

Large Language Models Meet Knowledge Graphs for Question Answering: Synthesis and Opportunities

Chuangtao Ma¹, Yongrui Chen², Tianxing Wu^{2*}, Arijit Khan^{3,1*}, Haofen Wang⁴

¹Aalborg University, Denmark, ²Southeast University, China

³Bowling Green State University, USA, ⁴Tongji University, China

chuma@cs.aau.dk, {yongruichen, tianxingwu}@seu.edu.cn, arijitk@bgsu.edu, haofenwang@tongji.edu.cn

Abstract

Large language models (LLMs) have demonstrated remarkable performance on question-answering (QA) tasks because of their superior capabilities in natural language understanding and generation. However, LLM-based QA struggles with complex QA tasks due to poor reasoning capacity, outdated knowledge, and hallucinations. Several recent works synthesize LLMs and knowledge graphs (KGs) for QA to address the above challenges. In this survey, we propose a new structured taxonomy that categorizes the methodology of synthesizing LLMs and KGs for QA according to the categories of QA and the KG's role when integrating with LLMs. We systematically survey state-of-the-art methods in synthesizing LLMs and KGs for QA and compare and analyze these approaches in terms of strength, limitations, and KG requirements. We then align the approaches with QA and discuss how these approaches address the main challenges of different complex QA. Finally, we summarize the advancements, evaluation metrics, and benchmark datasets and highlight open challenges and opportunities.

1 Introduction

Question answering (QA) plays a fundamental role in artificial intelligence, natural language processing, information retrieval, and data management areas since it has a wide range of applications, such as text generation, chatbots, dialog generation, web search, entity linking, natural language query, fact-checking, etc. The pre-trained language models (PLMs) and recent LLMs have shown superior performance in several QA tasks such as KBQA (Knowledge bases QA), KGQA (Knowledge graph QA), CDQA (Closed domain QA), etc. However, PLM and LLM-based methods are incapable of handling complex QA scenarios due to the following limitations. **(1) Limited complex reasoning capability:** LLMs encapsulate very limited reasoning

capability since they have been pre-trained with the tasks of predicting the next word in a text sequence. **(2) Lack of up-to-date and domain-specific knowledge:** LLMs are incapable of generating accurate and up-to-date responses for domain-specific QA because they have been pre-trained on the data and knowledge with a cutoff date. **(3) Tendency to generate hallucinated content:** LLMs usually generate hallucinated content due to the lack of factual verification and logical consistency checking.

Challenges. Retrieval augmented generation (RAG) (Mao et al., 2021) was proposed for open-domain QA by retrieving the relevant contexts from large documents, and several techniques, such as graph neural networks (GNNs) (Li et al., 2025b), have been investigated to enhance the retrieval coverage from passages. Although RAG-based QA can generate better responses in comparison to NoRAG-based QA, it still has limited capability for knowledge reasoning and understanding user interactions during complex QA. Complex QA usually involves knowledge interactions and fusion among data across modalities and sources, and an excellent understanding of complex queries and user interactions, whereas the RAG-based QA suffers a lot from the following technical challenges. **(1) Knowledge conflicts:** Conflicts occur due to the fusion of inconsistent and overlapping knowledge between LLMs and external sources in RAG-based QA that may further tend to generate inconsistent answers. **(2) Poor relevance and quality of retrieved context:** The accuracy of the generated answers in RAG-based QA largely depends on the relevance and quality of the retrieved context, where irrelevant context leads to incorrect results. **(3) Lack of iterative and multi-hop reasoning:** RAG-based QA struggles to generate accurate and explainable answers for questions requiring global and summarized contexts due to a lack of iterative.

The emergence of synthesizing LLMs+KGs provides a unique opportunity to address the above

*Corresponding authors.

challenges and limitations of LLMs for knowledge-intensive tasks, e.g., complex QA (Ma et al., 2025a). The graph retrieval augmented generation (GraphRAG) (Zhang et al., 2025; Peng et al., 2024; Han et al., 2024) and knowledge graph retrieval augmented generation (KG-RAG) (Sanmartin, 2024; Yang et al., 2024) have demonstrated the strengths of unifying LLMs with KGs for complex QA. These GraphRAG and KG-RAG based QA approaches introduce several modules, such as knowledge integration and fusion, reasoning guidelines, and knowledge validation and refinement, to mitigate the above challenges.

Motivation. In recent years, a rapidly increasing number of works on synthesizing LLMs and KGs for QA have been conducted to achieve complex QA under open-domain and long-context settings. This survey aims to address the aforementioned limitations by outlining the recent progress in integrating LLMs with KGs for complex QA, summarizing technical advances, and identifying open challenges and future research opportunities. Our survey differs from the existing surveys (**details in Appendix §A.1**), it provides a comprehensive overview of the recent advancements of QA from the perspective of the roles of KGs when synthesizing LLMs and KGs for complex QA.

Taxonomy. We categorize the methodology of synthesizing LLMs and KGs and complex QA from different perspectives, and a structured taxonomy (**details in Appendix §B**) is given in Figure 1. The taxonomy from different perspectives aims to highlight the alignments between various LLM+KG approaches and different complex QA by discussing how the LLM+KG approaches with different roles of KG can address the challenges of complex QA. Notably, these categories, divided from different perspectives, are non-exclusive, where a work might be classified into multiple categories from different perspectives.

Contributions. Considering the popularity and mainstream adoption of both LLMs and KGs, and their wide applications in QA, our survey is timely. The contributions of this work are summarized below. (1) We introduce the structured taxonomy that categorizes state-of-the-art (SOTA) works on synthesizing LLMs+KGs for QA. (2) We review the SOTA works on synthesizing LLMs and KGs for QA in various categories and discuss the recent advanced topics in this field. (3) We create the alignments between various approaches in syn-

thesizing LLMs and KGs and complex QA, and highlight how these approaches address the specific challenges of different complex QA; (4) We discuss the current challenges and opportunities in synthesizing LLMs and KGs for QA. The online resources of this survey are available on Github¹.

2 Complex QA

Complex QA usually involves question decomposition, knowledge fusion among data across modalities and sources, where complex knowledge reasoning is required to generate accurate answers. The methodology in synthesizing LLMs and KGs for complex QA has been exploited as follows.

2.1 Multi-document QA

Multi-document QA refers to the QA over contexts from multiple documents, while efficiently and effectively retrieving the relevant knowledge from multiple contexts is the main technical challenge. To reduce the retrieval latency and improve the quality of the retrieved context for multi-document QA, KGP (Wang et al., 2024d) introduces an LLM-based graph traversal agent for retrieving relevant knowledge from KG. Similarly, CuriousLLM (Yang and Zhu, 2025) integrates a knowledge graph prompting, reasoning-infused LLM agent, and graph traversal agent to augment LLMs for multi-document QA. VisDom (Suri et al., 2024) introduces a novel multimodal RAG for multi-document question answering by integrating and fusing the multi-modal knowledge and leveraging the (Chain-of-thought) CoT-based reasoning.

2.2 Multi-modal QA

Multi-modal QA refers to the QA over multi-modal data, and visual QA (VQA) is one of the typical multi-modal QA. To retrieve the most relevant knowledge from the external KG for enhancing VQA, MMJG (Wang et al., 2022) introduces an adaptive knowledge selection to jointly select knowledge from visual and text knowledge based on the knowledge-aware attention and multi-modal guidance. To effectively retrieve the evidence from multi-modal data, RAMQA (Bai et al., 2025) enhances multi-modal retrieval-augmented QA by integrating learning-to-rank with training of generative models via multi-task learning. KVQA (Dong et al., 2024b) integrates LLMs with multimodal knowledge by using a two-stage prompting and a

¹<https://github.com/machuangtao/LLM-KG4QA>

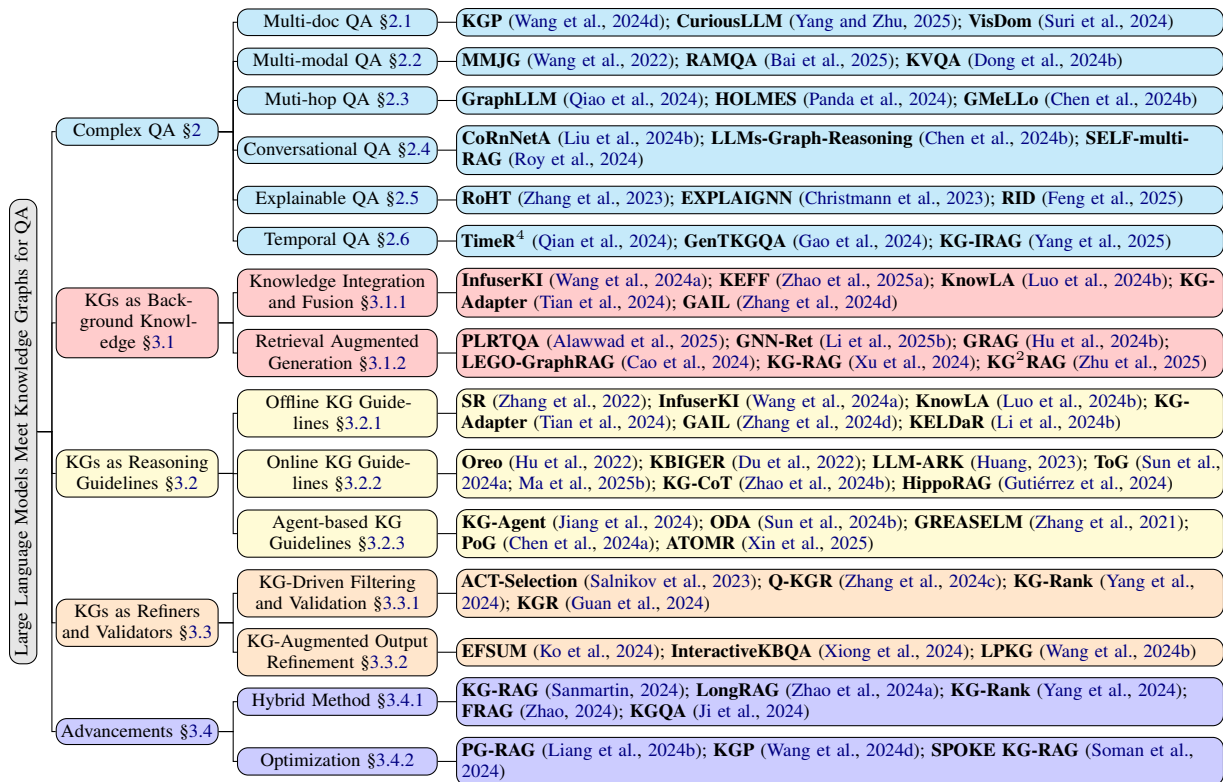


Figure 1: A Structured Taxonomy of Synthesizing LLMs and KGs for QA.

pseudo-siamese graph medium fusion to balance intra-modal and inter-modal reasoning.

2.3 Multi-hop QA

Multi-hop QA differs from simple QA, which usually involves multi-step reasoning to generate the final answers. The basic idea is to decompose the multi-hop questions into multiple single-hop questions, then generate the answers for each single-hop question, and finally integrate them (Linders and Tomczak, 2025). For instance, GraphLLM (Qiao et al., 2024) leverages LLMs to decompose the multi-hop question into several simple sub-questions and retrieve the sub-graphs via GNNs and LLMs to generate the answers for sub-questions based on graph reasoning. HOLME (Panda et al., 2024) utilizes a context-aware retrieved and pruned hyper-relational KG that is constructed based on the entity-document graph to enhance LLMs for generating the answers of multi-hop QA. To enable accurate factual knowledge retrieval and reasoning of LLMs for multi-hop QA, GMeLLO (Chen et al., 2024b) effectively integrates the explicit knowledge of KGs with the linguistic knowledge of LLMs by introducing the fact triple extraction, relation chain extraction, and query and answer generation.

2.4 Conversational QA

The challenge of conversational QA lies in how to make the language model (LM) easily understand the questions and intermediate interactions. To make user interactions easily understood by machines, CoRnNetA (Liu et al., 2024b) introduces an LLM-based question reformulation, reinforcement learning agent, and soft reward mechanism to improve the interpretation of multi-turn interactions with KGs. The conversational QA involves several multi-run QA to refine and get accurate answers through multiple rounds of interactions. The knowledge aggregation module and graph reasoning are introduced for joint reasoning between the graph and LLMs (Jain and Lapata, 2024) to address the challenges of understanding the question and context for conversational QA. To improve the contextual understanding and the answer quality for conversational QA, SELF-multi-RAG (Roy et al., 2024) leverages LLMs to retrieve from the summarized conversational history and reuse the retrieved knowledge for augmentation.

2.5 Explainable QA

Explainable QA (XQA) aims to provide explanations for the generated answers based on the reasoning over the factual KGs. To effectively inte-

grate the multiple sources of knowledge for XQA, RoHT (Zhang et al., 2023) introduces a two-stage XQA method that implements the probabilistic reasoning based on the constructed Hierarchical Question Decomposition Tree (HQDT) from the aggregated knowledge. To trace the provenance and improve the explainability of the answers, EXPLAIGNN (Christmann et al., 2023) constructs a heterogeneous graph from retrieved knowledge and user explanations and generates explanatory evidence based on GNN with question-level attention. RID (Feng et al., 2025) directly integrates the unsupervised retrieval with LLMs based on reinforcement learning-driven knowledge distillation.

2.6 Temporal QA

The challenges of temporal QA lie in fully understanding the implicit time constraints and effectively incorporating them with temporal knowledge for temporal reasoning. To improve the accuracy of LLMs in answering temporal questions, Timer⁴ (Qian et al., 2024) introduces a Retrieve-Retrieve-Rerank pipeline to augment the temporal reasoning of LLMs by temporal knowledge-based fine-tuning. Similarly, GenTKGQA (Gao et al., 2024) introduces a temporal GNN and virtual knowledge indicators to capture temporal knowledge embeddings, and further dynamically integrates retrieved subgraphs into LLMs for temporal reasoning. To facilitate reasoning of LLMs with KGs, KG-IRAG (Yang et al., 2025) enables LLMs to incrementally retrieve knowledge and evaluate its sufficiency for augmenting the capabilities of LLMs in answering time-sensitive and event-based queries involving temporal dependencies.

3 Approaches and Alignments

The strengths and limitations of the approach in synthesizing LLMs and KGs for QA are summarized (details in Table 2), and the research progress on the alignment of these approaches with the complex QA is discussed (details in Table 7).

3.1 KGs as Background Knowledge

KGs usually play the role of background knowledge when synthesizing LLMs for complex QA, where knowledge fusion and RAG are the main technical paradigms (comparison in Table 3).

3.1.1 Knowledge Integration and Fusion

Knowledge integration and fusion aim to enhance language models (LMs) by integrating unknown

knowledge into LMs for QA. As shown in Figure 2,




Figure 2: Knowledge Integration and Fusion.

the KGs and text are aligned via local subgraph extraction and entity linking, and then fed into the cross-model encoder to bidirectionally fuse text and KG to jointly train the language models for complex QA tasks (Yasunaga et al., 2022; Zhang et al., 2021). To address knowledge forgetting and noisy knowledge during knowledge integration, InfuserKI (Wang et al., 2024a) and KEFF (Zhao et al., 2025a) introduce the adaptive selection and knowledge enhancement filter, respectively, which selects the new knowledge and integrates it with LLMs. Fine-tuning LLMs with text and knowledge graphs can improve their performance on specified tasks. For instance, KG-Adapter (Tian et al., 2024) improves parameter-efficient fine-tuning of LLMs by introducing a knowledge adaptation layer to LLMs. GAIL (Zhang et al., 2024d) fine-tunes LLMs for lightweight KGQA models based on retrieved SPARQL-question pairs from KGs.

3.1.2 Retrieval Augmented Generation

RAG serves as a retrieval and augmentation mechanism, as shown in Figure 3. It first retrieves relevant knowledge from the text chunks based on vector-similarity retrieval, and then augments the LLMs by integrating the retrieved context with LLMs (Alawwad et al., 2025; Li et al., 2025b). However, the mainstream RAG methods retrieve

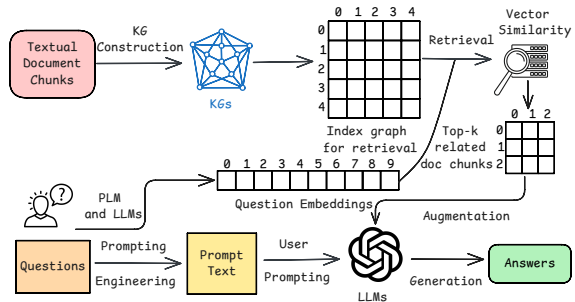


Figure 3: Retrieval Augmented Generation.

the relevant knowledge from the textual chunks,

which ignore the structured information and inter-relations of these textual chunks. To mitigate this limitation, Graph RAG (Hu et al., 2024b) and KG-RAG (Sanmartin, 2024; Xu et al., 2024; Linders and Tomczak, 2025) are proposed. Instead of retrieving the knowledge from textual chunks, Graph RAG directly retrieves the relevant knowledge from graph data. GRAG (Hu et al., 2024b) retrieves the top- k relevant subgraphs from the textual graph and then integrates the retrieved subgraphs with the query by aggregating and aligning the graph embeddings with text embeddings based on GNNs. LEGO-GraphRAG (Cao et al., 2024) decomposes the retrieval process into subgraph-extraction, path-filtering, and path-refinement modules, thereby improving the reasoning capabilities of LLMs with retrieved knowledge. KG²RAG (Zhu et al., 2025) retrieves the relevant subgraph from KG and expanding the textual chunks with the retrieved KG for augmenting the generation.

3.1.3 Aligning with Complex QA

In the approaches of synthesizing LLMs and KGs for QA, where KG plays the role of background knowledge, the retrieved relevant knowledge from the factual knowledge graph can reconcile knowledge conflicts from the multiple documents for multiple-doc QA (Wang et al., 2024d). The cross-modal reasoning can facilitate the cross-modal interaction and alignment for multi-modal QA (Suri et al., 2024). Additionally, the question decomposition of multi-hop QA can be augmented by fusing the knowledge from LLMs and KGs, which further facilitates iterative reasoning for generating the accurate final question (Saleh et al., 2024; Cao et al., 2024). The RAG (Roy et al., 2024) and KG-RAG (Sanmartin, 2024) can also improve the capabilities of LLMs in understanding the user’s interactions for generating accurate answers for conversational QA. However, the key technical challenge behind this methodology is *how to retrieve the relevant knowledge from large-scale KGs and then effectively fuse with LLMs without inducing knowledge conflicts?*

3.2 KGs as Reasoning Guidelines

KGs can provide reasoning guidelines for LLMs to access precise knowledge from factual evidence based on reasoning. Recent methods (**comparison in Table 4**) for integrating the reasoning of KG and LLMs can be classified into three categories.

3.2.1 Offline KG Guidelines

In offline KG guidelines, KG supplies potential paths or subgraphs before the reasoning process by the LLM, from which the LLM selects the most relevant path for reasoning. For instance, SR (Zhang et al., 2022) trains a subgraph retriever that operates independently from the downstream reasoning process, while Keqing (Wang et al., 2023) decomposes complex questions using predefined templates, retrieves candidate entities and triples from a KG. EtD (Liu et al., 2024a) initially uses a GNN to identify promising candidates and extracts relevant fine-grained knowledge pertinent to the questions, and then creates a knowledge-enhanced multiple-choice prompt to guide LLM for generating the final answer. Recent studies have explored the application of novel formats of guidelines. GCR (Luo et al., 2024a) first converts a KG into a KG-Trie and then develops a graph-constrained decoding and a lightweight LLM to generate multiple reasoning paths and candidate answers. KELDaR (Li et al., 2024b) introduces the question decomposition and atomic retrieval modules to extract implicit information and retrieves the relevant subgraphs from KG to augment LLM for QA.

3.2.2 Online KG Guidelines

This paradigm emphasizes that the guidance of the KG directly involves the LLM’s reasoning process, rather than merely relying on existing static knowledge. For example, Oreo (Hu et al., 2022) uses a contextualized random walk across a KG and conducts a single step of reasoning through the specific layers. KBIGER (Du et al., 2022) considers the $(k-1)$ -th reasoning graph to construct the k -th instruction for reasoning and corrects erroneous predictions of intermediate entities. LLM-ARK (Huang, 2023) approaches reasoning tasks as sequential decision-making processes and employs Proximal Policy Optimization (PPO) for optimization. ToG (Sun et al., 2024a) allows LLMs to iteratively perform beam search over KGs, by which the most promising reasoning paths and the most likely reasoning outcomes are generated. In contrast, ToG-2 (Ma et al., 2025b) utilizes entities as intermediaries to guide LLMs toward precise answers based on iterative retrieval between documents and KGs. KG-CoT (Zhao et al., 2024b) leverages external KGs to generate reasoning paths for joint reasoning of LLMs and KGs to enhance the reasoning capabilities of LLMs for QA. To identify the relevant KG subgraphs, HippoRAG (Gutiérrez

et al., 2024) integrates multi-hop reasoning with single-step multi-hop knowledge retrieval.

3.2.3 Agent-based KG Guidelines

KGs can also be integrated into the reasoning process of LLMs as a component within an Agent system, as shown in Figure 4. This integration al-

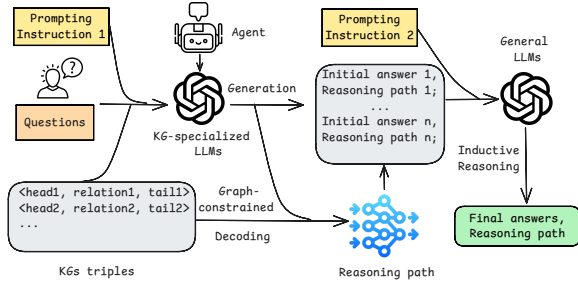


Figure 4: KGs as Reasoning Guidelines.

lows the Agent to leverage structured knowledge for augmenting the decision-making and problem-solving capabilities of LLMs. KG-Agent (Jiang et al., 2024) is a multifunctional toolbox integrating LLMs with a KG-based executor and a knowledge memory system that autonomously selects tools and updates the memory to enhance the reasoning of LLMs over KGs. ODA (Sun et al., 2024b) incorporates KG reasoning capabilities through a global observation approach, which improves reasoning abilities by employing a cyclical paradigm of observation, action, and reflection. GREASELM (Zhang et al., 2021) effectively integrates encoded representations from LMs and GNNs by introducing several modality interaction layers to seamlessly blend structured knowledge with language contexts. PoG (Chen et al., 2024a) integrates reflection and self-correction mechanisms to adaptively explore the reasoning paths over KG via an LLM agent, and then augments the LLM in complex reasoning and question answering based on the retrieved knowledge paths. ATOMR (Xin et al., 2025) leverages LLMs as a reasoning agent to retrieve and incorporate the knowledge across multiple knowledge sources for augmenting the reasoning capability of LLM in knowledge-intensive question answering.

3.2.4 Aligning with Complex QA

The approaches that incorporate KGs with LLMs can enable multi-hop and iterative reasoning over the factual KGs, further augmenting the reasoning capability of LLMs for complex QA. The challenges of knowledge retrieval and conflicts across modalities and knowledge sources, as well as com-

plex knowledge reasoning and question answering in multi-document QA, multi-modal QA, and multi-hop QA, can be mitigated through joint reasoning over factual KGs and LLMs (Suri et al., 2024; Liang et al., 2025; Qiao et al., 2024; Xin et al., 2025). Moreover, the joint reasoning over the factual KGs and LLMs provides the logical inference chains and anchor for LLMs to generate explainable answers with clear evidence from factual KGs (Zhao et al., 2024b). Although the joint reasoning over the factual KGs and LLMs can address the challenges of complex QA well, joint KG-LLM reasoning remains inefficient because large-scale graph traversal is computationally intensive and time-consuming. Moreover, the reasoning capabilities of KGs mainly depend on the completeness and knowledge coverage of KGs, where the incomplete, inconsistent, and outdated knowledge from KGs might induce noise or conflicts. The main challenge lies in *how to improve the reasoning efficiency over the large-scale graph and reasoning capabilities under incomplete KG?*

3.3 KGs as Refiners and Validators

The factual evidence from KGs enables LLMs to refine and verify the intermediate answers, as shown in Figure 5. In these methods (comparison in Table 5), KGs act as refiner and validator.

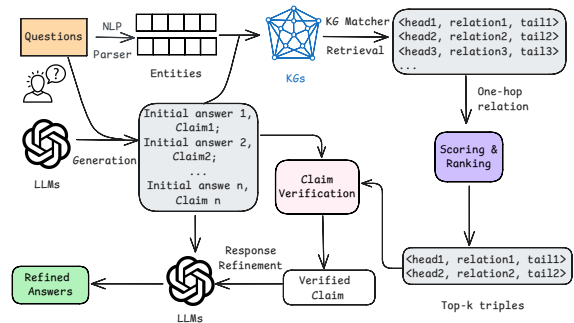


Figure 5: KGs as Refiners and Validators.

3.3.1 KG-Driven Filtering and Validation

KGs enhance the accuracy and reliability of LLM outputs by filtering and validating candidate answers with structured and verified knowledge. For instance, ACT-Selection (Salnikov et al., 2023) filters and re-ranks answer candidates based on their types extracted from Wikidata. Q-KGR (Zhang et al., 2024c) improves the reasoning capabilities of LLMs by filtering out the irrelevant knowledge based on the ranking of the relevance score between

the question and knowledge. KGs can improve the factual accuracy of generated answers, as demonstrated by KG-Rank (Yang et al., 2024), which integrates medical KGs with re-ranking techniques to enhance the credibility of generated responses. Moreover, KGR (Guan et al., 2024) autonomously extracts and validates factual statements, significantly boosting performance on factual QA.

3.3.2 KG-Augmented Output Refinement

Integrating KGs with LLMs is essential to refining the outputs of LLMs for greater clarity and accuracy. EFSUM (Ko et al., 2024) employs LLM as a fact summarizer to generate relevant summaries from KGs, thereby improving performance in zero-shot QA. InteractiveKBQA (Xiong et al., 2024) enables iterative knowledge interactions, allowing LLMs to generate logical forms and refine outputs based on user feedback. LPKG (Wang et al., 2024b) fine-tunes LLMs with KG-derived planning data to enhance the planning capabilities of LLMs in sophisticated reasoning for complex QA.

3.3.3 Aligning with Complex QA

The approaches that leverage the retrieved factual evidence from KGs for refinement and validation are designed to augment the capability of LLMs in understanding user interactions and verifying the intermediate reasoning for multi-hop QA (Chen et al., 2024b) and conversational QA (Xiong et al., 2024), respectively. However, the common case is that the knowledge from KGs is not available or not enough to verify the intermediate results, since the factual knowledge from KGs might be incomplete (Zhou et al., 2025). In particular, the knowledge conflicts between the intermediate answer and KG facts might induce irrelevant results due to the poorly verified intermediate results. Meanwhile, the refinement and validation of results largely depend on the correctness, timeliness, and completeness of factual knowledge in KGs. Thereby, the main challenges of this approach lie in *how to handle the knowledge conflict between intermediate answers and KG facts and incrementally update KGs to ensure factual knowledge in KGs is up-to-date and correct?*

3.4 Advancements

The advanced methods can be categorized into two subcategories: hybrid methods and optimization.

3.4.1 Hybrid Method

In addition to acting as a single role, the approaches (**comparison in Table 6**) where KGs serve multiple roles when synthesizing with LLMs for QA have been investigated. For instance, KG-RAG (Sanmartin, 2024) introduces the Chain-of-Explorations (CoE) to rank and obtain the top- k nodes or relationships based on vector similarity, by which the most relevant knowledge is selected and then fed into LLMs for the final answer. LongRAG (Zhao et al., 2024a) retrieves the top- k most relevant chunks based on a hybrid retriever and analyzes their relevance with the query by introducing a CoT-guided filter. Furthermore, the KG-augmented prompting is introduced to LLMs for augmenting the generation of the final answer. In KG-Rank (Yang et al., 2024), multiple ranking methods are introduced to refine the retrieved triples for augmenting the reasoning with the most relevant knowledge. FRAG (Zhao, 2024) introduces reasoning-aware and flexible-retrieval modules to retrieve reasoning paths from KGs, thereby guiding and augmenting LLM for efficient reasoning and answer generation. KGQA (Ji et al., 2024) combines the CoT-based prompting with graph retrieval to improve the retrieval quality and the multi-hop reasoning capability of LLMs in QA.

3.4.2 Optimization

To mitigate the low efficiency and high computing costs of existing methods in synthesizing LLMs and KGs for complex QA, several optimization techniques (**extension in Appendix § D**) are proposed to improve the efficiency of synthesizing LLMs and KGs. (1) *Index-based optimization*. It aims to accelerate the process of learning embeddings and vector retrieval for given questions and knowledge context. For instance, PG-RAG (Liang et al., 2024b) proposes dynamic and adaptable knowledge retrieval indexes based on LLMs that can effectively handle complex queries and improve the overall performance of RAG systems in QA tasks. (2) *Prompt-based optimization*. It mainly enhances the quality of the prompts and facilitates the knowledge fusion based on prompt engineering. For example, KGP (Wang et al., 2024d) proposes a KG prompting approach to enhance the prompt for LLMs and optimize the knowledge retrieval by introducing the KG construction module and LLM-based graph traversal agent. (3) *Cost-based optimization*. It aims to minimize computation costs by reducing the number of calls to LLMs

and accelerating knowledge retrieval. In particular, SPOKE KG-RAG (Soman et al., 2024) proposes a token-based optimized KG-RAG framework that integrates explicit and implicit knowledge from KG with LLM to enhance LLMs for cost-effective QA.

3.4.3 Aligning with Complex QA

The hybrid method that involves the multiple roles of KG when synthesizing LLMs and KGs can address the limitations and challenges of complex QA in knowledge interactions and complex reasoning across modalities and knowledge sources. The combination of knowledge fusion, RAG, CoT-based reasoning, and ranking-based refinement is capable of accelerating the complex question decomposition (Wang et al., 2023) for multi-hop QA, enhancing the context understanding for conversation QA (Roy et al., 2024), facilitating the interactions across modalities for multi-modal QA (Dong et al., 2024b), and improving the explainability of the generated answers (Christmann et al., 2023). However, the remaining challenge is *how to achieve efficient vector indexing and search over large-scale KGs and maintain a balance between the cost and performance?*

4 Evaluation and Application

We summarize the metrics & dataset and showcase the related applications and demonstrations.

4.1 Evaluation and Benchmark Dataset

The metrics and benchmark datasets are summarized and compared (**details in Appendix §E.1**).

Metrics. The metrics for evaluating the approach in synthesizing LLMs with KGs for QA are summarized: **(1) Answer Quality (AnsQ)**; **(2) Retrieval Quality (RetQ)**; **(3) Reasoning Quality (ReaQ)**.

Benchmark Dataset. The recent benchmark datasets are summarized and compared (**details in Table 8**) with the focus on *Answer Quality*, *Retrieval Quality*, and *Reasoning Quality*.

4.2 Applications

We showcase the industrial and scientific applications and demonstrations (**details in Appendix §E.2**) in synthesizing LLMs with KGs for QA.

5 Open Challenges and Opportunities

Even though the synthesis of LLM and KG can take advantage of their strengths and be mutually beneficial for complex question answering, the current

LLM+KG QA systems may still meet challenges. We summarize the challenges by highlighting the opportunities and discussing the future directions.

Scaling to both Effectiveness and Efficiency.

LLM+KG systems retrieve the factual knowledge and perform multi-hop reasoning under tight latency and memory budgets. Three bottlenecks are emerging: **(1) Structure-aware retrieval:** Vanilla dense or sparse retrieval treats a KG as an unordered triple, thereby discarding topological cues that are vital for pruning the search space (Tian et al., 2025). Hierarchical graph partitioning, dynamic neighbourhood expansion, and learned path-prior proposal networks are promising ways to expose structure to the retriever while keeping the vector index sublinear. **(2) Amortized reasoning:** Current retrieval and prompting pipelines repeatedly query the KG for every Beam search or CoT step. Caching subgraphs, reusing intermediate embeddings, and exploiting incremental-computation friendly hardware can mitigate the quadratic blow-up of iterative reasoning. **(3) Lightweight answer validation:** Most guardrails rely on LLMs, while probabilistic logic programs and bloom filter sketches with KG-based fact-checking offer a lightweight solution. An opportunity is to design a retriever and validator that estimates the uncertainty of retrieval results and then guides the selective execution of validation.

Knowledge Alignment and Dynamic Integration.

Once a KG snapshot is injected into an LLM, it starts to become outdated, just like real-world KGs usually involve adding new entities, deleting relations, and resolving contradictions. Future work should focus on: **(1) Quantify alignment:** We lack metrics that score not only semantic overlap but also *structural* compatibility between parametric knowledge in the LLM and symbolic knowledge in the KG. Contrastive probing with synthetic counterfactuals or topology-aware alignment losses may fill this gap. **(2) Facilitate real-time updates:** Parameter-efficient tuning (e.g. LoRA modules keyed by graph deltas) and retrieval-time patching (streaming KGs with temporal indices) are early steps toward *stream-time* knowledge alignment. **(3) Detect and resolve conflicts:** Bayesian trust networks, source-aware knowledge distillation, and multi-agent debate protocols can estimate and reconcile confidence scores across modalities and sources. Incorporating these into the decoding objective is an open challenge with high pay-off.

Explainable and Fairness-Aware QA. The scale of LLMs poses challenges to explainability and fairness in QA. While integrating KGs offers a path toward interpretable reasoning, it also introduces computational challenges and fairness concerns. Future work may consider the following directions: **(1) Reasoning over subgraphs:** Retrieving subgraphs from large-scale KGs is computationally expensive and often results in overly complex or incomprehensible explanations. Structure-aware retrieval and reranking methods should be employed to identify subgraphs consistent with the gold subgraphs. Furthermore, CoT-based prompting can be used to guide LLMs in generating explicit reasoning steps grounded in the retrieved subgraphs. **(2) Fairness-aware knowledge retrieval:** LLMs can capture social biases from training data, but KGs may contain incomplete or biased knowledge. As a result, the fairness concerns remain in RAG (Wu et al., 2024b). Incorporating fairness-aware techniques into KG retrieval (e.g., reranking based on bias detection) and integrating them with counterfactual prompting can mitigate bias. **(3) Explainable conversational QA:** Single-turn QA restricts the exploration of diverse perspectives and limits the reasoning processes. Developing conversational QA with retrieval strategies can dynamically detect and adjust knowledge biases and further improve the explainability of the RAG-based QA system through multi-turn user interactions.

6 Conclusion

This survey has systematically examined the synthesis of LLMs and KGs in question answering, presenting a novel taxonomy that categorizes methodologies based on QA types and the role of KGs. Our analysis highlights the strengths and limitations of current approaches, emphasizing the potential opportunities of leveraging the KGs to augment LLMs to overcome challenges such as hallucinations, limited reasoning capabilities, and knowledge conflicts in complex QA scenarios. Despite significant advancements, several remaining challenges include efficient knowledge retrieval, dynamic knowledge integration, effective reasoning over knowledge at scale, and explainable and fairness-aware QA. Future research should focus on developing adaptive frameworks that dynamically integrate up-to-date knowledge with LLMs, as well as establishing efficient methods for scaling reasoning, explainability, and fairness.

Limitations

This survey covers the taxonomy and survey of the recent advancements in synthesizing LLMs and KGs for QA, as well as the discussion of its challenges and opportunities. However, we are aware that this survey may miss some newly released works due to the rapid expansion of works on this topic. Moreover, the survey mainly highlights the alignments between the recent methodologies of incorporating LLMs and KGs for QA and the challenges of the various complex QA tasks, while these taxonomies from different perspectives are non-exclusive, and the overlap between the two taxonomies may arise. Furthermore, this survey underemphasizes the quantitative and experimental evaluation of different methodologies due to the various implementation details, the diversity of the benchmark datasets, and non-standardized evaluation metrics.

Acknowledgments

This work was supported by the Novo Nordisk Foundation grant (NNF22OC0072415), the National Natural Science Foundation of China (U23B2057, 62176185, 62376058), and the GOBLIN COST Action (CA23147).

References

- Hessa Abdulrahman Alawwad, Areej Alhothali, Usman Naseem, Ali Alkathlan, and Amani Jamal. 2025. Enhancing textbook question answering task with large language models and retrieval augmented generation. *Pattern Recognit.*, 162:111332.
- Yang Bai, Christan Earl Grant, and Daisy Zhe Wang. 2025. RAMQA: A unified framework for retrieval-augmented multi-modal question answering. *arXiv:2501.13297*.
- Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. 2023. Prompting is programming: A query language for large language models. *Proc. ACM Program. Lang.*, 7:1946–1969.
- Yuzheng Cai, Zhenyue Guo, Yiwen Pei, Wanrui Bian, and Weiguo Zheng. 2024. SimGRAG: Leveraging similar subgraphs for knowledge graphs driven retrieval-augmented generation. *arXiv:2412.15272*.
- Yukun Cao, Zengyi Gao, Zhiyang Li, Xike Xie, and S Kevin Zhou. 2024. LEGO-GraphRAG: Modularizing graph-based retrieval-augmented generation for design space exploration. *arXiv:2411.05844*.
- Liyi Chen, Panrong Tong, Zhongming Jin, Ying Sun, Jieping Ye, and Hui Xiong. 2024a. Plan-on-Graph:

- Self-correcting adaptive planning of large language model on knowledge graphs. In *NeurIPS*, pages 37665–37691.
- Ruirui Chen, Weifeng Jiang, Chengwei Qin, Ishaan Rawal, Cheston Tan, Dongkyu Choi, Bo Xiong, and Bo Ai. 2024b. LLM-based multi-hop question answering with knowledge graph integration in evolving environments. In *EMNLP*, pages 14438–14451.
- Weijie Chen, Ting Bai, Jinbo Su, Jian Luan, Wei Liu, and Chuan Shi. 2024c. KG-Retriever: Efficient knowledge indexing for retrieval-augmented large language models. *arXiv:2412.05547*.
- Zichen Chen, Jianda Chen, Ambuj Singh, and Misha Sra. 2024d. XplainLLM: A knowledge-augmented dataset for reliable grounded explanations in LLMs. In *EMNLP*, pages 7578–7596.
- Philipp Christmann, Rishiraj Saha Roy, and Gerhard Weikum. 2023. Explainable conversational question answering over heterogeneous sources via iterative graph neural networks. In *SIGIR*, pages 643–653.
- Philipp Christmann and Gerhard Weikum. 2024. RAG-based question answering over heterogeneous data and text. *arXiv:2412.07420*.
- Xavier Daull, Patrice Bellot, Emmanuel Bruno, Vincent Martin, and Elisabeth Murisasco. 2023. Complex QA and language models hybrid architectures, survey. *arXiv:2302.09051*.
- Junnan Dong, Qinggang Zhang, Chuang Zhou, Hao Chen, Daochen Zha, and Xiao Huang. 2024a. Cost-efficient knowledge-based question answering with large language models. In *NeurIPS*.
- Junnan Dong, Qinggang Zhang, Huachi Zhou, Daochen Zha, Pai Zheng, and Xiao Huang. 2024b. Modality-aware integration with large language models for knowledge-based visual question answering. In *ACL*, pages 2417–2429.
- Haowei Du, Quzhe Huang, Chen Zhang, and Dongyan Zhao. 2022. Knowledge-enhanced iterative instruction generation and reasoning for knowledge base question answering. In *NLPCC*, pages 431–444.
- Shahul Es, Jithin James, Luis Espinosa Anke, and Steven Schockaert. 2024. RAGAs: Automated evaluation of retrieval augmented generation. In *EACL*, pages 150–158.
- Tianyu Fan, Jingyuan Wang, Xubin Ren, and Chao Huang. 2025. MiniRAG: Towards extremely simple retrieval-augmented generation. *arXiv:2501.06713*.
- Jianzhou Feng, Qin Wang, Huaxiao Qiu, and Lirong Liu. 2025. Retrieval in decoder benefits generative models for explainable complex question answering. *Neural Networks*, 181:106833.
- Yifu Gao, Linbo Qiao, Zhigang Kan, Zhihua Wen, Yongquan He, and Dongsheng Li. 2024. Two-stage generative question answering on temporal knowledge graph using large language models. In *ACL Findings*, pages 6719–6734.
- Hengrui Gu, Kaixiong Zhou, Xiaotian Han, Ninghao Liu, Ruobing Wang, and Xin Wang. 2024. PokeMQA: Programmable knowledge editing for multi-hop question answering. In *ACL*, pages 8069–8083.
- Xinyan Guan, Yanjiang Liu, Hongyu Lin, Yaojie Lu, Ben He, Xianpei Han, and Le Sun. 2024. Mitigating large language model hallucinations via autonomous knowledge graph-based retrofitting. In *AAAI*, pages 18126–18134.
- Willis Guo, Armin Toroghi, and Scott Sanner. 2024. CR-LT-KGQA: A knowledge graph question answering dataset requiring commonsense reasoning and long-tail knowledge. *arXiv:2403.01395*.
- Vivek Gupta, Pranshu Kandoi, Mahek Vora, Shuo Zhang, Yujie He, Ridho Reinanda, and Vivek Sriku-mar. 2023. TempTabQA: Temporal question answering for semi-structured tables. In *EMNLP*, pages 2431–2453.
- Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. HippoRAG: Neurobiologically inspired long-term memory for large language models. In *NeurIPS*, pages 59532–59569.
- Haoyu Han, Yu Wang, Harry Shomer, Kai Guo, Jiayuan Ding, Yongjia Lei, Mahantesh Halappanavar, Ryan A Rossi, Subhabrata Mukherjee, Xianfeng Tang, and 1 others. 2024. Retrieval-augmented generation with graphs (GraphRAG). *arXiv:2501.00309*.
- Jie He, Nan Hu, Wanqiu Long, Jiaoyan Chen, and Jeff Z Pan. 2024. MINTQA: A multi-hop question answering benchmark for evaluating LLMs on new and tail knowledge. *arXiv:2412.17032*.
- Aidan Hogan, Xin Luna Dong, Denny Vrandečić, and Gerhard Weikum. 2025. Large language models, knowledge graphs and search engines: A crossroads for answering users’ questions. *arXiv:2501.06699*.
- Nan Hu, Jiaoyan Chen, Yike Wu, Guilin Qi, Sheng Bi, Tongtong Wu, and Jeff Z Pan. 2024a. Benchmarking large language models in complex question answering attribution using knowledge graphs. *arXiv:2401.14640*.
- Yuntong Hu, Zhihan Lei, Zheng Zhang, Bo Pan, Chen Ling, and Liang Zhao. 2024b. GRAG: Graph retrieval-augmented generation. *arXiv:2405.16506*.
- Ziniu Hu, Yichong Xu, Wenhao Yu, Shuhang Wang, Ziyi Yang, Chenguang Zhu, Kai-Wei Chang, and Yizhou Sun. 2022. Empowering language models with knowledge graph reasoning for question answering. In *EMNLP*, pages 9562–9581.

- Pengcheng Huang, Zhenghao Liu, Yukun Yan, Xiaoyuan Yi, Hao Chen, Zhiyuan Liu, Maosong Sun, Tong Xiao, Ge Yu, and Chenyan Xiong. 2025. PIP-KAG: Mitigating knowledge conflicts in knowledge-augmented generation via parametric pruning. *arXiv:2502.15543*.
- Yuxuan Huang. 2023. Evaluating and enhancing large language models for conversational reasoning on knowledge graphs. *arXiv:2312.11282*.
- Nourhan Ibrahim, Samar Aboulela, Ahmed Ibrahim, and Rasha Kashef. 2024. A survey on augmenting knowledge graphs (KGs) with large language models (LLMs): models, evaluation metrics, benchmarks, and challenges. *Discov. Artif. Intell.*, 4(1):76.
- Parag Jain and Mirella Lapata. 2024. Integrating large language models with graph-based reasoning for conversational question answering. *arXiv:2407.09506*.
- Yixin Ji, Kaixin Wu, Juntao Li, Wei Chen, Mingjie Zhong, Xu Jia, and Min Zhang. 2024. Retrieval and reasoning on KGs: Integrate knowledge graphs into large language models for complex question answering. In *EMNLP*, pages 7598–7610.
- Zhen Jia, Soumajit Pramanik, Rishiraj Saha Roy, and Gerhard Weikum. 2021. Complex temporal question answering on knowledge graphs. In *CIKM*, pages 792–802.
- Jinhao Jiang, Kun Zhou, Wayne Xin Zhao, Yang Song, Chen Zhu, Hengshu Zhu, and Ji-Rong Wen. 2024. Kg-Agent: An efficient autonomous agent framework for complex reasoning over knowledge graph. *arXiv:2402.11163*.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Appl. Sci.*, 11(14):6421.
- Amanda Kau, Xuzeng He, Aishwarya Nambissan, Aland Astudillo, Hui Yin, and Amir Aryani. 2024. Combining knowledge graphs and large language models. *arXiv:2407.06564*.
- Sungho Ko, Hyunjin Cho, Hyungjoo Chae, Jinyoung Yeo, and Dongha Lee. 2024. Evidence-focused fact summarization for knowledge-augmented zero-shot question answering. In *EMNLP*, pages 10636–10651.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. BioASQ-QA: A manually curated corpus for biomedical question answering. *Sci. Data*, 10(1):170.
- Chuhan Li, Ziyao Shangguan, Yilun Zhao, Deyuan Li, Yixin Liu, and Arman Cohan. 2024a. M3SciQA: A multi-modal multi-document scientific QA benchmark for evaluating foundation models. In *EMNLP*, pages 15419–15446.
- Feiyang Li, Peng Fang, Zhan Shi, Arijit Khan, Fang Wang, Dan Feng, Weihao Wang, Xin Zhang, and Yongjian Cui. 2025a. CoT-RAG: Integrating chain of thought and retrieval-augmented generation to enhance reasoning in large language models. *arXiv:2504.13534*.
- Yading Li, Dandan Song, Changzhi Zhou, Yuhang Tian, Hao Wang, Ziyi Yang, and Shuhao Zhang. 2024b. A framework of knowledge graph-enhanced large language model based on question decomposition and atomic retrieval. In *EMNLP Findings*, pages 11472–11485.
- Zijian Li, Qingyan Guo, Jiawei Shao, Lei Song, Jiang Bian, Jun Zhang, and Rui Wang. 2025b. Graph neural network enhanced retrieval for question answering of large language models. In *NAACL*, pages 6612–6633.
- Jianxin Liang, Xiaojun Meng, Huishuai Zhang, Yueqian Wang, Jiansheng Wei, and Dongyan Zhao. 2025. ReasVQA: Advancing VideoQA with imperfect reasoning process. *arXiv:2501.13536*.
- Lei Liang, Mengshu Sun, Zhengke Gui, Zhongshu Zhu, Zhouyu Jiang, Ling Zhong, Yuan Qu, Peilong Zhao, Zhongpu Bo, Jin Yang, and 1 others. 2024a. KAG: Boosting LLMs in professional domains via knowledge augmented generation. *arXiv:2409.13731*.
- Xun Liang, Simin Niu, Sensen Zhang, Shichao Song, Hanyu Wang, Jiawei Yang, Feiyu Xiong, Bo Tang, and Chenyang Xi. 2024b. Empowering large language models to set up a knowledge retrieval indexer via self-learning. *arXiv:2405.16933*.
- Jasper Linders and Jakub M Tomczak. 2025. Knowledge graph-extended retrieval augmented generation for question answering. *arXiv:2504.08893*.
- Guangyi Liu, Yongqi Zhang, Yong Li, and Quanming Yao. 2024a. Explore then determine: A GNN-LLM synergy framework for reasoning over knowledge graph. *arXiv:2406.01145*.
- Lihui Liu, Blaine Hill, Boxin Du, Fei Wang, and Hanghang Tong. 2024b. Conversational question answering with language models generated reformulations over knowledge graph. In *ACL*, pages 839–850.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *NeurIPS*, 35:2507–2521.
- Linhao Luo, Zicheng Zhao, Chen Gong, Gholamreza Haffari, and Shirui Pan. 2024a. Graph-constrained reasoning: Faithful reasoning on knowledge graphs with large language models. *arXiv:2410.13080*.
- Xindi Luo, Zequn Sun, Jing Zhao, Zhe Zhao, and Wei Hu. 2024b. KnowLA: Enhancing parameter-efficient finetuning with knowledgeable adaptation. In *NAACL*, pages 7146–7159.

- Chuangtao Ma, Yongrui Chen, Tianxing Wu, Arijit Khan, and Haofen Wang. 2025a. Unifying large language models and knowledge graphs for question answering: Recent advances and opportunities. In *EDBT*, pages 1174–1177.
- Shengjie Ma, Chengjin Xu, Xuhui Jiang, Muzhi Li, Huaren Qu, and Jian Guo. 2025b. Think-on-Graph 2.0: Deep and interpretable large language model reasoning with knowledge graph-guided retrieval. In *ICLR*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. Generation-augmented retrieval for open-domain question answering. In *ACL/IJCNLP*, pages 4089–4100.
- Jannat Ara Meem, Muhammad Shihab Rashid, Yue Dong, and Vagelis Hristidis. 2024. PAT-questions: A self-updating benchmark for present-anchored temporal question-answering. In *ACL*, pages 13129–13148.
- Lars-Peter Meyer, Johannes Frey, Kurt Junghanns, Felix Brei, Kirill Bulert, Sabine Gründer-Fahrer, and Michael Martin. 2023. Developing a scalable benchmark for assessing large language models in knowledge graph engineering. In *SEMANTICS*.
- Linyong Nan, Weining Fang, Aylin Rasteh, Pouya Lahabi, Weijin Zou, Yilun Zhao, and Arman Cohen. 2024. OMG-QA: Building open-domain multi-modal generative question answering systems. In *EMNLP*, pages 1001–1015.
- Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhania, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omeljanenko, Wen Zhang, Matteo Lissandrini, Russa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, and Damien Graux. 2023. Large language models and knowledge graphs: Opportunities and challenges. *Trans. Graph Data Knowl.*, 1(1):1–38.
- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiaapu Wang, and Xindong Wu. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Trans. Knowl. Data Eng.*, 36(7):3580–3599.
- Pranoy Panda, Ankush Agarwal, Chaitanya Devaguptapu, Manohar Kaul, and Prathosh A P. 2024. HOLMES: Hyper-relational knowledge graphs for multi-hop question answering using LLMs. In *ACL*, pages 13263–13282.
- Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2024. Graph retrieval-augmented generation: A survey. *arXiv:2408.08921*.
- Aleksandr Perevalov, Andreas Both, and Axel-Cyrille Ngonga Ngomo. 2024. Multilingual question answering systems for knowledge graphs—a survey. *Semant. Web*, 15(5):2089–2124.
- Xinying Qian, Ying Zhang, Yu Zhao, Baohang Zhou, Xuhui Sui, Li Zhang, and Kehui Song. 2024. TimeR⁴: Time-aware retrieval-augmented large language models for temporal knowledge graph question answering. In *EMNLP*, pages 6942–6952.
- Zijian Qiao, Nan Li, Chenxi Huang, Gangliang Wang, Shenglin Liang, Hui Lin, and Qinglang Guo. 2024. GraphLLM: A general framework for multi-hop question answering over knowledge graphs using large language models. In *NLPCC*, pages 136–148.
- Nirmal Roy, Leonardo Ribeiro, Rexhina Blloshmi, and Kevin Small. 2024. Learning when to retrieve, what to rewrite, and how to respond in conversational QA. In *EMNLP*, pages 10604–10625.
- Ahmmad OM Saleh, Gökhan Tür, and Yucel Saygin. 2024. SG-RAG: Multi-hop question answering with large language models through knowledge graphs. In *ICNLSP*, pages 439–448.
- Mikhail Salnikov, Maria Lysyuk, Pavel Braslavski, Anton Razzhigaev, Valentin A Malykh, and Alexander Panchenko. 2023. Answer candidate type selection: Text-to-text language model for closed book question answering meets knowledge graphs. In *KONVENS*, pages 155–164.
- Diego Sanmartin. 2024. KG-RAG: Bridging the gap between knowledge and creativity. *arXiv:2405.12035*.
- Juan Sequeda, Dean Allemang, and Bryon Jacob. 2024. A benchmark to understand the role of knowledge graphs on large language model’s accuracy for question answering on enterprise SQL databases. In *GRADES-NDA@SIGMOD/PODS*, pages 1–12.
- Zhengliang Shi, Lingyong Yan, Weiwei Sun, Yue Feng, Pengjie Ren, Xinyu Ma, Shuaiqiang Wang, Dawei Yin, Maarten de Rijke, and Zhaochun Ren. 2025. Direct retrieval-augmented optimization: Synergizing knowledge selection and language models. *arXiv:2505.03075*.
- Moein Shirdel, Joel Rorseth, Parke Godfrey, Lukasz Golab, Divesh Srivastava, and Jarek Szlichta. 2025. AprèsCoT: Explaining LLM answers with knowledge graphs and chain of thought. In *EDBT*, pages 1142–1145.
- Karthik Soman, Peter W Rose, John H Morris, Rabia E Akbas, Brett Smith, Braian Peetoom, Catalina Villouta-Reyes, Gabriel Ceron, Yongmei Shi, Angela Rizk-Jackson, and 1 others. 2024. Biomedical knowledge graph-optimized prompt generation for large language models. *Bioinformatics*, 40(9):btac560.
- Zirui Song, Bin Yan, Yuhua Liu, Miao Fang, Mingzhe Li, Rui Yan, and Xiuying Chen. 2025. Injecting domain-specific knowledge into large language models: a comprehensive survey. *arXiv:2502.10708*.

- Daniel Steinigen, Roman Teucher, Timm Heine Ruland, Max Rudat, Nicolas Flores-Herr, Peter Fischer, Nikola Milosevic, Christopher Schymura, and Angelo Ziletti. 2024. Fact Finder—enhancing domain expertise of large language models by incorporating knowledge graphs. *arXiv:2408.03010*.
- Miao Su, ZiXuan Li, Zhuo Chen, Long Bai, Xiaolong Jin, and Jiafeng Guo. 2024. Temporal knowledge graph question answering: A survey. *arXiv:2406.14191*.
- Yuan Sui and Bryan Hooi. 2024. Can knowledge graphs make large language models more trustworthy? An empirical study over open-ended question answering. *arXiv:2410.08085*.
- Jiashuo Sun, Chengjin Xu, Luminyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, and Jian Guo. 2024a. Think-on-Graph: Deep and responsible reasoning of large language model with knowledge graph. In *ICLR*.
- Lei Sun, Zhengwei Tao, Youdi Li, and Hiroshi Arakawa. 2024b. ODA: Observation-driven agent for integrating LLMs and knowledge graphs. In *ACL*, pages 7417–7431.
- Manan Suri, Puneet Mathur, Franck Dernoncourt, Kanika Goswami, Ryan A Rossi, and Dinesh Manocha. 2024. VisDoM: Multi-document QA with visually rich elements using multimodal retrieval-augmented generation. *arXiv:2412.10704*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *NAACL*, pages 4149–4158.
- Dehao Tao, Feng Huang, Yongfeng Huang, and Minghu Jiang. 2024. Clue-Guided Path Exploration: Optimizing knowledge graph retrieval with large language models to address the information black box challenge. *arXiv:2401.13444*.
- Wen tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *ACL*, pages 201–206.
- Shiyu Tian, Yangyang Luo, Tianze Xu, Caixia Yuan, Huixing Jiang, Chen Wei, and Xiaojie Wang. 2024. KG-Adapter: Enabling knowledge graph integration in large language models through parameter-efficient fine-tuning. In *ACL Findings*, pages 3813–3828.
- Shiyu Tian, Shuyue Xing, Xingrui Li, Yangyang Luo, Caixia Yuan, Wei Chen, Huixing Jiang, and Xiaojie Wang. 2025. A systematic exploration of knowledge graph alignment with large language models in retrieval augmented generation. In *AAAI*, pages 25291–25299.
- Chaojie Wang, Yishi Xu, Zhong Peng, Chenxi Zhang, Bo Chen, Xinrun Wang, Lei Feng, and Bo An. 2023. keqing: knowledge-based question answering is a nature chain-of-thought mentor of LLM. *arXiv:2401.00426*.
- Fali Wang, Runxue Bao, Suhang Wang, Wenchao Yu, Yanchi Liu, Wei Cheng, and Haifeng Chen. 2024a. Infuserki: Enhancing large language models with knowledge graphs via infuser-guided knowledge integration. In *EMNLP*, pages 3675–3688.
- Jianfeng Wang, Anda Zhang, Huifang Du, Haofen Wang, and Wenqiang Zhang. 2022. Knowledge-enhanced visual question answering with multimodal joint guidance. In *IJCKG*, pages 115–120.
- Jinyu Wang, Jingjing Fu, Rui Wang, Lei Song, and Jiang Bian. 2025. PIKE-RAG: sPecialized KnowledgeE and Rationale Augmented Generation. *arXiv:2501.11551*.
- Junjie Wang, Mingyang Chen, Binbin Hu, Dan Yang, Ziqi Liu, Yue Shen, Peng Wei, Zhiqiang Zhang, Jinjie Gu, Jun Zhou, Jeff Z. Pan, Wen Zhang, and Hua-jun Chen. 2024b. Learning to plan for retrieval-augmented large language models from knowledge graphs. In *EMNLP*, pages 7813–7835.
- Minzheng Wang, Longze Chen, Fu Cheng, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang, and Yongbin Li. 2024c. Leave no document behind: Benchmarking long-context LLMs with extended multi-doc QA. In *EMNLP*, pages 5627–5646.
- Yu Wang, Nedim Lipka, Ryan A. Rossi, Alexa Siu, Ruiyi Zhang, and Tyler Derr. 2024d. Knowledge graph prompting for multi-document question answering. In *AAAI*, volume 38, pages 19206–19214.
- Yifan Wei, Yisong Su, Huanhuan Ma, Xiaoyan Yu, Fangyu Lei, Yuanzhe Zhang, Jun Zhao, and Kang Liu. 2023. MenatQA: A new dataset for testing the temporal comprehension and reasoning abilities of large language models. In *EMNLP*, pages 1434–1447.
- Min Wu, Xinglu Yi, Hui Yu, Yu Liu, and Yujue Wang. 2022. Nebula graph: An open source distributed graph database. *arXiv:2206.07278*.
- Shirley Wu, Shiyu Zhao, Michihiro Yasunaga, Kexin Huang, Kaidi Cao, Qian Huang, Vassilis Ioannidis, Karthik Subbian, James Y Zou, and Jure Leskovec. 2024a. Stark: Benchmarking LLM retrieval on textual and relational knowledge bases. *NeurIPS*, 37:127129–127153.
- Xuyang Wu, Shuwei Li, Hsin-Tai Wu, Zhiqiang Tao, and Yi Fang. 2024b. Does RAG introduce unfairness in LLMs? evaluating fairness in retrieval-augmented generation systems. *arXiv:2409.19804*.
- Yurong Wu, Yan Gao, Bin Zhu, Zineng Zhou, Xiaodi Sun, Sheng Yang, Jian-Guang Lou, Zhiming Ding, and Linjun Yang. 2024c. StraGo: Harnessing strategic guidance for prompt optimization. In *EMNLP*, pages 10043–10061.

- Amy Xin, Jinxin Liu, Zijun Yao, Zhicheng Lee, Shulin Cao, Lei Hou, and Juanzi Li. 2025. Atomr: Atomic operator-empowered large language models for heterogeneous knowledge reasoning. In *KDD*, pages 3344–3355.
- Guanming Xiong, Junwei Bao, and Wen Zhao. 2024. Interactive-KBQA: Multi-turn interactions for knowledge base question answering with large language models. In *ACL*, pages 10561–10582.
- Chuan Xu, Qiaosheng Chen, Yutong Feng, and Cheng Gong. 2025a. mmRAG: A modular benchmark for retrieval-augmented generation over text, tables, and knowledge graphs. *arXiv:2505.11180*.
- Tianyang Xu, Haojie Zheng, Chengze Li, Haoxiang Chen, Yixin Liu, Ruoxi Chen, and Lichao Sun. 2025b. NodeRAG: Structuring graph-based rag with heterogeneous nodes. *arXiv:2504.11544*.
- Zhentao Xu, Mark Jerome Cruz, Matthew Guevara, Tie Wang, Manasi Deshpande, Xiaofeng Wang, and Zheng Li. 2024. Retrieval-augmented generation with knowledge graphs for customer service question answering. In *SIGIR*, pages 2905–2909.
- Rui Yang, Haoran Liu, Edison Marrese-Taylor, Qingcheng Zeng, Yuhe Ke, Wanxin Li, Lechao Cheng, Qingyu Chen, James Caverlee, Yutaka Matsuo, and Irene Li. 2024. KG-Rank: Enhancing large language models for medical QA with knowledge graphs and ranking techniques. In *BioNLP Workshop@ACL*, pages 155–166.
- Ruiyi Yang, Hao Xue, Imran Razzak, Hakim Hacid, and Flora D Salim. 2025. KG-IRAG: A knowledge graph-based iterative retrieval-augmented generation framework for temporal reasoning. *arXiv:2503.14234*.
- Zukang Yang and Zixuan Zhu. 2025. CuriousLLM: Elevating multi-document QA with reasoning-infused knowledge graph prompting. In *NAACL*, pages 274–286.
- Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy S Liang, and Jure Leskovec. 2022. Deep bidirectional language-knowledge graph pretraining. In *NeurIPS*, pages 37309–37323.
- Hao Yu, Aoran Gan, Kai Zhang, Shiwei Tong, Qi Liu, and Zhaofeng Liu. 2024. Evaluation of retrieval-augmented generation: A survey. In *Big Data*, pages 102–120.
- Hong Qing Yu and Frank McQuade. 2025. RAG-KG-IL: A multi-agent hybrid framework for reducing hallucinations and enhancing llm reasoning through rag and incremental knowledge graph learning integration. *arXiv:2503.13514*.
- Gongbo Zhang, Zihan Xu, Qiao Jin, Fangyi Chen, Yilu Fang, Yi Liu, Justin F Rousseau, Ziyang Xu, Zhiyong Lu, Chunhua Weng, and Yifan Peng. 2024a. A mapreduce approach to effectively utilize long context information in retrieval augmented language models. *arXiv:2412.15271*.
- Haozhen Zhang, Tao Feng, and Jiaxuan You. 2024b. Graph of records: Boosting retrieval augmented generation for long-context summarization with graphs. *arXiv:2410.11001*.
- Jiajie Zhang, Shulin Cao, Tingjian Zhang, Xin Lv, Juanzi Li, Lei Hou, Jiaxin Shi, and Qi Tian. 2023. Reasoning over hierarchical question decomposition tree for explainable question answering. In *ACL*, pages 14556–14570.
- Jing Zhang, Xiaokang Zhang, Jifan Yu, Jian Tang, Jie Tang, Cuiping Li, and Hong Chen. 2022. Subgraph retrieval enhanced model for multi-hop knowledge base question answering. In *ACL*.
- Michael Zhang and Eunsol Choi. 2021. SituatedQA: Incorporating extra-linguistic contexts into QA. In *EMNLP*, pages 7371–7387.
- Qinggang Zhang, Shengyuan Chen, Yuanchen Bei, Zheng Yuan, Huachi Zhou, Zijin Hong, Junnan Dong, Hao Chen, Yi Chang, and Xiao Huang. 2025. A survey of graph retrieval-augmented generation for customized large language models. *arXiv:2501.13958*.
- Xikun Zhang, Antoine Bosselut, Michihiro Yasunaga, Hongyu Ren, Percy Liang, Christopher D Manning, and Jure Leskovec. 2021. GreaseLM: Graph reasoning enhanced language models for question answering. In *ICLR*.
- Yu Zhang, Kehai Chen, Xuefeng Bai, Zhao Kang, Qianjiang Guo, and Min Zhang. 2024c. Question-guided knowledge graph re-scoring and injection for knowledge graph question answering. In *EMNLP*, pages 8972–8985.
- Zhiqiang Zhang, Liqiang Wen, and Wen Zhao. 2024d. A GAIL fine-tuned LLM enhanced framework for low-resource knowledge graph question answering. In *CIKM*, pages 3300–3309.
- Dongfang Zhao. 2024. FRAG: Toward federated vector database management for collaborative and secure retrieval-augmented generation. *arXiv:2410.13272*.
- Qi Zhao, Qi Song, Tian Xie, Haiyue Zhang, Hongyu Yang, and Xiangyang Li. 2025a. Improving pre-trained language models with knowledge enhancement and filtering framework. In *NAACL Findings*, pages 3860–3871.
- Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong, and Jie Tang. 2024a. LongRAG: A dual-perspective retrieval-augmented generation paradigm for long-context question answering. In *EMNLP*, pages 22600–22632.
- Ruilin Zhao, Feng Zhao, Long Wang, Xianzhi Wang, and Guandong Xu. 2024b. KG-CoT: Chain-of-thought prompting of large language models over knowledge graphs for knowledge-aware question answering. In *IJCAI*, pages 6642–6650.

Xuejiao Zhao, Siyan Liu, Su-Yin Yang, and Chunyan Miao. 2025b. MedRAG: Enhancing retrieval-augmented generation with knowledge graph-elicited reasoning for healthcare copilot. In *WWW*, pages 4442–4457.

Shangshang Zheng, He Bai, Yizhe Zhang, Yi Su, Xiaochuan Niu, and Navdeep Jaitly. 2024a. KGLens: Towards efficient and effective knowledge probing of large language models with knowledge graphs. *arXiv:2312.11539*.

Xinxin Zheng, Feihu Che, Jinyang Wu, Shuai Zhang, Shuai Nie, Kang Liu, and Jianhua Tao. 2024b. KS-LLM: Knowledge selection of large language models with evidence document for question answering. *arXiv:2404.15660*.

Dongzhuoran Zhou, Yuqicheng Zhu, Yuan He, Jiaoyan Chen, Evgeny Kharlamov, and Steffen Staab. 2025. Evaluating knowledge graph based retrieval augmented generation methods under knowledge incompleteness. *arXiv:2504.05163*.

Andrew Zhu, Alyssa Hwang, Liam Dugan, and Chris Callison-Burch. 2024. FanOutQA: A multi-hop, multi-document question answering benchmark for large language models. In *ACL*, pages 18–37.

Xiangrong Zhu, Yuexiang Xie, Yi Liu, Yaliang Li, and Wei Hu. 2025. Knowledge graph-guided retrieval augmented generation. *arXiv:2502.06864*.

A Related Survey and Paper Selection

A.1 Related Survey

Previous surveys have drawn the roadmap of unifying LLMs+KGs (Pan et al., 2024), discussed the opportunities and challenges (Pan et al., 2023) of leveraging LLMs for knowledge extraction, ontologies and KGs constructions, summarized the integration and synthesis paradigms of LLMs and KGs (Kau et al., 2024; Ibrahim et al., 2024), and gave an overview of the knowledge injection methods between LLMs and domain-specific knowledge (Song et al., 2025). Additionally, the existing surveys give an overview of multilingual knowledge graph question answering (Perevalov et al., 2024), review temporal knowledge graph QA (Su et al., 2024), complex QA (Daull et al., 2023), and discuss search engines, KGs and LLMs from the perspective of user information seeking for QA (Hogan et al., 2025). The comparative analysis of these surveys across LLMs, KGs, LLMs+KGs, GraphRAG, and QA is given in Table 1.

We summarize that the previous surveys on synthesizing LLMs and KGs for QA show limitations in terms of the scope and tasks. (1) The scope

of the surveys is dedicated for general knowledge-intensive tasks, such as knowledge extraction, entity linking, KG construction and completion, and text generation, etc; (2) the QA task in the surveys is limited to the close-domain QA over the knowledge graph; (3) the survey on integrating LLMs, KGs, and search engines for answering user’s questions is from the perspective of user-centric information needs.

A.2 Paper Selection

We first retrieve research papers since 2021 by using Google Scholar and PaSa² using the phrases “knowledge graph and language model for question answering” and “KG and LLM for QA”, and extend the search scope of the benchmark dataset paper to 2016, and then screen and select them based on their relevancy and publication venue quality.

B Taxonomy

B.1 Complex QA

We classify the complex QA into the following categories according to the technical challenges.

- **Multi-document QA.** Multi-document QA (Multi-doc QA) retrieves and synthesizes relevant information from various sources to provide a comprehensive answer.
- **Multi-modal QA.** It refers to the QA over data and knowledge involving multiple modalities such as text, audio, images, video, etc.
- **Muti-hop QA.** It usually involves decomposing the complex question and generating the final answers based on multi-step and iterative reasoning over a factual KG.
- **Conversational QA.** It involves user engagement in multi-turn QA to understand the given context, determine the final answers, and satisfy their information needs.
- **Explainable QA.** XQA (Explainable QA) not only gives the answers to the question but also provides explanations for the given answers.
- **Temporal QA.** It refers to the questions with temporal intent over the temporal KGs that have entities, relations, and associated temporal conditions.

B.2 Approaches of Synthesizing LLMs and KGs for QA

We categorize the methodology of synthesizing LLMs and KGs for QA based on the role of KGs.

²<https://pasa-agent.ai/>

| Survey | KGs | LLMs | LLMs+KGs | Graph RAG | QA |
|---|-----|------|----------|-----------|----|
| Unifying LLMs and KGs: Roadmap (Pan et al., 2024) | ✓ | ✓ | ✓ | ✓ | ✗ |
| LLMs+KGs: Opportunities and Challenges (Pan et al., 2023) | ✓ | ✓ | ✓ | ✓ | ✗ |
| Combing KGs and LLMs (Kau et al., 2024) | ✓ | ✓ | ✓ | ✓ | ✗ |
| mQAKG (Perevalov et al., 2024) | ✓ | ✗ | ✗ | ✗ | ✓ |
| TKGQA (Su et al., 2024) | ✓ | ✗ | ✗ | ✗ | ✓ |
| LLMs, KGs, and SEs (Hogan et al., 2025) | ✓ | ✓ | ✓ | ✗ | ✓ |
| Complex QA (Daull et al., 2023) | ✓ | ✓ | ✗ | ✗ | ✓ |
| Knowledge Injection (Song et al., 2025) | ✓ | ✓ | ✗ | ✗ | ✓ |
| [This work] | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 1: Comparison of Existing Survey Across LLMs, KGs, LLMs+KGs, GraphRAG, and QA. mKGQA: multi-lingual Question Answering for Knowledge Graphs, TKGQA: Temporal Knowledge Graph QA. ✓ Covered or Dicussed, ✗ Not Covered or Dicussed.

- **KG as Background Knowledge.** When KGs are used as background knowledge to enhance LLMs for QA, the questions are parsed to identify the relevant subgraphs from KGs, which are then incorporated with LLMs based on knowledge integration and fusion.
- **KGs as Reasoning Guidelines.** KGs can serve as reasoning guidelines for LLMs for QA tasks by providing structured real-world facts. This factual knowledge and their reliable reasoning paths can improve the explainability of the generated answers.
- **KGs as Refiners and Validators.** KGs can be treated as refiners and validators for LLMs in QA tasks, where LLMs can verify initial answers with factual knowledge and filter out inaccurate responses by integrating KG to ensure that the final responses are precise.
- **Hybrid Methods.** The hybrid methods involve a KG’s multiple roles, namely, background knowledge, reasoning guidelines, refiner and validator, when synthesizing LLMs and KGs for QA.

We summarize the strengths and limitations of each method of synthesizing LLMs and KGs for QA in Table 2 according to our previous taxonomy.

C Summary and Alignment

C.1 Summary Tables of Approaches

The detailed summarization and comparison of the different categories of approaches in terms of main techniques, language model (LM), knowledge graph (KG), dataset, QA types, and evaluation metrics are given in the following tables.

KG as Background Knowledge. Table 3 gives

a summary and comparison of approaches in the category of KG as background knowledge.

KG as Reasoning Guidelines. Table 4 gives a summary and comparison of approaches in the category of KG as reasoning guidelines.

KG as Refiners and Validators. Table 5 gives a summary and comparison of approaches in the category KG as refiners and validators.

Hybrid Approach. Table 6 gives a summary and comparison of hybrid approaches.

C.2 Alignment of Approaches to Complex QA

The alignment of the existing approaches in synthesizing LLMs and KGs to diverse complex QA is presented in Table 7.

D Extended Optimization

The extended optimization techniques are summarized and discussed as follows.

Index-based Optimization. It is a time-consuming and complex task to create vector indexes from long-range facts and retrieve the relevant knowledge from large-scale graphs. To address this issue, GoR (Zhang et al., 2024b) leverages GNN and BERT score-based objectives to optimize node embeddings during graph indexing. KG-Retriever (Chen et al., 2024c) leverages a hierarchical index graph to enhance knowledge correlations and improve information retrieval for efficient knowledge indexing. NodeRAG (Xu et al., 2025b) integrates the heterogeneous graphs and fine-grained retrieval with RAG by optimizing the indexing of graph structures. DRO (Shi et al., 2025) proposes a direct retrieval-agumentation method to directly estimate the distribution of document permutations from the selected model for jointly

| Approach | Strength | Limitation | KG Requirement |
|------------------------------|-------------------------|------------------------|-------------------------|
| KG as Background Knowledge | Broad Coverage | Static Knowledge | High Domain Coverage |
| KG as Reasoning Guidelines | Multi-hop Capabilities | Computational Overhead | Rich Relational Paths |
| KG as Refiners and Validator | Hallucination Reduction | Validation Latency | High Accuracy & Recency |
| Hybrid Approach | Limitation Mitigation | High Computing Cost | Dynamic Adaptation |

Table 2: Comparison of Approaches with Different Roles of KGs.

| Methods | Techniques | LM(s) | KG(s) | Dataset(s) | QA Type | Metric(s) |
|--|---|----------------------------------|-------------------------------|---------------------------------|------------------------|--|
| InfuserKI (Wang et al., 2024a) | Knowledge-based Fine Tuning | Llama-2-7B | UMLS, Movie KG (MetaQA) | PubMedQA, MetaQA-1HopQA | KGQA | NR, RR, F1 |
| KnowLA (Luo et al., 2024b) | Knowledgeable Adaptation | Llama2-7B, Alpaca2 | WordNet, ConceptNet, Wikidata | CSQA, SIQA, BBH, WQSP, TriviaQA | MCQA, CBQA, TruthfulQA | Acc, CE Score, BLEU, ROUGE |
| KG-Adapter (Tian et al., 2024) | Parameter Efficient Fine Tuning & Joint Reasoning | Llama-2-7B-base, Zephyr-7B-alpha | ConceptNet, Freebase | OBQA, CSQA, WQSP, CWQ | KGQA, MCQA, OBQA, CWQ | Acc, Hits@1 |
| GAIL (Zhang et al., 2024d) | GAIL Fine-tuning | Llama-2-7B, BERTa | Freebase | WQSP, CWQ, GraiQA | KGQA | EM, F1, Hits@1 |
| GRAG (Hu et al., 2024b) | Textual Graph RAG | Llama-2-7B | WebQSP, Expla-Graphs | WQSP, CWQ, GraiQA | KGQA | F1, Hits@1, Acc |
| LEGO-GraphRAG (Cao et al., 2024) | Modular Graph RAG | Qwen2-72B, Sentence Transformer | Freebase | WQSP, CWQ, GraiQA | KBQA, CWQ | R, F1, Hits@1 |
| KG²-RAG (Zhu et al., 2025) | Graph-guided Chunks Expansion | Llama-3-8B | Dataset Inherent KGs | HotpotQA | Multi-hop QA | F1, P, R |
| KG-RAG (Xu et al., 2024) | Vector-based Subgraph Retrieval | GPT-4 | Self-constructed KGs | Curated Dataset | KGQA | MRR, Recall@K, NDCG@K, BLEU, ROUGE, METEOR |

Table 3: Summary and Comparison of Methods in the Category of KG as Background Knowledge. NR: Newly-learned Rate, RR: Remembering Rate, CE Score: Cross Entropy Score, EM: Exact Match, MRR: Mean Reciprocal Rank, BLEU: Bilingual Evaluation Understudy, ROUGE: Recall-Oriented Understudy, NDCG: Normalized Discounted Cumulative Gain, METEOR: Metric for Evaluation of Translation with Explicit Ordering.

learning the selection model and generative LM model.

Prompt-based Optimization. To facilitate the deep fusion between the retrieved knowledge from KGs and the internal knowledge of LLMs, several prompt-based optimization approaches are proposed. For instance, StraGo (Wu et al., 2024c) enhances the quality and stability of the prompts based on the insights and strategic guidance learned from the historical prompts by using in-context learning. Meanwhile, several ranking strategies have been exploited to retrieve the top relevant knowledge from diverse knowledge bases, and the most relevant contexts are further fed to LLMs together with the prompts. For example, KG-Rank (Yang et al., 2024) leverages re-ranking techniques based on the score measuring relevance and redundancy to rank the top- k triples from KGs and then combine them with the prompt to generate the answers for QA. Similarly, KS-LLM (Zheng et al., 2024b) introduces the evidence sentence selection module to retrieve the most relevant evidence sentences based on the ranking of the Euclidean distance between the triples and each evidence sentence. Instead of direct ranking of context-based relevance, BriefContext (Zhang et al., 2024a) in-

troduces the preflight check to predict the relevance ranking between the user query and the retrieved documents, which is divided into multiple chunks for multiple RAG subtasks based on the map-reduce strategy. To leverage the multi-source knowledge for RAG-based QA, QUASAR (Christmann and Weikum, 2024) enhances RAG-based QA by effectively integrating unstructured text, structured tables, and knowledge graphs for evidence retrieval and re-ranking and filtering the relevant evidence from the retrieved evidence.

Cost-based Optimization. To implement efficient and effective knowledge probing of LLMs, GLens (Zheng et al., 2024a) initially leverages the Thompson sampling strategy to measure the alignment between KGs and LLMs for addressing the knowledge blind spots of LLMs, and then designs a graph-guided question generator to convert KGs to text, together with a sampling strategy on the parameterized KG structure to accelerate KG traversal. Similarly, Coke (Dong et al., 2024a) minimizes the calls of LLMs for KGQA by introducing a cluster-level Thompson sampling to formulate the accuracy expectation and an optimized context-aware policy to distinguish the expert model based on question semantics. LMQL (Beurer-Kellner

| Methods | Techniques | LM(s) | KG(s) | Dataset(s) | QA Type | Metric(s) | |
|-------------------------------|---|--|---------------------------------|------------|---|--|----------------|
| SR (Zhang et al., 2022) | Trainable Subgraph Retriever, Fine-tuning | RoBERTa-base | Dataset KGs | Inherent | WQSP, CWQ | KBQA | Hits@1, F1 |
| KELDaR (Li et al., 2024b) | Question Decomposition Tree, Atomic KG Retrieval | GPT-3.5-Turbo, GPT-4-Turbo | Dataset KGs | Inherent | WQSP, CWQ | KGQA, Multi-hop QA | EM |
| Oreo (Hu et al., 2022) | Knowledge Interaction and Injection, KG Random Walk | RoBERTa-base, T5-base | Dataset KGs, Wikidata | Inherent | NQ, WQ, WQSP, TriviaQA, CWQ, HotpotQA | CBQA, OBQA, Multi-hop QA | Acc |
| KBIGER (Du et al., 2022) | Iterative Instruction Reasoning | LSTM based Pre-trained Model | Dataset KGs, Freebase | Inherent | WQSP, CWQ, GrailQA | Multi-hop KBQA | Hits@1, F1 |
| ToG (Sun et al., 2024a) | Beam-search base Retrieval, LLM Agent | GPPT-3.5-Turbo, GPT-4, Llama-2-70B-Chat Sentence Transformer | Dataset KGs, Freebase, Wikidata | Inherent | CWQ, WQSP, GrailQA, QALD10-en, WQ | KBQA, Open-Domain QA | Hits@1 |
| ToG-2 (Ma et al., 2025b) | Hibrid RAG, Knowledge-guided Context Retrieval | GPT-3.5-Turbo, GPT-4o, Llama-3-8B, Qwen2-7B | Wikipedia, Wikidata | | WQSP, QALD10-en, AdvHotpotQA, HotpotQA, ToG-FinQA | Multi-hop KBQA & Document QA, Domain-specific QA | Acc, EM, R, F1 |
| KG-CoT (Zhao et al., 2024b) | CoT-based KG and LLM Joint Reasoning | GPT-4, GPT-3.5-Turbo, Llama-7B, Llama-13B | Dataset KGs, Freebase | Inherent | WQSP, CWQ, SQ, WQ | KBQA, Multi-hop QA | Acc, Hit@K |
| KG-Agent (Jiang et al., 2024) | KG-Agent based Instruction Tuning | Dvinci-003, GPT-4, Llama-2-7B | NQ-Wiki, TQ-Wiki, WQ-Freebas | | WQSP, CWQ, GrailQA | KGQA, ODQA | Hits@1, F1 |
| ODA (Sun et al., 2024b) | ODA-based Knowledge Graph Retrieval | GPT-4, GPT-3.5 | Wikidata | | QALD10-en | KBQA | Hits@1, Acc |
| GREASELM (Zhang et al., 2021) | Mint-based KG and LM Cross-modal Fusion and Pretraing | RoBERTa-Large, AristoRoBERTA, SapBERT, PubmedBERT | ConcepNet, DrugBank | UMLS | CQA, OBQA, MedQAt | Multiple-choice QA | Acc |

Table 4: Summary and Comparison of Methods in the Category of KG as Reasoning Guidelines. EM: Exact Match, Acc: Accuracy, R: Recall.

et al., 2023) minimizes the call of LLMs by generating an efficient inference procedure based on the LMP (Language Model Programming) constraints and control flow. It differs from the above work in that CGPE (Tao et al., 2024) optimizes the knowledge retrieval based on clue-guided path exploration and information matching from knowledge bases to enhance the capabilities of LLMs for unfamiliar questions and reduce the costs of LLMs.

To summarize, even if several optimizations and ranking strategies have been recently investigated to reduce the costs of graph retrieval, graph reasoning, and the length of the context of LLMs. However, the relevant subgraphs extraction, graph reasoning, and vector-based retrieval remain a computationally costly task.

E Evaluations and Applications

The details of evaluation metrics, benchmark datasets, industrial and scientific applications, and demonstrations of synthesizing LLMs and KGs for QA are summarized and compared as follows.

E.1 Evaluations

Metrics. The evaluation metrics in synthesizing LLMs with KGs for QA are: **(1) the metrics measuring the answer quality**, BERTScore (Peng et al., 2024), answer relevance (AR), hallucination (HAL) (Yang et al., 2025), accuracy matching, human-verified completeness (Yu and McQuade,

2025); **(2) the metrics measuring the retrieval quality of RAG**, context relevance (Es et al., 2024), faithfulness score (FS) (Yang et al., 2024), precision, context recall (Yu et al., 2024; Huang et al., 2025), mean reciprocal rank (MRR) (Xu et al., 2024), normalized discounted cumulative gain (NDCG) (Xu et al., 2024); **(3) the metrics measuring the reasoning quality for multi-hop QA**, Hop-Acc (Gu et al., 2024), reasoning accuracy (RA) (Li et al., 2025a).

Benchmark Dataset. The recent benchmark datasets are summarized and compared in Table 8 with the focus on the following evaluations. **(1) Answer Quality (AnsQ):** the correctness of the generated answer with supported evidence and retrieved context in comparison to the ground-truth answer; **(2) Retrieval Quality (RetQ):** the relevance of the provided or retrieved context in comparison to the human-validated relevant context; **(3) Reasoning Quality (ReaQ):** the correctness of generated reasoning chains and intermediate steps that explain how the final answer is derived.

E.2 Applications

We showcase the applications and demonstrations in synthesizing LLMs with KGs for QA.

KAG (Liang et al., 2024a) (by Antgroup)³ is a domain-knowledge augmented generation framework that leverages KGs and vector retrieval to bidirectionally enhance LLMs for knowledge-intensive

³<https://github.com/OpenSPG/KAG>

| Methods | Techniques | LM(s) | KG(s) | Dataset(s) | QA Type | Metric(s) |
|--|---|---|--|---|--------------------------------|---|
| ACT-Selection (Salnikov et al., 2023) | ACT-based Answer Selection and Ranking | T5-Large-SSM | Wikidata | SQ, RuBQ, Mintaka | KGQA, CBQA, Multi-lingual KGQA | Hit@1 |
| Q-KGR (Zhang et al., 2024c) | Question-guided KG Re-scoring, FNN-based Knowledge Injection | FLAN-T5-XL, RoBERTa-Large, Llama-2-7B | ConceptNet | OBQA, PIQA | KGQA | Acc |
| KG-Rank (Yang et al., 2024) | Similarity and MMR based Ranking | GPT-4, Llama-2-7B, Llama-2-13B | UMLS, DBpedia | LiveQA, ExpertQA-Bio, ExpertQA-Med, MedQA | Domain-specific QA | ROUGE-L, BERTScore, Mover-Score, BLEURT |
| KGR (Guan et al., 2024) | Refine-then-Retrieve, Knowledge Truthfulness Verification | GPT-4, Llama-2-7B, Vanilla Llama-2-7B, Transformer | CKG, PrimeKG | MedQuAD | Domain-specific QA | Truthfulness Score |
| EFSUM (Ko et al., 2024) | KG Fact Summarization, KG Helpfulness and Faithfulness Filter | GPT-3.5-Turbo, FLan-T5-XL, Llama-2-7B-Chat | Dataset Inherent KGs (Freebase, Wikidata) | WQSP, Mintaka | KGQA, Multi-hop QA | Acc |
| InteractiveKBQA (Xiong et al., 2024) | Multi-turn Interaction for Observation and Thinking | GPT-4-Turbo, Mistral-7B, Llama-2-13B | Freebase, Wikidata, Movie KG | WQSP, CWQ, KQA Pro, MetaQA | KBQA, Domain-specific QA | F1, Hits@1, EM, Acc |
| LPKG (Wang et al., 2024b) | Planning LLM Tuning, Inference, and Execution | GPT-3.5-Turbo, CodeQwen1.5-7B-Chat, Llama-3-8B-Instruct | Dataset Inherent KGs (Wikidata), Wikidata15K | HotpotQA, 2WikiMQA, Bamboogle, MuSiQue, CLQA-Wiki | KGQA, Multi-hop QA | EM, P, R |

Table 5: Summary and Comparison of Methods in the Category of KG as Refiners and Validators. Acc: Accuracy, EM: Exact Match, BLEURT: Bilingual Evaluation Understudy with Representations from Transformers.

| Methods | Techniques | LM(s) | KG(s) | Dataset(s) | QA Type | Metric(s) |
|---|--|---|---------------------------------|---|-------------------------------------|-----------------------------|
| LongRAG (Zhao et al., 2024a) | Domain-specific Fine-Tuning for RAG & CoT-guided Filter | ChatGLM3-6B, Qwen1.5-7B, Vicuna-v1.5-7B, Llama-3-8B, GPT-3.5-Turbo, GLM-4 | Wikidata | HotpotQA, 2WikiMQA, MusiQue | KBQA, Multi-hop QA | F1 |
| SimGRAG (Cai et al., 2024) | Instruction Fine-Tuning for RAG with Filter | Llama-3-8B-Instruct, Gemma-2-27B-it | Wikipedia, PubMed | PubMedQA, BioASQm, MedQA, MedMCQA, LiveQA, MedicationQA | Domain-specific QA, Multi-choice QA | Acc, Rouge-L, MAUVE, EM, F1 |
| KGQA (Ji et al., 2024) | KG-related Instruction Tuning with CoT Reasoning | Llama-2-7B-Chat, BGE-1.5-en-base | Dataset Inherent KGs | WQAP, CWQ | KGQA | Hits@1, F1 |
| KG-IRAG (Yang et al., 2025) | Incremental Retrieval and Iterative Reasoning | Llama-3-8B-Instruct, GPT-3.5-Turbo, GPT-4o-mini, GPT-4o | Self-constructed KGs | TFNSW | Temporal QA | EM, F1, HR, HAL |
| PIP-KAG (Huang et al., 2025) | Parameteric Pruning for KAG | Llama-3-8B-Instruct | Dataset Inherent KGs | CoConflictQA | KGQA, Multi-hop QA | EM, ConR, MR |
| RAG-KG-IL (Yu and McQuade, 2025) | Agent-based Incremental Learning and Knowledge Dynamic Update | GPT-4o | Self-constructed KGs | Curated Dataset from NHS | Domain-specific QA | AM, HVC |
| CoT-RAG (Li et al., 2025a) | KG-driven CoT Generation and Knowledge-aware RAG with Pseudo-program KGs | ERNIE-Speed-128K, GPT-4o-mini | Self-curated Pseudo-Program KGs | HotpotQA, CSQA, SIQA | KGQA, Multi-hop QA | RA, Robustness |

Table 6: Summary and Comparison of Methods in the Category of Hybrid Approach. EM: Exact Match, HR: Hit Rate, HAL: Hallucination, ConR: Context Recall, MR: Memorization Ratio, AM: Accuracy Matching, HVC: Human-verified Completeness, RA: Reasoning Accuracy.

tasks, such as QA.

PIKE-RAG (Wang et al., 2025) (by Microsoft)⁴ is a specialized knowledge and rationale augmented generation system with a focus on extracting, understanding, and applying domain-specific knowledge to guide LLMs toward accurate responses.

GraphRAG-QA (by NebulaGraph)⁵ is an industrial demo of GraphRAG integrating several query engines for augmenting QA, NLP2Cypher-based KG query engine (Wu et al., 2022), vector RAG query engine, and Graph vector RAG query engine.

MedRAG (Zhao et al., 2025b) (by Nanyang Tech-

nological University et al.)⁶ is a KG-elicited reasoning enhanced RAG-based healthcare copilot that generates diagnoses and treatment recommendations based on the input patient manifestations.

Fact Finder (Steinigen et al., 2024) (by Fraunhofer IAIS and Bayer)⁷ augments LLMs with the query-based retrieval from medical KG to improve the completeness and correctness of answers.

AprèsCoT (Shirdel et al., 2025) (by University of Waterloo) is a visualization tool for explaining the answers generated by LLMs with CoT-based path exploration and inference over KGs.

⁴<https://github.com/microsoft/PIKE-RAG>

⁵<https://github.com/wey-gu/demo-kg-build>

⁶<https://github.com/SNOWTEAM2023/MedRAG>

⁷<https://github.com/chrschy/fact-finder>

| Approach | Multi-doc QA | Multi-modal QA | Multi-hop QA | Conversational QA | XQA | Temporal QA |
|------------------------------|--------------|----------------|--------------|-------------------|-----|-------------|
| KG as Background Knowledge | ✓ | ✓ | ✓ | ✓ | △ | ✓ |
| KG as Reasoning Guidelines | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ |
| KG as Refiners and Validator | △ | ✗ | △ | ✓ | ✗ | △ |
| Hybrid Methods | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 7: Research Progress on Alignment of Existing Approaches of Synthesizing LLMs and KGs with Complex QA. ✓ Fully Investigated. △ Partially Investigated. ✗ Not Yet Investigated.

| Dataset | Category | AnsQ | RetQ | ReaQ | Brief Description |
|-------------------------------------|-------------------|------|------|------|--|
| WebQSP (tau Yih et al., 2016) | KBQA | ✓ | ✓ | △ | Includes SPARQL queries for knowledge-based QA. |
| BioASQ-QA (Krithara et al., 2023) | KBQA | ✓ | ✓ | ✗ | Includes exact and ideal answers of question over multi-doc summarization. |
| CAQA (Hu et al., 2024a) | KBQA | △ | ✓ | △ | Evaluates complex reasoning and attribution. |
| CR-LT KGQA (Guo et al., 2024) | KGQA | ✓ | △ | ✓ | Supports long-tail entities and common-sense reasoning. |
| EXAQT (Jia et al., 2021) | KGQA | ✓ | ✓ | △ | Supports temporal QA with multiple entities, predicates, and conditions. |
| CommonsenseQA (Talmor et al., 2019) | Multi-choice QA | ✓ | △ | ✗ | Includes questions with the correct answer and four distractor answers. |
| MedQA (Jin et al., 2021) | Multi-choice QA | ✓ | △ | △ | Medical multi-choice QA dataset with multilingual medical examination text. |
| MINTQA (He et al., 2024) | Multi-hop QA | ✓ | ✓ | ✓ | Supports LLMs evaluation on new and tail knowledge. |
| FanOutQA (Zhu et al., 2024) | Multi-Hop QA | △ | △ | ✓ | Includes multi-hop question over multi-document. |
| M3SciQA (Li et al., 2024a) | Multi-modal QA | ✓ | ✓ | ✓ | QA over the context across multiple document. |
| OMG-QA (Nan et al., 2024) | Multi-modal QA | ✓ | ✓ | △ | Evaluates the retrieval and reasoning of QA across modalities. |
| ScienceQA (Lu et al., 2022) | Multi-modal QA | ✓ | ✗ | ✓ | Supports multi-choice questions across disciplines in science topics. |
| TempTabQA (Gupta et al., 2023) | Temporal QA | ✓ | △ | ✓ | Supports temporal reasoning over semi-structured tables. |
| PATQA (Meem et al., 2024) | Temporal QA | ✓ | △ | ✓ | Provides Wikidata questions for Present-anchored temporal QA. |
| MenatQA (Wei et al., 2023) | Temporal QA | △ | △ | ✓ | Evaluates the temporal reasoning capability of LLMs. |
| Loong (Wang et al., 2024c) | Long-context QA | ✓ | ✗ | △ | QA over multi-doc with a relevant document for the final answer. |
| SituatedQA (Zhang and Choi, 2021) | Open-retrieval QA | ✓ | ✓ | ✗ | Includes correct question with temporal or geographical context. |
| ChatData (Sequeda et al., 2024) | LLM-KG QA | ✓ | ✓ | ✗ | LLMs+KGs QA over the enterprise SQL database. |
| LLM-KG-Bench (Meyer et al., 2023) | LLM-KG QA | ✓ | ✗ | ✗ | Evaluates capabilities of LLMs in knowledge graph engineering. |
| XplainLLM (Chen et al., 2024d) | LLM-KG QA | ✓ | △ | ✓ | Focuses on QA explainability and reasoning. |
| OKGQA (Sui and Hooi, 2024) | LLM-KG QA | ✓ | ✓ | ✓ | Evaluates LLMs+KGs for open-ended QA. |
| LiHua-World (Fan et al., 2025) | LLM-KG QA | ✓ | ✓ | ✓ | Evaluates the capability of LLMs on multi-hop QA in the scenario of RAG. |
| STaRK (Wu et al., 2024a) | LLM-KG QA | ✓ | ✓ | ✓ | Evaluates the performance of LLMs-driven RAG for QA. |
| CoConflictQA (Huang et al., 2025) | LLM-KG QA | ✓ | △ | ✓ | Evaluates contextual faithfulness for QA in the scenario of KAG. |
| mmRAG (Xu et al., 2025a) | LLM-KG QA | ✓ | ✓ | ✓ | Evaluates multi-modal RAG including QA dataset across text, tables, and KGs. |
| BlendQA (Xin et al., 2025) | LLM-KG QA | ✓ | △ | ✓ | Evaluates cross-knowledge source reasoning capabilities of RAG for QA. |

Table 8: Comparison of LLMs+KGs for QA Benchmark Dataset. ✓ The dataset was primarily designed to support this evaluation. △ The dataset can be adapted for this evaluation, but this evaluation is not their main focus. ✗ The dataset does not support this evaluation.