research*, 2024.
- [210] H. Chen, Y. Bei, Q. Shen, Y. Xu, S. Zhou, W. Huang, F. Huang, S. Wang, and X. Huang, “Macro graph neural networks for online billion-scale recommender systems,” in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 3598–3608.
- [211] X. Wan, K. Xu, X. Liao, Y. Jin, K. Chen, and X. Jin, “Scalable and efficient full-graph gnn training for large graphs,” *Proceedings of the ACM on Management of Data*, vol. 1, no. 2, pp. 1–23, 2023.
- [212] B. Bi, S. Liu, Y. Wang, L. Mei, and X. Cheng, “Lpnl: Scalable link prediction with large language models,” in *Findings of the Association for Computational Linguistics ACL 2024*, 2024, pp. 3615–3625.
- [213] M. Yang, Y. Wang, and Y. Gu, “Language-based reasoning graph neural network for commonsense question answering,” *Neural Networks*, p. 106816, 2024.
- [214] A. O. Saleh, G. Tür, and Y. Saygin, “Sg-rag: Multi-hop question answering with large language models through knowledge graphs,” in *Proceedings of the 7th International Conference on Natural Language and Speech Processing (ICNLSP 2024)*, 2024, pp. 439–448.
- [215] Z. Zhang, L. Wen, and W. Zhao, “A gail fine-tuned llm enhanced framework for low-resource knowledge graph question answering,” in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 3300–3309.
- [216] K. Zuo, Y. Jiang, F. Mo, and P. Lio, “Kg4diagnosis: A hierarchical multi-agent llm framework with knowledge graph enhancement for medical diagnosis,” *arXiv preprint arXiv:2412.16833*, 2024.
- [217] X. Zhao, M. Blum, R. Yang, B. Yang, L. M. Carpintero, M. Pina-Navarro, T. Wang, X. Li, H. Li, Y. Fu *et al.*, “Agentigraph: An interactive knowledge graph platform for llm-based chatbots utilizing private data,” *arXiv preprint arXiv:2410.11531*, 2024.
- [218] R. Yang, B. Yang, A. Feng, S. Ouyang, M. Blum, T. She, Y. Jiang, F. Lecue, J. Lu, and I. Li, “Graphusion: A rag framework for knowledge graph construction with a global perspective,” *arXiv preprint arXiv:2410.17600*, 2024.
- [219] Z. Sepasdar, S. Gautam, C. Midoglu, M. A. Riegler, and P. Halvorsen, “Soccer-graphrag: Applications of graphrag in soccer,” in *International Workshop on Graph-Based Approaches in Information Retrieval*. Springer, 2024, pp. 1–10.
- [220] A. Bordes, N. Usunier, S. Chopra, and J. Weston, “Large-scale simple question answering with memory networks,” *arXiv preprint arXiv:1506.02075*, 2015.
- [221] J. Berant, A. Chou, R. Frostig, and P. Liang, “Semantic parsing on freebase from question-answer pairs,” in *Proceedings of the 2013 conference on empirical methods in natural language processing*, 2013, pp. 1533–1544.
- [222] A. Talmor and J. Berant, “The web as a knowledge-base for answering complex questions,” *arXiv preprint arXiv:1803.06643*, 2018.
- [223] Y. Zhang, H. Dai, Z. Kozareva, A. Smola, and L. Song, “Variational reasoning for question answering with knowledge graph,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [224] P. Trivedi, G. Maheshwari, M. Dubey, and J. Lehmann, “Lc-quad: A corpus for complex question answering over knowledge graphs,” in *International Semantic Web Conference*. Springer, 2017, pp. 210–218.
- [225] M. Dubey, D. Banerjee, A. Abdelkawi, and J. Lehmann, “Lc-quad 2.0: A large dataset for complex question answering over wikidata and dbpedia,” in *Proceedings of the 18th International Semantic Web Conference (ISWC)*. Springer, 2019.
- [226] S. Cao, J. Shi, L. Pan, L. Nie, Y. Xiang, L. Hou, J. Li, B. He, and H. Zhang, “KQA pro: A dataset with explicit compositional programs for complex question answering over knowledge base,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2022.
- [227] P. Sen, A. F. Aji, and A. Saffari, “Mintaka: A complex, natural, and multilingual dataset for end-to-end question answering,” *arXiv preprint arXiv:2210.01613*, 2022.
- [228] J. Kim, S. Park, Y. Kwon, Y. Jo, J. Thorne, and E. Choi, “Factkg: Fact verification via reasoning on knowledge graphs,” *arXiv preprint arXiv:2305.06590*, 2023.
- [229] W.-t. Yih, M. Richardson, C. Meek, M.-W. Chang, and J. Suh, “The value of semantic parse labeling for knowledge base question answering,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2016, pp. 201–206.
- [230] Y. Gu, S. Kase, M. Vanni, B. Sadler, P. Liang, X. Yan, and Y. Su, “Beyond iid: three levels of generalization for question answering on knowledge bases,” in *Proceedings of the Web Conference 2021*, 2021, pp. 3477–3488.
- [231] Y. Lyu, Z. Li, S. Niu, F. Xiong, B. Tang, W. Wang, H. Wu, H. Liu, T. Xu, and E. Chen, “Crud-rag: A comprehensive chinese benchmark for retrieval-augmented generation of large language models,” *ACM Transactions on Information Systems*, 2024.
- [232] H. Qian, P. Zhang, Z. Liu, K. Mao, and Z. Dou, “Memorag: Moving towards next-gen rag via memory-inspired knowledge discovery,” *arXiv preprint arXiv:2409.05591*, 2024.