Think-on-Graph 3.0: Efficient and Adaptive LLM Reasoning on Heterogeneous Graphs via Multi-Agent Dual-Evolving Context Retrieval

Anonymous ACL submission

Abstract

Graph-based Retrieval-Augmented Generation (GraphRAG) has become an important paradigm for enhancing Large Language Models (LLMs) with external knowledge. However, existing approaches are constrained by their reliance on high-quality knowledge graphs: manually built ones are not scalable, while automatically extracted ones are limited by the performance of LLM extractors, especially when using smaller, locally deployed models. To address this, we introduce Think-on-Graph 3.0 (ToG-3), a novel framework featuring a Multi-Agent Context Evolution and Retrieval (MACER) mechanism. Its core contribution is the dynamic construction and iterative refinement of a Chunk–Triplet–Community heterogeneous graph index, powered by a dual-evolution process that adaptively evolves both the query and the retrieved subgraph during reasoning. ToG-3 dynamically builds a targeted graph index tailored to the query, enabling precise evidence retrieval and reasoning even with lightweight LLMs. Extensive experiments demonstrate that ToG-3 outperforms compared baselines on both deep and broad reasoning benchmarks, and ablation studies confirm the efficacy of the MACER framework.

1 Introduction

The rapid advancement of both commercial (OpenAI, 2025; AI, 2025a; Comanici et al., 2025) and open-source Large Language Models (LLMs) (Yang et al., 2025; AI, 2025b; Liu et al., 2024; Zeng et al., 2025; Gan et al., 2023) has significantly enhanced the accessibility of generative AI capabilities for both end-users and developers. Retrieval-augmented generation (RAG) (Gao et al., 2023) has become a popular method for grounding Large Language Models (LLMs) with external knowledge, addressing issues like knowledge cutoff and hallucination. ToG (Sun et al., 2023; Ma et al., 2024) pioneered an iterative hybrid RAG framework that tightly couples text and KGs retrieval, though their approach relies on pre-existing structured KGs such as Freebase and Wikidata. On the other hand, methods like GraphRAG (Edge et al., 2024) and LightRAG (Guo et al., 2024) address this issue by constructing a graph directly from the input documents. They create an entity-based graph to enhance information retrieval and summarization. However, corpus-built graph RAG remains bottlenecked by extraction quality: lightweight, locally deployed LLMs can miss triplets, produce coarse relations, or fail to parse structured outputs, yielding incomplete or noisy graphs. Existing systems also often split local entity-level retrieval from global community-level reasoning. Appendix A provides concrete examples and parsing-failure statistics.

To address these limitations, we introduce **Think-on-Graph 3.0** (ToG-3), a RAG framework that integrates graph-based retrieval with iterative reasoning. Our core contribution is a Chunk–Triplet–Community heterogeneous graph architecture together with MACER (Multi-Agent Context Evolution and Retrieval), a dual-evolution mechanism that jointly evolves the **query** and the retrieved **subgraph** for precise evidence retrieval. Unlike methods that add retrieval layers to a fixed index, ToG-3 treats weak initial graph evidence as repairable at query time: MACER decomposes under-supported queries, expands or re-extracts missing evidence, and prunes irrelevant subgraphs until the context approaches a minimal sufficient evidence subgraph. Figure 1 illustrates the key distinctions between ToG-3 and classical RAG paradigms such as NaiveRAG and GraphRAG. ToG-3 dynamically refines both the query representation and the graph structure, while prior RAG methods typically construct a static graph index in a single pass without adapting to the actual query.

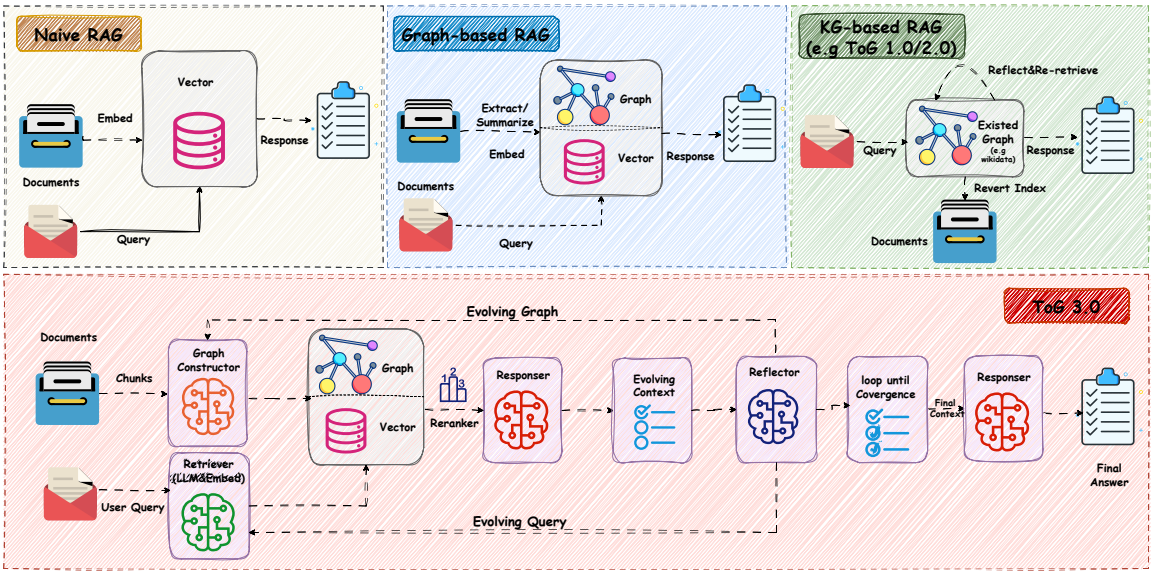

Figure 1: Evolution of Retrieval-Augmented Generation Paradigms. (**a**) Naive RAG embeds raw documents and performs single-shot retrieval. (**b**) Graph-based RAG pre-builds a static graph once and retrieves from it. (**c**) ToG-3 introduces a *five-agent* loop, including Retriever, Constructor, Reflector, Reranker, and Response agents, where the graph and query sub-tasks *co-evolve* at runtime, yielding dynamic, query-adaptive context that converges to a minimal sufficient subgraph.

Our key contributions are summarized as follows:

1. We propose **MACER** (Multi-Agent Context Evolution and Retrieval), a multi-agent framework that introduces a dual-evolution mechanism integrating evolving queries and evolving subgraphs within graph-based RAG. This design explicitly repairs weak initial graph construction and improves retrieval for complex reasoning, especially when lightweight open-source LLMs serve as the backbone.

2. We present **ToG-3**, a unified reasoning system that combines chunk-, triplet-, and community-level evidence in a single heterogeneous graph index, enabling both deep multi-hop inference and broad community-level reasoning.

3. We conduct extensive experiments on both **Deep and Broad Reasoning Tasks**, together with ablations, efficiency analysis, additional baseline comparisons, case studies, and theoretical analysis, demonstrating the effectiveness and practical trade-offs of the proposed framework.

2 Related Work

2.1 Graph-Based Retrieval-Augmented Generation

Recent advances in retrieval-augmented generation (RAG) have increasingly emphasized structural awareness to improve reasoning depth and contextual coherence. Edge et al. (2024) propose GraphRAG, which builds a knowledge graph (KG) from documents via LLM-based entity and relation extraction, then applies community detection to generate hierarchical summaries for global sensemaking. Guo et al. (2024) introduce LightRAG, which employs a dual-level retrieval system combining low-level fact retrieval and high-level semantic discovery using a compact KG, improving both efficiency and coverage. Further building on this idea, Gutiérrez et al. (2024, 2025) present a non-parametric continual learning framework that uses Personalized PageRank over an open KG to enable associative, multi-hop reasoning. Other structure-augmented RAG methods include RAPTOR (Sarthi et al., 2024), which recursively organizes retrieved chunks through abstractive summaries, and memory-augmented approaches (Chen et al., 2023) that enhance sensemaking through associative context. These methods improve broad context integration, but they do not explicitly repair missing entity-relation evidence during query-time multi-hop reasoning. Note that RL-based frameworks such as

GraphRAG-R1 (Yu et al., 2025) and Graph-R1 (Luo et al., 2025a) utilize existing Graph-based RAG methods as their retrieval components and train an end-to-end agentic framework. Our work, in contrast, proposes a complementary approach to improve the underlying retrieval paradigm itself. Consequently, ToG-3 could also serve as a plug-in component to enhance such RL frameworks.

Positioning of ToG-3. ToG-3 differs from prior multi-layer or graph-enhanced RAG systems along three dimensions. First, its Chunk–Triplet–Community schema supports both fine-grained entity-relation traversal and coarse-grained community reasoning within one retrieval substrate, whereas many prior systems such as RAPTOR, E^{2}GraphRAG, KET-RAG, and KGP (Sarthi et al., 2024; Zhao et al., 2025; Huang et al., 2025; Wang et al., 2024) are optimized mainly for either multi-hop factual QA or global sense-making. Second, the MACER loop makes graph retrieval adaptive: when the current evidence is insufficient, it evolves sub-queries and repairs the retrieved subgraph through expansion, re-extraction, and pruning, instead of relying only on a fixed graph built during indexing. Third, the online refinement process is formalized as an MDP, providing the basis for the convergence analysis in Appendix L. This positioning makes ToG-3 especially relevant for private or offline deployments where the graph extractor is a lightweight LLM and the initial graph is expected to be imperfect.

2.2 Knowledge Graphs in RAG and Hybrid Approaches

The integration of structured knowledge into LLM reasoning has long been pursued to improve faithfulness and interpretability. Early KG-augmented RAG systems retrieve triples from static external knowledge bases such as Wikidata or Freebase to ground model outputs (Sun et al., 2023). However, these sources are often incomplete, outdated, or misaligned with domain-specific content. To overcome this, hybrid RAG frameworks (Ma et al., 2024) combine unstructured text and structured KGs to balance breadth and precision. Chain-of-Knowledge (CoK) (Li et al., 2024) retrieves from multiple structured sources including Wikipedia, Wikidata, and Wikitable to ground LLM responses. HybridRAG (Sarmah et al., 2024) fuses vector-based and KG-based retrievers, demonstrating superior reasoning performance compared to either modality alone.

2.3 Iterative and Reflective Reasoning in LLMs

Enabling LLMs to reason iteratively has been shown to improve accuracy and faithfulness. ITER-RETGEN (Shao et al., 2023) introduces an iterative loop that alternates between retrieval and generation, using generated hypotheses to guide further search. IRCoT (Trivedi et al., 2023) combines Chain-of-Thought (CoT) with retrieval, interleaving reasoning steps with evidence gathering and significantly improving performance on multi-hop QA. Self-RAG (Asai et al., 2023) equips LLMs with reflection tokens to decide when to retrieve and whether the output is hallucinated. ReAct (Yao et al., 2023a) combines reasoning traces with external actions, enabling task decomposition and environment interaction. Other efforts focus on continual learning for LLMs, where RAG serves as a non-parametric alternative to fine-tuning (Shi et al., 2024). Continual pretraining (Jin et al., 2022) and instruction tuning (Zhang et al., 2023) can update model parameters but suffer from catastrophic forgetting (Huang et al., 2024). Model editing methods (Yao et al., 2023b) offer fine-grained updates but struggle with generalization.

3 Methodology

Think-on-Graph 3.0 (ToG-3) introduces a novel *Multi-Agent Context Evolution and Retrieval (MACER)* framework for open-domain question answering.

3.1 Problem Formulation

Let $\mathcal{D} = \{d_i\}_{i=1}^N$ be a text corpus. The objective is to answer a user query q with an answer a^* that is both accurate and *faithful* to the source corpus, derived from a *minimal, sufficient subgraph* $\mathcal{G}_q^*$ of a heterogeneous graph $\mathcal{G}$ constructed from $\mathcal{D}$:

$$\mathcal{G}_q^* = \underset{\mathcal{G}' \subseteq \mathcal{G}}{\operatorname{argmin}} |\mathcal{G}'| \quad \text{subject to} \quad \texttt{Suff}(q, \mathcal{G}') = 1, \tag{1}$$

where $\texttt{Suff}(\cdot,\cdot) \in \{0,1\}$ is a function judging the sufficiency of a subgraph for answering the query.

Existing methods face a critical dilemma: **(1)** Systems like ToG-1 or 2 rely on high-quality, pre-constructed KGs, limiting their applicability

to private or specialized domains. **(2)** Corpus-based GraphRAG methods (e.g., GraphRAG, LightRAG) build a static graph from $\mathcal{D}$ in one go. Their performance is bottlenecked by the quality of this initial graph, which in turn depends heavily on the capability of the LLM used for information extraction.

3.2 Heterogeneous Graph Index: Schema and Construction

3.2.1 Node and Edge Schema

The Constructor Agent builds a heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with three node types:

- **Chunks** ($\mathcal{C}$): Sentence-level text passages from the corpus.
- **Triplets** ($\mathcal{T}$): Semantic triples (s, p, o) extracted from chunks, annotated with entity and relation types ($\texttt{type}_s$, $\texttt{type}_p$, $\texttt{type}_o$).
- **Communities** ($\mathcal{M}$): Summaries of entity clusters obtained via Leiden clustering on the entity co-occurrence graph, each condensed into an abstract.

Edges are defined by three type relations:

- **OpenRel**(s, p, o): Connects entities s and o via predicate p extracted by the LLM, forming an open-domain semantic triple.
- **MentionedIn**(t, c): Connects a triplet t to the chunk c from which it was extracted.
- **SummaryFor**(m, e): Connects a community summary node m to an entity e that belongs to that community.

This unified schema allows both fine-grained (chunk/triplet) and high-level (community) information to be retrieved seamlessly within a single vector space, effectively addressing the local/global retrieval dichotomy of prior GraphRAG systems.

3.2.2 Offline Index Construction

Algorithm 1 in Appendix. C details the one-time construction of the universal index $\mathcal{G}$. A key design choice is the use of a single frozen encoder E_θ (e.g., jina-embeddings-v3 (Sturua et al., 2024)) to embed all nodes, regardless of type, into a unified 1024-dimensional dense vector space. This enables efficient vector search across all node types during retrieval.

3.3 The MACER Process: Multi-Agent Context Evolution and Retrieval

The core of ToG-3 is the online MACER loop (Algorithm 2), an iterative process of retrieval, generation, and reflection that dynamically evolves the context subgraph $\mathcal{G}_k$. We formalize this process as an episodic Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r)$.

State Space ($\mathcal{S}$) : At each step k, the state $s_k = (q, \mathcal{G}_k, \mathcal{H}_k)$ captures the complete reasoning context, including the original query q, the current evidence subgraph $\mathcal{G}_k$ retrieved by Retriever Agent π_{ret} and reranked by Reranker Agent π_{rer}, and the trajectory history $\mathcal{H}_k = \{(q'_i, a_i, r_i, \mathcal{G}_i)\}_{i=0}^{k-1}$ of all previous sub-queries, answers, rewards, and subgraphs.

Action Space ($\mathcal{A}$) : The Reflector Agent π_{ref} serves as the policy network. Its action a_k at state s_k is either to generate a targeted refinement sub-query q'_k (to continue the reasoning process) or to output the STOP action (to terminate the episode).

Reward Function (r) : Upon the Response Agent generating an answer a_k, the Reflector immediately provides a sparse, binary reward r_k:

$$r_k = \begin{cases} 1 & \text{if } \texttt{Suff}(q, \mathcal{G}_k, a_k) = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

This reward signal is produced by the Reflector Agent to determine if the current context evidence is sufficient to answer the user's query.

Transition Dynamics (P) Given the current state s_k and an action a_k (which corresponds to issuing a sub-query q'_k), the transition to the next state s_{k+1} occurs deterministically according to the following update rules: The constructor agent π_{const} applies the transition operator using the generated sub-query q'_k and the current graph state $\mathcal{G}_k$ to produce an updated graph $\mathcal{G}_{k+1}$. This step includes an iterative sequence of *evolving queries* and *evolving subgraphs*, reflecting the structural evolution of the graph based on the agent's reasoning action and formally defined by the recurrence:

$$q'_k = \pi_{\text{ref}}^{\text{evolve}}(q, \mathcal{G}_k), \tag{3}$$

$$\mathcal{G}_{k+1} = \pi_{\text{const}}^{\text{evolve}}(q'_k, \mathcal{G}_k), \tag{4}$$

The action history $\mathcal{H}_{k+1}$ is augmented with a new tuple recording the executed sub-query q'_k, the corresponding action a_k, the reward r_k received,

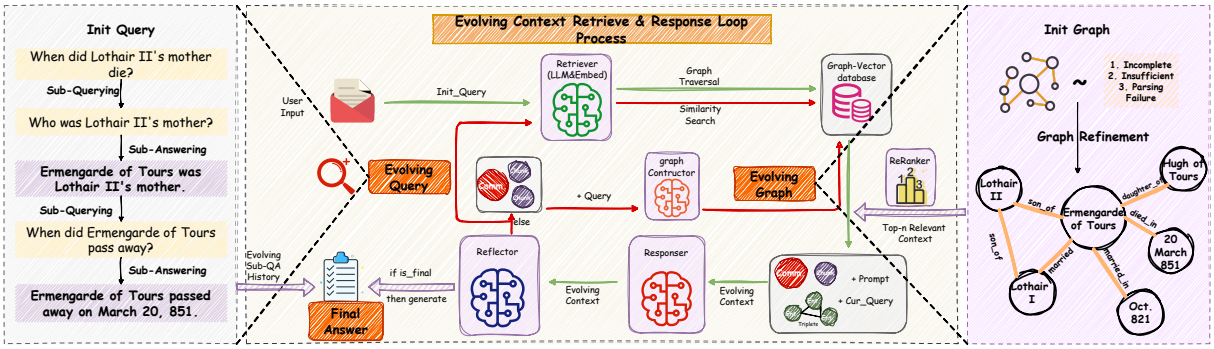

Figure 2: Multi-Agent Dual-Evolving Context Retrieval-Response Loop. The Retriever fetches an initial chunk–triplet–community subgraph, and the Reranker selects the top-n most relevant pieces of evidence. The Response Agent produces an answer; the Reflector Agent judges sufficiency (reward=1/0). If insufficient (reward=0), the Reflector evolves the query into sub-queries while the Constructor evolves the subgraph through refinement. The loop repeats until the context becomes sufficient or the horizon is reached, after which the Response Agent synthesizes the final answer from the full trajectory.

and the resulting graph state $\mathcal{G}_{k+1}$. This ensures a comprehensive trace of the reasoning trajectory, which is essential for credit assignment and subsequent learning.

$$\mathcal{H}_{k+1} = \mathcal{H}_k \cup \{(q'_k, a_k, r_k, \mathcal{G}_{k+1})\} \tag{5}$$

$$a^* \leftarrow \pi_{\text{resp}}^{\text{final}}(q, \mathcal{H}_k) \tag{6}$$

The complete MACER process, now cast as an MDP, is summarized in Algorithm 2. The loop continues until the Reflector's policy π_{ref} outputs the STOP action (via $r_k = 1$) or a maximum horizon K is reached. The final answer a^* is synthesized from the full trajectory $\mathcal{H}_k$ of states and actions, ensuring faithfulness to the evolved evidence. This MDP formulation provides the formal foundation for establishing the convergence of the MACER process under mild assumptions, as detailed in Appendix. L. Note that the offline index $\mathcal{G}$ constructed in Section 3.2.2 remains static and intentionally lightweight. What evolves online is the retrieved evidence subgraph $\mathcal{G}_k$. This evolution is not a simple union of retrieved results across turns: the Constructor Agent expands missing evidence, optionally re-extracts or normalizes triples from newly retrieved chunks, and prunes irrelevant nodes according to the evolved sub-query and Reflector feedback. This iterative refinement allows ToG-3 to start from a potentially weak initial graph but *specialize* it towards the reasoning path of the specific query, converging on a high-quality evidence subgraph $\mathcal{G}_q^*$. This evolution and refinement mechanism alleviates three fundamental weaknesses of small LMs in static GraphRAG, including incomplete triplet recall, insufficient knowledge details, and parsing failures in LLM outputs, as mentioned in Section 1.

4 Experiment

4.1 Experimental Setup

Datasets To comprehensively evaluate the reasoning capabilities of RAG systems, we conduct experiments on two distinct categories of tasks: **Deep Reasoning Tasks** including HotpotQA (Yang et al., 2018), 2WikiMultiHopQA (Ho et al., 2020) and Musique (Trivedi et al., 2022) and **Broad Reasoning Tasks** including 4 subsets of UltraDomain (Qian et al., 2025) benchmark. Detailed statistics for all datasets are provided in Table 3 and Appendix. D.

Evaluation Metrics For **Deep Reasoning Tasks**, we follow standard QA evaluation practices with **Exact Match (EM)** (following ToG and ToG-2 (Sun et al., 2023; Ma et al., 2024), we employ a substring-based Exact Match metric) and **F1 Score**. For **Broad Reasoning Tasks**, we adopt a multi-dimensional LLM-based evaluation approach including **Comprehensiveness**, **Diversity**, and **Empowerment** following (Guo et al., 2024). Metric details are provided in Appendix F.

Baselines We compare ToG-3 against the following state-of-the-art RAG methods across all datasets, including NaiveRAG (Gao et al., 2023), ToG-2 (Ma et al., 2024), GraphRAG (Edge et al., 2024), LightRAG (Guo et al., 2024), MiniRAG (Fan et al., 2025) and HippoRAG-2 (Gutiérrez et al., 2025). Baselines details can be found in Appendix.E. For graph-based methods, we maintain identical chunk sizes (1024 tokens) and use the same LLM (Qwen2.5-32B-

Instruct (Yang et al., 2024)) for all extraction and generation tasks to eliminate model capability variations. Implementation details are provided in Appendix B.

4.2 Result of Deep Reasoning Benchmark

Result Analysis from a Method Perspective. Results shown in Table 1 represent the average of three independent reasoning experiments. Previous graph-based methods such as GraphRAG, which rely on LLM-based graph construction, show limited performance. Their performance is the lowest, particularly in terms of F1 scores as shown in Figure 3b, which can be attributed to a lack of focus on deep factual reasoning and a tendency to produce verbose responses, resulting in low token-level recall. More detailed precision and recall results are provided in Appendix. G.1. ToG-2, without leveraging well-curated knowledge graphs like Freebase and Wikidata, demonstrates moderate performance in open-domain settings. NaiveRAG achieves competitive third-place results by avoiding graph construction limitations and relying solely on retrieved documents for response generation. HippoRAG-2 emerges as the strongest baseline, employing an efficient embedding model with Personalized PageRank algorithm and LLM-based triple filtering to achieve second-best performance. However, our proposed method consistently outperforms all competitors, achieving the highest average EM (0.474) and F1 (0.345) scores across all three benchmarks. This superior performance is attributed to our novel Chunk-Triplets-Community heterogeneous graph architecture and the Multi-Agent Context Evolution and Retrieval (MACER) framework, which enables adaptive subgraph refinement and evolving query decomposition for complex reasoning tasks and overcomes the graph construction challenges that plague other graph-based RAG systems. Additional baselines are provided in Appendix I, including graph-structured KGP and reflective non-graph IRCoT under identical settings. KGP reaches an average EM/F1 of 0.275/0.125, and IRCoT reaches 0.435/0.301, showing that iterative reflection helps but does not fully close the gap to ToG-3.

Result Analysis from a Dataset Perspective. As shown in Figure 3, the average performance of the baselines and our method across the HotpotQA, 2WikiMultiHopQA, and Musique datasets generally follows a descending trend. This pattern can be attributed to the following reasons: HotpotQA (Yang et al., 2018): Although widely used, this dataset has been shown to provide a weaker test of multi-hop reasoning due to the presence of numerous spurious cues and shortcut signals (Trivedi et al., 2022; Gutiérrez et al., 2024). Musique (Trivedi et al., 2022): A challenging multi-hop QA dataset comprising approximately requiring 2–4 hops, which emphasizes a comprehensive evaluation of multi-step reasoning abilities. Musique is designed to feature diverse and complex reasoning paths, necessitating the integration of information across multiple hops to arrive at correct answers.

4.3 Result of Broad Reasoning Tasks

As shown in Figure 4, the four heatmaps demonstrate that the five methods can be divided into two clusters: the upper-right region (predominantly red, indicating superior performance) and the lower-left region (predominantly blue, indicating inferior performance). Specifically, ToG-3, GraphRAG, and LightRAG exhibit significantly higher win rates compared to NaiveRAG and HippoRAG-2. Detailed win rates (%) of baselines versus ToG-3 across four datasets are provided in Table 5 of Appendix G. Our framework outperforms NaiveRAG by substantial margins (up 75.0% average win rate on all four datasets), highlighting the limitations of chunk-based retrieval for complex queries. While GraphRAG shows competitive performance in comprehensiveness due to its extensive community summarization and retrival, ToG-3 achieves better balance across all metrics, particularly excelling in diversity and empowerment through its heterogeneous graph architecture that integrates chunk-level, triplet-level, and community-level information. Detailed ELO rating calculation for broad reasoning tasks can be found in Appendix. G.3. The multi-agent dual-evolving context retrieval mechanism enables both deep knowledge reasoning through entity-relation exploration and broad community reasoning. Our analysis reveals that, on average, 20% of the samples require one evolving-context iteration, 32% require two iterations, and 48% require three iterations.

Detailed comparisons of time and token consumption across different methods are provided in Appendix H. Case studies of ToG-3 retrieval and

Table 1: Exact Match (EM) and F1 scores on Deep Reasoning datasets.We highlight the best, second-best, and third-best methods with different background color shades.

Method	HotpotQA		2WikiMultihopQA		Musique		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
NaiveRAG	0.634	0.365	0.382	0.189	0.230	0.143	0.415	0.232
ToG-2	0.308	0.153	0.401	0.194	0.103	0.105	0.271	0.151
GraphRAG	0.337	0.011	0.439	0.018	0.109	0.008	0.295	0.012
LightRAG	0.308	0.013	0.420	0.023	0.082	0.009	0.270	0.015
MiniRAG	0.213	0.012	0.125	0.018	0.067	0.007	0.135	0.012
HippoRAG-2	0.612	0.534	0.491	0.254	0.212	0.145	0.438	0.311
Ours	0.654	0.569	0.527	0.291	0.241	0.174	0.474 ↑8.2%	0.345 ↑10.9%

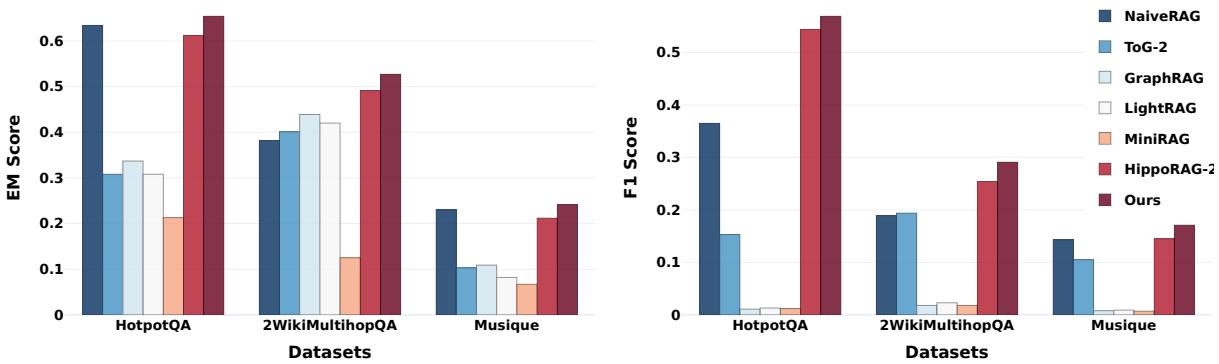

(a) Exact Match (EM) Score Comparison **(b)** F1 Score Comparison

Figure 3: Performance comparison of different RAG methods on multi-hop QA datasets. (a) Exact Match scores measure the percentage of questions where the model's answer matches the ground truth. (b) F1 scores provide a harmonic mean of precision and recall for token-level answer matching.

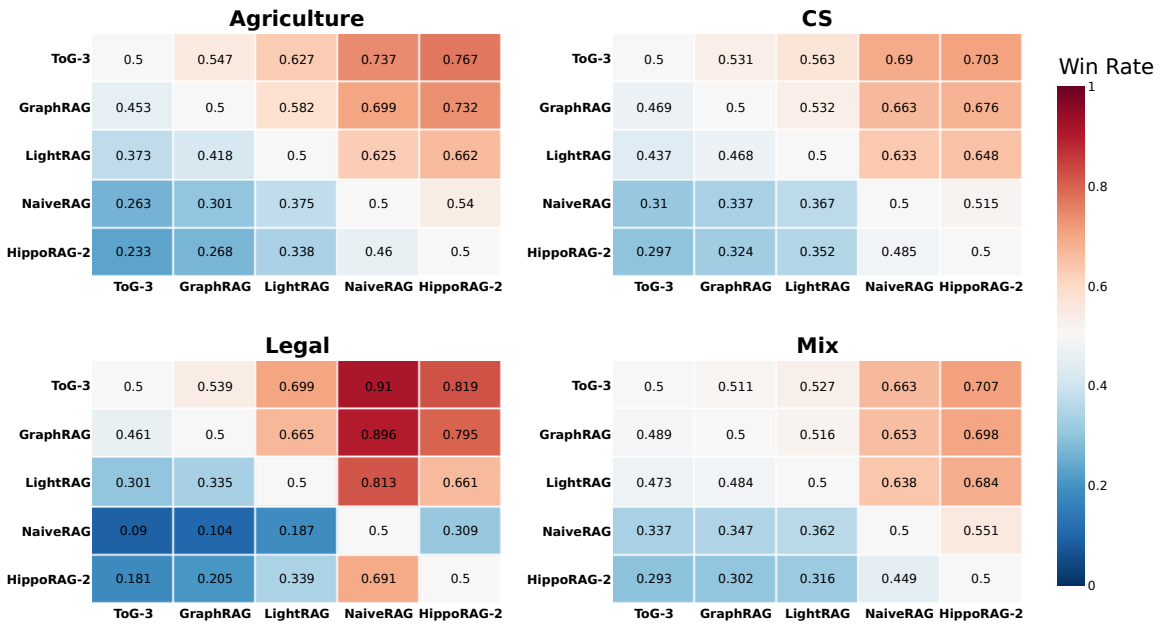

Figure 4: ELO-based Pairwise Win Rate Matrices Across Four Benchmark Datasets. Each heatmap visualizes win probabilities derived from direct head-to-head experimental comparisons, transformed through the ELO framework to ensure transitive consistency. The diagonal of the heatmap is set to a default value of 0.5, indicating self-comparison of the method.

Table 2: Ablation studies of MACER components and foundation model scaling. Standard ToG-3 settings incorporates all MACER components, employs the Qwen2.5-32B-instruct as the backbone LLM, and utilizes the Jina-v3-embedding model for representation encoding and Jina-reranker-v2 for reranking the retrieved evidence.

Ablation Setting	HotpotQA		2WikiMultihopQA		Musique		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
MACER Components Ablation								
w/o Evolving Query	0.614	0.495	0.440	0.227	0.198	0.141	0.417	0.288
w/o Evolving Subgraph	0.629	0.525	0.486	0.258	0.223	0.158	0.446	0.314
w/o Community Node	0.656	0.572	0.514	0.283	0.236	0.169	0.469	0.341
Foundation Model Scaling Ablation								
LLM Model								
Qwen2.5-14B	0.587	0.521	0.480	0.255	0.218	0.154	0.428	0.310
Qwen2.5-72B	0.683	0.592	0.550	0.305	0.255	0.182	0.496	0.360
Embedding Model								
Qwen3-Embed-0.6B	0.653	0.571	0.532	0.294	0.244	0.176	0.476	0.347
Qwen3-Embed-4B	0.658	0.577	0.535	0.296	0.247	0.179	0.480	0.351

response output are provided in Appendix. J.

4.4 Ablation Study

Ablation Study of MACER Components Our ablation study reveals the relative importance of each MACER component for deep reasoning performance. The most significant performance degradation occurs when removing the evolving query mechanism (average performance drop of 12.0% in EM and 16.5% in F1), underscoring its critical role in complex question answering, especially when using smaller LLMs. Removing subgraph refinement causes a moderate performance decrease (average drop of 6.0% in EM and 9.0% in F1), indicating its importance in adapting the knowledge structure to the specific reasoning context. Interestingly, community nodes show the smallest impact on deep reasoning tasks (a slight drop in the average EM and F1 scores), suggesting that while they contribute to performance, the chunk and triplet representations carry most of the relevant information for precise answer generation. However, in broad reasoning tasks, community nodes are essential for comprehensive coverage and diversity, highlighting the complementary roles of different node types in our heterogeneous graph architecture. Note that the reranker agent also delivers a 4.6% improvement in EM and a 10.6% improvement in F1. This is because, during multi-turn RAG processes, an excessive amount of retrieved evidence can otherwise impair the response quality of the Response Agent.

Ablation Study of Foundation Models The foundation model scaling analysis reveals several important patterns. First, LLM capacity has a substantially greater impact on performance than embedding model size. Scaling from Qwen2.5-14B to Qwen2.5-72B yields a 15.9% average improvement in EM scores, highlighting the critical role of reasoning capability in complex QA tasks. Second, larger embedding models provide consistent but more modest improvements. Qwen3-Embed-0.6B shows a slight average EM improvement over jina-embeddings-v3, while Qwen3-Embed-4B provides a 1.7% improvement. This suggests that while retrieval quality matters and larger embedding models contribute to better performance, the LLM's reasoning capacity remains the primary bottleneck for complex reasoning tasks. These findings provide practical guidance for resource allocation in real-world deployments.

5 Conclusion

In this work, we introduced Think-on-Graph 3.0, a novel framework that fundamentally rethinks the paradigm of RAG for complex reasoning. By proposing the Multi-Agent Context Evolution Retrieval (MACER) mechanism and a dynamic Chunk-Triplets-Community heterogeneous graph architecture, we address critical limitations in both existing graph-based RAG methods and knowledge-graph approaches. Our comprehensive experimental evaluation demonstrates that ToG-3 achieves state-of-the-art performance across multiple challenging benchmarks. This adaptive capability proves particularly valuable for overcoming the quality constraints of static graph construction and the domain limitations of pre-existing knowledge bases. The framework's ability to work with light LLMs also opens possibilities for more efficient and deployable AI systems.

Limitations

Our work has several limitations and clear applicability boundaries. First, constrained by GPU resources, our experiments are primarily conducted with LLMs up to 72B parameters and embedding models up to 4B parameters, though these sizes are practical for many local deployment scenarios. Second, the evolving query and subgraph refinement components increase inference latency, typically 2–3$\times$ slower than baseline methods. ToG-3 is therefore better suited to accuracy-critical and knowledge-sensitive applications, rather than latency-critical single-hop retrieval. Third, the same mechanisms result in longer context inputs, which demand larger GPU memory capacity for efficient processing. Fourth, MACER relies on an LLM-based sufficiency decision to stop iteration. Although the maximum horizon and trajectory history reduce unbounded looping, premature stopping, over-iteration, and query drift remain possible failure modes that deserve more systematic evaluation against fixed-hop or marginal-gain stopping rules. Fifth, ToG-3 starts from an automatically extracted graph. The dual-evolving loop can repair missing or under-specified evidence in many cases, but systematic extractor errors or entities completely absent from the retrievable corpus may not be recoverable. Sixth, our evaluation focuses on English text corpora and knowledge-intensive QA settings; extending the framework to multilingual, multimodal, or rapidly changing corpora remains future work. These limitations could be mitigated through model distillation, optimized graph traversal algorithms, dynamic context pruning, controlled robustness training, and more efficient stopping policies.

Ethical Considerations

This research focuses on improving the technical performance of knowledge-enhanced language models. This work utilizes only public benchmark datasets and adheres to strict reproducibility standards. While our framework improves text generation capability, we acknowledge potential risks of generating misleading content and note that performance may reflect biases inherent in base models. We follow the ACL ethical guidelines when conducting the research in this paper.

Information About Use Of AI Assistants

In the preparation of this work, the author used AI-assisted technology (specifically, large language models such as GPT-5 and Deepseek-V4) exclusively for text refinement purposes. The AI was employed to assist in proofreading, correcting grammatical errors, and polishing linguistic expressions to improve the clarity and readability of the manuscript.

References

Anthropic AI. 2025a. Introducing claude 4.

Meta AI. 2025b. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection.

Howard Chen, Ramakanth Pasunuru, Jason Weston, and Asli Celikyilmaz. 2023. Walking down the memory maze: Beyond context limit through interactive reading.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Benoit Dherin, Michael Munn, Hanna Mazzawi, Michael Wunder, and Javier Gonzalvo. 2025. Learning without training: The implicit dynamics of in-context learning. *arXiv preprint arXiv:2507.16003*.

Junnan Dong, Siyu An, Yifei Yu, Qian-Wen Zhang, Linhao Luo, Xiao Huang, Yunsheng Wu, Di Yin, and Xing Sun. 2025. Youtu-graphrag: Vertically unified agents for graph retrieval-augmented complex reasoning. *arXiv preprint arXiv:2508.19855*.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.

Tianyu Fan, Jingyuan Wang, Xubin Ren, and Chao Huang. 2025. Minirag: Towards extremely simple retrieval-augmented generation. *arXiv preprint arXiv:2501.06713*.

Ruyi Gan, Ziwei Wu, Renliang Sun, Junyu Lu, Xiaojun Wu, Dixiang Zhang, Kunhao Pan, Ping Yang, Qi Yang, Jiaxing Zhang, et al. 2023. Ziya2: Data-centric learning is all llms need. *arXiv preprint arXiv:2311.03301*.

Junqi Gao, Xiang Zou, YIng Ai, Dong Li, Yichen Niu, Biqing Qi, and Jianxing Liu. 2025. Graph counselor: Adaptive graph exploration via multi-agent synergy to enhance llm reasoning. *arXiv preprint arXiv:2506.03939*.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yixin Dai, Jiawei Sun, Haofen Wang, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2(1).

Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2024. Lightrag: Simple and fast retrieval-augmented generation.

Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. Hipporag: Neurobiologically inspired long-term memory for large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. 2025. From rag to memory: Non-parametric continual learning for large language models.

Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing A multi-hop QA dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 6609–6625. International Committee on Computational Linguistics.

Jianheng Huang, Leyang Cui, Ante Wang, Chengyi Yang, Xinting Liao, Linfeng Song, Junfeng Yao, and Jinsong Su. 2024. Mitigating catastrophic forgetting in large language models with self-synthesized rehearsal. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1416–1428, Bangkok, Thailand. Association for Computational Linguistics.

Yiqian Huang, Shiqi Zhang, and Xiaokui Xiao. 2025. Ket-rag: A cost-efficient multi-granular indexing framework for graph-rag. *arXiv preprint arXiv:2502.09304*.

Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, and Xiang Ren. 2022. Lifelong pretraining: Continually adapting language models to emerging corpora. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4764–4780, Seattle, United States. Association for Computational Linguistics.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing. 2024. Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources. In *International Conference on Learning Representations*.

Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.

Haoran Luo, Guanting Chen, Qika Lin, Yikai Guo, Fangzhi Xu, Zemin Kuang, Meina Song, Xiaobao Wu, Yifan Zhu, Luu Anh Tuan, et al. 2025a. Graph-r1: Towards agentic graphrag framework via end-to-end reinforcement learning. *arXiv preprint arXiv:2507.21892*.

Haoran Luo, Guanting Chen, Yandan Zheng, Xiaobao Wu, Yikai Guo, Qika Lin, Yu Feng, Zemin Kuang, Meina Song, Yifan Zhu, et al. 2025b. Hypergraphrag: Retrieval-augmented generation via hypergraph-structured knowledge representation. *arXiv preprint arXiv:2503.21322*.

Shengjie Ma, Chengjin Xu, Xuhui Jiang, Muzhi Li, Huaren Qu, Cehao Yang, Jiaxin Mao, and Jian Guo. 2024. Think-on-graph 2.0: Deep and faithful large language model reasoning with knowledge-guided retrieval augmented generation.

OpenAI. 2025. Introducing gpt-5.

Hongjin Qian, Zheng Liu, Peitian Zhang, Kelong Mao, Defu Lian, Zhicheng Dou, and Tiejun Huang. 2025. Memorag: Boosting long context processing with global memory-enhanced retrieval augmentation. In *Proceedings of the ACM Web Conference 2025 (TheWebConf 2025)*, Sydney, Australia. ACM. ArXiv:2409.05591.

Bhaskarjit Sarmah, Benika Hall, Rohan Rao, Sunil Patel, Stefano Pasquali, and Dhagash Mehta. 2024. Hybridrag: Integrating knowledge graphs and vector retrieval augmented generation for efficient information extraction. *arXiv preprint arXiv:2408.04948*.

Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. 2024. RAPTOR: recursive abstractive processing for tree-organized retrieval. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy.

Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. 2024. Continual learning of large language models: A comprehensive survey. *arXiv preprint arXiv:2404.16789*.

Saba Sturua, Isabelle Mohr, Mohammad Kalim Akram, Michael Günther, Bo Wang, Markus Krimmel, Feng Wang, Georgios Mastrapas, Andreas Koukounas, Andreas Koukounas, Nan Wang, and Han Xiao. 2024. jina-embeddings-v3: Multilingual embeddings with task lora.

Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M Ni, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph. *arXiv preprint arXiv:2307.07697*.

Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. 2019. From louvain to leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):1–12.

Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. MuSiQue: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554.

Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions.

Feng Wang, Yuqing Li, and Han Xiao. 2025. jina-reranker-v3: Last but not late interaction for listwise document reranking. *arXiv preprint arXiv:2509.25085*.

Yu Wang, Nedim Lipka, Ryan A. Rossi, Alexa Siu, Ruiyi Zhang, and Tyler Derr. 2024. Knowledge graph prompting for multi-document question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19206–19214.

Derong Xu, Pengyue Jia, Xiaopeng Li, Yingyi Zhang, Maolin Wang, Qidong Liu, Xiangyu Zhao, Yichao Wang, Huifeng Guo, Ruiming Tang, et al. 2025. Align-grag: Reasoning-guided dual alignment for graph retrieval-augmented generation. *arXiv preprint arXiv:2505.16237*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023a. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.

Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023b. Editing large language models: Problems, methods, and opportunities. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10222–10240, Singapore. Association for Computational Linguistics.

Chuanyue Yu, Kuo Zhao, Yuhan Li, Heng Chang, Mingjian Feng, Xiangzhe Jiang, Yufei Sun, Jia Li, Yuzhi Zhang, Jianxin Li, et al. 2025. Graphrag-r1: Graph retrieval-augmented generation with process-constrained reinforcement learning. *arXiv preprint arXiv:2507.23581*.

Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, et al. 2025. Glm-4.5: Agentic, reasoning, and coding (arc) foundation models. *arXiv preprint arXiv:2508.06471*.

Zihan Zhang, Meng Fang, Ling Chen, and Mohammad-Reza Namazi-Rad. 2023. CITB: A benchmark for continual instruction tuning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9443–9455, Singapore. Association for Computational Linguistics.

Yibo Zhao, Jiapeng Zhu, Ye Guo, Kangkang He, and Xiang Li. 2025. Eˆ 2graphrag: Streamlining graph-based rag for high efficiency and effectiveness. *arXiv preprint arXiv:2505.24226*.

Appendices

Within this supplementary material, we elaborate on the following aspects:

- Appendix A: Preliminaries
- Appendix B: Implementation Details and Hyperparameters
- Appendix C: Detailed ToG-3 Algorithms
- Appendix D: Datasets Statistics and Details
- Appendix E: Baselines Details
- Appendix F: Evaluation Metrics
- Appendix G: More Experiment Results and Details
- Appendix H: Comparison of Time and Token Consumption
- Appendix I: Additional Baselines
- Appendix J: Case Study for ToG-3
- Appendix K: Graph Visualization Examples
- Appendix L: Theoretical Support for ToG-3
- Appendix M: LLM Prompts

A Preliminaries

In practical private or on-premises deployments, graph construction is often performed by lightweight open-source LLMs rather than frontier proprietary models. This setting changes the failure profile of graph-based RAG: the initial graph may omit key triplets, collapse fine-grained relations into coarse edges, or return malformed structured outputs. Since static GraphRAG pipelines reuse this graph for downstream retrieval, early extraction errors can prevent the retriever from exposing the evidence required for multi-hop reasoning even when the source corpus contains the answer. Figure 5 illustrates this bottleneck and motivates MACER's query-time repair strategy, which revisits insufficient evidence by evolving both the query decomposition and the retrieved subgraph.

The example in Figure 5 highlights three common evidence failures. First, extraction can be incomplete, such as missing the triplet `Ermengarde of Tours -- died_in -- 20 March 852`. Second, the extracted relation can be under-specified: an output fragment such as `he -- son_of` should instead resolve to `Lothair II -- son_of -- Lothair I`. Third, structured extraction itself can fail at the parsing stage; in our observed JSON outputs, the parsing failure rates are 26.43% on HotpotQA, 23.67% on 2WikiMultiHopQA, and 22.51% on Musique. These failures directly reduce the usable evidence available to a static graph retriever.

B Implementation Details

We implement ToG-3 experiments with the following configuration: **Data Processing**: Chunk size is set to 1024 tokens with 20-token overlap between consecutive chunks to maintain contextual continuity. **Multi-Agent hyperparameter**: Constructor Agent extracts a maximum of 2 knowledge triplets per chunk and employs hierarchical Leiden clustering (Traag et al., 2019) with maximum cluster size of 5 for community detection. Retriever Agent retrieves top-5 most relevant nodes using hybrid vector-graph similarity matching. Then, the Reranker reranks the top-2 relevant evidence nodes (or triples) within this retrieved subgraph. Reflector/Response Agent utilizes the top-2 retrieved passages as context for answer generation. **Backend Infrastructure**: LLM service is based on Qwen2.5-32B-Instruct (Yang et al., 2024) deployed with vLLM (Kwon et al., 2023) engine using bfloat16 precision and prefix caching enabled and greedy-search generation method, which is more stable than the Qwen3 model in mixed reasoning mode in our task; embeddings are generated using Jina-embeddings-v3 (1024-dimensional) (Sturua et al., 2024); we use jina-reranker-v2 (Wang et al., 2025) as the reranker model; Our server is equipped with 8 A100 40GB cards, AMD EPYC 256-core Processor, 2TB memory, and Ubuntu 20.04.1 system. and the hybrid vector-graph storage is implemented using Neo4j community edition [1] for efficient knowledge representation and retrieval, see Appendix.K for visualized graph example.

C ToG-3 Algorithms

Algorithms 1 and 2 present the two-stage pipeline of ToG-3. The first stage constructs a heterogeneous graph index comprising chunks, triplets, and communities, while the second stage implements a Multi-Agent Context Evolution and

[1] https://neo4j.com/product/community-edition

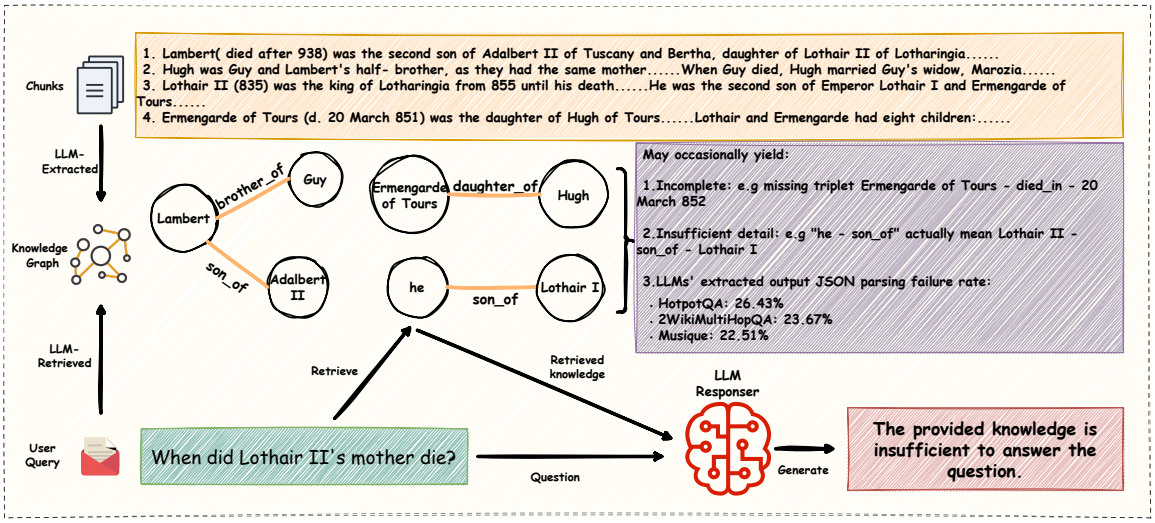

Figure 5: Performance limitations of graph-based RAG systems under resource-constrained and locally deployed scenarios. In such settings, developers often adopt open-source models such as Llama or Qwen as backbone LLMs. Incomplete extracted triplets, insufficient extraction details, and parsing failures can lead to insufficient knowledge provision, ultimately causing the system to fail to answer the query adequately.

Retrieval (MACER) loop featuring a novel dual-evolution mechanism—Evolving Query and Evolving Subgraph—that dynamically refines both the query representation and the graph structure through iterative interaction.

D Dataset Detail

This section presents a comprehensive statistical overview of the **Deep and Broad datasets** we use in this paper, including detailed statistics metadata and licensing information, as summarized in Table 3. Additionally, we provide individual descriptions of each dataset to elucidate their respective characteristics and intended use cases.

D.1 Deep Reasoning Datasets

- **HotpotQA** (Yang et al., 2018): A crowd-sourced question answering dataset built on English Wikipedia, comprising approximately 113K questions. Each question is constructed to require the combination of information from the introductory sections of two Wikipedia articles for answering. The dataset provides two gold paragraphs per question, along with a list of sentences identified as supporting facts necessary to answer the question. HotpotQA includes various reasoning strategies such as bridge questions (involving missing entities), intersection questions (e.g., "what satisfies both property A and property B?"), and comparison questions (comparing two entities

Algorithm 1 Offline Construction of Heterogeneous Index Graph $\mathcal{G}$

Require: Corpus $\mathcal{D} = \{d_i\}_{i=1}^{N}$, lightweight LM $\mathcal{L}_{\text{light}}$, encoder E_θ
Ensure: Heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

$\mathcal{V} \leftarrow \emptyset, \mathcal{E} \leftarrow \emptyset$
$\mathcal{C} \leftarrow \texttt{SplitIntoChunks}(\mathcal{D})$ ▷ Sentence-level segmentation
$\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{C}$
for each chunk $c \in \mathcal{C}$ **do**
  $\mathcal{T}_c \leftarrow \mathcal{L}_{\text{light}}(c)$ ▷ Extract semantic triplets $(s, p, o, \texttt{type}_s, \texttt{type}_p, \texttt{type}_o)$
  $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{T}_c$
  for each triplet $t \in \mathcal{T}_c$ **do**
    $\mathcal{E} \leftarrow \mathcal{E} \cup \{\textsc{MentionedIn}(t, c)\}$
  end for
end for
$G_e \leftarrow \texttt{BuildEntityCoOccurrenceGraph}(\mathcal{T})$ ▷ $\mathcal{T}$ is all triplets
$\{M_\ell\}_\ell \leftarrow \texttt{LeidenClustering}(G_e)$
for each community M_ℓ **do**
  $m_\ell \leftarrow \mathcal{L}_{\text{light}}(M_\ell)$ ▷ Generate community summary
  $\mathcal{V} \leftarrow \mathcal{V} \cup \{m_\ell\}$
  for each entity $e \in M_\ell$ **do**
    $\mathcal{E} \leftarrow \mathcal{E} \cup \{\textsc{SummaryFor}(m_\ell, e)\}$
  end for
end for
Encode every node $v \in \mathcal{V}$ using E_θ ▷ Unified dense encoding
return $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

Table 3: Statistics of Deep Reasoning and Broad Reasoning Datasets. Metrics abbreviations: Comp. (Comprehensiveness), Div. (Diversity), Emp. (Empowerment).

Dataset	Corpus Size	Chunks	Entities/Relations	Communities	Metrics	License
Deep Reasoning Tasks						
HotpotQA	9,809	9,812	37,358/30,987	5,041	EM, F1	Apache-2.0
2WikiMultihopQA	6,119	6122	19,311/21,077	3,417	EM, F1	Apache-2.0
Musique	11,254	11,300	32,842/39,134	6,258	EM, F1	CC-BY-4.0
Broad Reasoning Tasks						
CS	10	2,134	3,530/33,507	1,166	Comp., Div., Emp.	Apache-2.0
Agriculture	12	2,025	6,043/12,571	1,039		
Legal	94	5,900	26,180/44,334	1,359		
Mix	61	658	2,784/5,089	425		

Algorithm 2 ToG-3: Multi-Agent Context Evolution and Retrieval (MACER) Loop

Require: Query q, heterogeneous graph $\mathcal{G}$, LLM $\mathcal{L}$, max rounds K
Ensure: Final answer a^*

$k \leftarrow 0, \mathcal{G}_0 \leftarrow \texttt{Retriever}(q, \mathcal{G})$ ▷ Initial retrieval
$\mathcal{H}_0 \leftarrow \{(q, \mathcal{G}_0, \text{init})\}$ ▷ Initialize trajectory history
repeat
  $G_k \leftarrow \pi_{\text{rer}}(q, \mathcal{G}_{k-1})$ ▷ Reranker Agent reranks and selects the subgraph
  $a_k \leftarrow \pi_{\text{resp}}(q, \mathcal{G}_k, \mathcal{H}_k)$ ▷ Response Agent generates answer
  $r_k \leftarrow \pi_{\text{ref}}^{\text{suff}}(q, \mathcal{G}_k, a_k)$ ▷ Reflector judges sufficiency
  if $r_k = 1$ **then break**
  end if
  $q'_k \leftarrow \pi_{\text{ref}}^{\text{evolve}}(q, \mathcal{G}_k)$ ▷ Reflector evolves query
  $\mathcal{G}_{k+1} \leftarrow \pi_{\text{const}}^{\text{evolve}}(q'_k, \mathcal{G}_k)$ ▷ Constructor evolves subgraph
  $\mathcal{H}_{k+1} \leftarrow \mathcal{H}_k \cup \{(q'_k, a_k, r_k, \mathcal{G}_{k+1})\}$
  $k \leftarrow k + 1$
until $k = K$
$a^* \leftarrow \pi_{\text{resp}}^{\text{final}}(q, \mathcal{H}_k)$ ▷ Synthesize answer from full trajectory
return a^*

through a common attribute). It is available in two settings: a *few-shot distractor setting* where models are provided with 10 paragraphs including the gold ones, and an *open-domain full-wiki setting* where models must retrieve relevant passages from the entire Wikipedia corpus given only the question.

- **2WikiMultihopQA** (Ho et al., 2020): A multi-hop question answering dataset that contains complex questions requiring reasoning over multiple Wikipedia paragraphs. Each question is designed to necessitate logical connections across different pieces of information to arrive at the correct answer.

- **Musique** (Trivedi et al., 2022): A challenging multi-hop QA dataset containing approximately 25K 2–4 hop questions, constructed by composing single-hop questions from five existing single-hop QA datasets. It is designed to feature diverse and complex reasoning paths, requiring models to integrate information from multiple hops to generate correct answers. The dataset emphasizes comprehensive evaluation of multi-step reasoning capabilities.

D.2 Broad Reasoning Datasets

The following datasets are curated from the UltraDomain (Qian et al., 2025) benchmark. The benchmark construction leverages financial reports, legal contracts, and 428 college textbooks across 18 distinct domains to evaluate model versatility and adaptability in specialized and broad application scenarios:

- **CS**: Computer science domain focusing on data science, software engineering, and programming topics, requiring technical comprehension and analytical reasoning.

- **Agriculture**: Covers agricultural practices including beekeeping, crop production, and disease prevention, demanding domain-specific knowledge integration.
- **Legal**: Derived from legal contracts and documents, focusing on corporate legal practices, regulatory compliance, and governance, requiring precise interpretation of nuanced legal language.
- **Mix**: Contains diverse contexts from college textbooks spanning natural sciences, humanities, and social sciences, testing generalization capabilities across interdisciplinary topics.

E Baselines

This section presents the baseline methods evaluated in this paper, encompassing both classical algorithms such as NaiveRAG and GraphRAG, as well as recently proposed approaches including LightRAG, ToG-2, and HippoRAG-2. Baselines are as follows:

- **NaiveRAG** (Gao et al., 2023): A standard chunk-based retrieval baseline that segments raw texts into chunks and stores them in a vector database using text embeddings. For queries, it generates vectorized representations to directly retrieve text chunks based on semantic similarity.
- **GraphRAG** (Edge et al., 2024): A graph-enhanced RAG system that utilizes an LLM to extract entities and relationships from text, representing them as nodes and edges. It generates community summaries through graph clustering and employs both local (entity-based) and global (community-based) retrieval strategies for comprehensive information access.
- **LightRAG** (Guo et al., 2024): A graph-structured RAG framework that employs a dual-level retrieval system combining low-level entity retrieval with high-level knowledge discovery. It integrates graph structures with vector representations for efficient retrieval of related entities and their relationships.
- **ToG-2** (Ma et al., 2024): A knowledge graph-based framework implements a tight-coupling hybrid RAG paradigm that iteratively retrieves information from both unstructured texts and structured knowledge sources. It alternates between graph retrieval and context retrieval for in-depth knowledge exploration.
- **HippoRAG-2** (Gutiérrez et al., 2025): A non-parametric continual learning framework that leverages Personalized PageRank algorithm over an open knowledge graph constructed using LLM-extracted triples. It enhances multi-hop reasoning capabilities through sophisticated graph traversal and passage integration mechanisms.

F Metrics

We employ different evaluation protocols for the two task categories:

For **Deep Reasoning Tasks**, we follow standard QA evaluation practices as ToG (Sun et al., 2023; Ma et al., 2024) and HippoRAG (Gutiérrez et al., 2024, 2025):

- **Exact Match (EM)**: Measures the percentage of predictions that exactly match the ground truth answer. Specifically, we follows the Substring-based EM metric (used in ToG/ToG-2 (Sun et al., 2023; Ma et al., 2024)) to robustly assess answer accuracy in longer, natural-language response generated by LLMs, which goes through the whole response to check whether the answer is in.
- **F1 Score**: Computes word-level overlap between predictions and ground truth answers.

For **Broad Reasoning Tasks**, we adopt a multi-dimensional LLM-based evaluation approach due to the complexity and open-ended nature of these queries following LightRAG (Guo et al., 2024):

- **Comprehensiveness (Comp.)**: Measures how thoroughly the answer addresses all aspects of the question.
- **Diversity (Div.)**: Assesses the variety of perspectives and insights provided in the answer.
- **Empowerment (Emp.)**: Evaluates how well the answer enables informed understanding and judgment.

The LLM-based evaluation uses GPT-4o-mini as judge, with careful attention to prompt design and answer ordering to avoid positional bias. The LLM evaluation prompt is shown in Appendix.M

G More Experiment Results and Details

This section presents extended experimental results, including detailed precision and recall metrics on Deep Reasoning tasks, as well as one-to-one win rates from Broad Reasoning tasks. The pairwise win rates are converted into a unified ELO rating system, with the resulting ratings visualized in the heatmap shown in Figure 4.

G.1 Precision and Recall Rate Results

Table 4 reveals the underlying reason for the relatively low F1 scores of GraphRAG and LightRAG: these methods are not specifically designed for deep reasoning tasks. By examining both precision/recall metrics and output cases, we observe that excessively long or unfocused responses tend to substantially reduce recall, thereby diminishing overall F1 performance.

G.2 Result Detail in Braod Reasoning Tasks

Table 5 presents the pairwise win rates (%) of baseline methods against ToG-3 across four datasets and four evaluation dimensions. The results demonstrate that ToG-3 consistently outperforms all compared baselines.

G.3 ELO Rating Calculation for Broad Reasoning Tasks

This appendix details the mathematical framework and computational process for deriving ELO ratings from pairwise comparison data across four benchmark datasets. The ELO rating system provides a mathematically consistent approach to quantify relative performance differences between retrieval-augmented generation methods. The ELO rating system transforms raw win rates into a logarithmic scale that ensures transitive consistency in performance rankings. The core transformation is defined as follows:

For a given method i with win rate w_i against the reference method (ToG-3), the ELO rating difference is calculated as:

$$\Delta R_i = 400 \cdot \log_{10}\left(\frac{1}{w_i} - 1\right)$$

The absolute ELO rating for method i is then:

$$R_i = R_{\text{ref}} - \Delta R_i$$

where $R_{\text{ref}} = 1600$ is the reference rating for ToG-3.

The win probability between any two methods i and j with ratings R_i and R_j is given by:

$$P(i \text{ beats } j) = \frac{1}{1 + 10^{(R_j - R_i)/400}}$$

H Analysis of Computation Cost

H.1 Comparison of Time Consumption

Table 6 reveals a consistent accuracy-efficiency trade-off across all datasets. We observed that during the indexing phase, GraphRAG required the longest processing time, averaging 13.10 hours. This is primarily due to its need to extract a large number of triplets and generate community summaries. In comparison, both ToG-3 and LightRAG showed similar indexing times, at 10.13 and 10.06 hours respectively. Although ToG-3 also involves community summary generation, it constructs the graph more efficiently by extracting fewer relational structures during graph initialization compared to both LightRAG and GraphRAG. NaiveRAG is the fastest baseline because it does not construct or refine a graph, but its lower average EM indicates that speed comes at a substantial reasoning-quality cost. While LightRAG achieves faster inference times, it suffers from lower accuracy due to redundant graph elements or simpler retrieval mechanisms. While HippoRAG-2 achieves competitive performance and faster inference speed, it still falls short of the EM scores attained by ToG-3. GraphRAG's expensive two-stage indexing yields suboptimal results despite longer processing times. ToG-3 demonstrates an effective balance: its efficient heterogeneous graph construction produces refined knowledge bases across all datasets, and while its multi-agent reasoning requires higher inference time, this cost is directly justified by its best performance on all benchmarks, making it ideal for quality-sensitive applications requiring reliable reasoning capabilities. Note that the reranker model is relatively small and reduces the input length to the LLM, thus having minimal impact on inference time. Detailed token consumption for graph construction and inference across different methods is provided in Appendix H.2.

H.2 Comparison of Token Consumption

Our proposed ToG-3 framework achieves a more favorable balance between inference efficiency

Table 4: Comprehensive Evaluation Metrics of five RAG methods across three deep reasoning datasets. The best results of each dataset are marked in **bold**.

Method	HotpotQA			2WikiMultihopQA			Musique		
	F1	R	P	F1	R	P	F1	R	P
NaiveRAG	0.365	0.593	0.346	0.189	0.345	0.168	0.143	0.280	0.126
GraphRAG	0.011	0.423	0.006	0.018	0.456	0.009	0.008	0.266	0.004
LightRAG	0.013	0.393	0.007	0.023	0.429	0.012	0.009	0.224	0.005
MiniRAG	0.012	0.372	0.006	0.018	0.403	0.009	0.007	0.203	0.003
ToG-3	**0.569**	**0.675**	**0.492**	**0.291**	**0.496**	**0.208**	**0.174**	**0.302**	**0.122**

P: Precision, R: Recall. ToG-3 achieves best F1 while maintaining high precision-recall balance.

Table 5: Win rates (%) of baselines v.s. ToG-3 across four datasets and four evaluation dimensions. The better results of each dataset are marked in **bold**.

Metrics	**Agriculture**		**CS**		**Legal**		**Mix**	
	NaiveRAG	**ToG-3**	NaiveRAG	**ToG-3**	NaiveRAG	**ToG-3**	NaiveRAG	**ToG-3**
Comprehensiveness	26.1%	**73.9%**	30.1%	**69.9%**	10.1%	**89.9%**	32.5%	**67.5%**
Diversity	16.9%	**83.1%**	29.7%	**70.3%**	7.3%	**92.7%**	25.9%	**74.1%**
Empowerment	27.2%	**72.8%**	30.5%	**69.5%**	10.1%	**89.9%**	36.2%	**63.8%**
Overall	26.3%	**73.7%**	31.0%	**69.0%**	9.0%	**91.0%**	33.7%	**66.3%**
	GraphRAG	**ToG-3**	GraphRAG	**ToG-3**	GraphRAG	**ToG-3**	GraphRAG	**ToG-3**
Comprehensiveness	44.5%	**55.5%**	47.3%	**52.7%**	47.3%	**52.7%**	49.3%	**50.7%**
Diversity	42.1%	**57.9%**	46.1%	**53.9%**	44.5%	**55.5%**	49.7%	**50.3%**
Empowerment	22.9%	**77.1%**	40.9%	**59.1%**	27.3%	**72.7%**	36.1%	**63.9%**
Overall	45.3%	**54.7%**	46.9%	**53.1%**	46.1%	**53.9%**	48.9%	**51.1%**
	LightRAG	**ToG-3**	LightRAG	**ToG-3**	LightRAG	**ToG-3**	LightRAG	**ToG-3**
Comprehensiveness	36.6%	**63.4%**	43.3%	**56.7%**	31.3%	**68.7%**	45.3%	**54.7%**
Diversity	29.7%	**70.3%**	39.7%	**60.3%**	25.7%	**74.3%**	37.0%	**63.0%**
Empowerment	38.2%	**61.8%**	43.7%	**56.3%**	31.3%	**68.7%**	49.7%	**50.3%**
Overall	37.3%	**62.7%**	43.7%	**56.3%**	30.1%	**69.9%**	47.3%	**52.7%**
	HippoRAG-2	**ToG-3**	HippoRAG-2	**ToG-3**	HippoRAG-2	**ToG-3**	HippoRAG-2	**ToG-3**
Comprehensiveness	22.2%	**77.8%**	29.3%	**70.7%**	19.3%	**80.7%**	27.3%	**72.7%**
Diversity	16.5%	**83.5%**	25.7%	**74.3%**	15.0%	**85.0%**	21.4%	**78.6%**
Empowerment	25.5%	**74.5%**	30.6%	**69.4%**	19.3%	**80.7%**	31.7%	**68.3%**
Overall	23.3%	**76.7%**	29.7%	**70.3%**	18.1%	**81.9%**	29.3%	**70.7%**

and performance. As shown in Table 7, compared to GraphRAG, ToG-3 saves approximately 60% of token consumption during the graph construction phase (an average of 5.03 vs. 12.82 million tokens), which benefits from the dynamic graph construction mechanism that avoids the overhead of pre-building large-scale static knowledge graphs. Although ToG-3's average inference token consumption per sample (72.1 tokens) is higher than that of GraphRAG (32.3 tokens) and LightRAG (23.1 tokens), this increased inference overhead is the necessary cost for achieving precise multi-hop reasoning: our multi-agent evolution mechanism effectively decomposes complex questions and focuses on critical evidence through deep iterative query and subgraph evolution, ultimately translating into superior answer quality (as demonstrated by the performance gains in Table 2). Note that, since LLM inference speed is comparable across methods, token consumption is directly proportional to the primary time overhead.

I Additional Baselines

As shown in Table 8, under the same experimental setup, we conduct a comprehensive comparison with a range of graph-enhanced RAG baselines proposed in recent years . Across all three multi-hop reasoning benchmarks, ToG-3 significantly outperforms all compared methods on every metric. Specifically, on the HotpotQA dataset, ToG-3 achieves an EM score of 0.654, surpassing the next best performers, Youtu-GraphRAG (0.600) and Graph Counselor (0.580). A similar trend of superior performance is observed on

Table 6: Computational cost comparison across datasets between Graph-based methods. The best EM score of each dataset are marked in **bold**. ToG-3 achieves the best accuracy with efficient indexing and justified inference cost.

Dataset	Method	Graph Statistics			Indexing	Inference	Avg.
		Entities	Relations	Communities	Time (h)	Time (s/q)	EM
HotpotQA	ToG-3	37,358	30,987	5,041	12.5	17.13	**0.645**
	HippoRAG-2	92,145	22,047	-	11.2	4.85	0.612
	GraphRAG	94,376	73,265	10,981	15.8	8.91	0.337
	LightRAG	94,578	76,157	-	12.1	6.54	0.308
2WikiMultihopQA	ToG-3	19,311	21,077	3,417	8.2	15.07	**0.527**
	HippoRAG-2	48,251	11,540	-	7.6	4.12	0.491
	GraphRAG	50,556	37,840	6,261	10.3	7.45	0.439
	LightRAG	50,177	37,995	-	7.8	5.23	0.420
Musique	ToG-3	32,842	39,134	6,258	9.7	13.34	**0.291**
	HippoRAG-2	112,270	26,581	-	10.1	4.92	0.212
	GraphRAG	106,042	83,139	9,407	13.2	9.37	0.109
	LightRAG	94,621	75,923	-	10.3	7.12	0.082
Average	ToG-3	29,837	30,399	4,905	10.13	15.18	**0.474**
	HippoRAG-2	84,222	20,056	-	9.63	4.63	0.438
	GraphRAG	83,658	64,748	8,883	13.10	8.58	0.295
	LightRAG	79,792	63,358	-	10.06	6.30	0.270
	NaiveRAG	-	-	-	0.52	1.03	0.415

Table 7: Comparison of token consumption for graph construction and inference across different methods.M means Millions.

Method	Avg. Graph Construction Tokens	Avg. Inference Tokens per Sample
ToG-3	5.03M	72.1
GraphRAG	12.82M	32.3
LightRAG	4.92M	23.1
HippoRAG-2	5.01M	20.6

the 2WikiMultihopQA and Musique datasets. The consistent and comprehensive lead of ToG-3 in both EM and F1 scores demonstrates that our proposed dynamic heterogeneous graph evolution and multi-agent collaboration mechanism can more effectively support complex, deep multi-hop reasoning tasks.

We further add two reviewer-suggested comparisons that isolate different aspects of the design. KGP (Wang et al., 2024) provides a stronger graph-structured prompting baseline, while IRCoT (Trivedi et al., 2023) provides a reflective non-graph retrieval baseline. As shown in Table 9, IRCoT substantially improves over NaiveRAG, confirming the value of iterative reflection, but ToG-3 remains stronger. KGP performs worse under the same lightweight-LLM setting, suggesting that ToG-3's query-time subgraph refinement is important when the initial graph extraction is imperfect.

J Case Study for ToG-3

This section provides a detailed case study of ToG-3 in deep reasoning task (Figure 6) and broad reasoning task (Figure 7 and Figure 8), offering an intuitive demonstration of the execution dynamics of its dual-evolution mechanism—comprising Evolving Query and Evolving Subgraph—across multi-step reasoning processes.

K Graph Visualization Examples

This section details two constructed graph used in our study: the 2WikiMultihopQA subset (exemplifying deep reasoning) and the computer science domain graph from UltraDomain (exemplifying broad reasoning), which are visualized with Neo4j community edition [2].

2WikiMultihopQA Dataset: Exemplar of Depth Reasoning 2WikiMultihopQA is designed to test depth reasoning—the ability to perform multi-step, sequential inference over entity-relation paths. Each question requires traversing at least two "hops" (e.g., first identifying a person's profession, then linking that profession to a historical event, and finally combining both to answer a causal query). This structure forces models to engage in complex semantic chaining, where errors in early steps propagate, challenging robustness in long-range dependency handling. The dataset's sparse yet densely connected knowledge

[2] https://neo4j.com/product/community-edition

Table 8: Comparison of additional Graph-based RAG methods across multi-hop reasoning benchmarks. The best performance in each column is marked in **bold**.

Method	HotpotQA		2WikiMultihopQA		Musique	
	EM	F1	EM	F1	EM	F1
Youtu-GraphRAG (Dong et al., 2025)	0.600	0.450	0.470	0.230	0.205	0.135
Graph Counselor (Gao et al., 2025)	0.580	0.434	0.464	0.219	0.203	0.137
RAPTOR (Sarthi et al., 2024)	0.580	0.400	0.420	0.200	0.190	0.120
HyperGraphRAG (Luo et al., 2025b)	0.538	0.337	0.456	0.265	0.195	0.124
E²GraphRAG (Zhao et al., 2025)	0.420	0.080	0.450	0.075	0.130	0.040
Align-GRAG (Xu et al., 2025)	0.442	0.222	0.432	0.251	0.172	0.116
KET-RAG (Huang et al., 2025)	0.452	0.328	0.425	0.221	0.160	0.102
ToG-3 (Ours)	**0.654**	**0.569**	**0.527**	**0.291**	**0.241**	**0.174**

Table 9: Reviewer-suggested baseline comparisons under identical settings.

Method	**Avg. EM**	**Avg. F1**
KGP (Wang et al., 2024)	0.275	0.125
IRCoT (Trivedi et al., 2023)	0.435	0.301
HippoRAG-2 (Gutiérrez et al., 2025)	0.438	0.311
ToG-3 (Ours)	**0.474**	**0.345**

graphs emphasize precision in step-by-step reasoning over surface-level pattern matching. A structural overview highlighting its multi-hop nature is shown in Figure 9.

Computer Science Domain Graph in UltraDomain: Exemplar of Breadth Reasoning The computer science domain graph from UltraDomain represents breadth reasoning—focused on expansive coverage of concepts and their interrelations. It includes a wide range of CS entities (from foundational data structures/algorithms to applied distributed systems/cloud services) and diverse relationship types (e.g., `implements`, `runs_on`, `contains`). This breadth challenges models to navigate a large, heterogeneous concept space, where connections span disparate subfields (e.g., linking a programming language to a database, or an algorithm to hardware). For instance, understanding how Spark relates to Hadoop, Kafka, and multiple programming languages requires integrating knowledge across multiple domains, reflecting the need for broad, cross-concept awareness. A visualization of this graph, illustrating its extensive node and edge diversity, is provided in Figure 10.

L Theoretical Support: Implicit Dynamics of In-Context Learning

The iterative refinement process in MACER and dual-evolving mechanism is not merely heuristic but possesses theoretical grounding through the lens of implicit in-context learning dynamics. Recent work by (Dherin et al., 2025) demonstrates that transformer-based models can perform in-context learning by implicitly modifying their MLP weights through attention mechanisms. We extend this theoretical framework to explain the convergence properties of our multi-agent reasoning process.

Implicit Weight Updates via Attention Dynamics The trajectory history $\mathcal{H}_k$ serves as an *in-context prompt* that induces implicit low-rank updates to the frozen LLM's parameters. Specifically, for a transformer module with MLP layer weights W, the context $\mathcal{H}_k$ generates an implicit weight update ΔW_k through the attention mechanism:

$$\Delta W_k = \frac{(W \Delta A_k) A(q)^\top}{\|A(q)\|^2}, \quad \text{where} \quad \Delta A_k = A(\mathcal{H}_k, q) - A(q). \tag{7}$$

Here, $A(\cdot)$ denotes the activation pattern from the attention layer, $A(q)$ represents the baseline activation without context, and $A(\mathcal{H}_k, q)$ captures the contextualized activation with the full reasoning history. The term ΔA_k quantifies the information injected by the evolving context $\mathcal{H}_k$. The low-rank nature of ΔW_k ensures efficient and targeted parameter updates without catastrophic forgetting of pre-trained knowledge.

Evolving Query and Subgraph Case Study

Question:
What nationality is the performer of the song When The Stars Go Blue?

Initial Evidence (Subgraph):
When The Stars Go Blue -¿ performed_by -¿ Ryan Adams
Ryan Adams -¿ known_for -¿ alternative country, rock, folk
Ryan Adams -¿ also_includes_genre -¿ indie rock, Americana
... (no nationality information present)

Evolving Query Decomposition 1 (Reflector Agent)
Who is the performer of the song When The Stars Go Blue?
Ryan Adams.

Evolving Subgraph Refinement (Constructor Agent):
Sammy Hagar -¿ is -¿ American rock vocalist
Chuck Berry -¿ is -¿ American singer and songwriter
Elvis Presley -¿ pioneered -¿ Rock and Roll (American origin)
Ryan Adams -¿ nationality -¿ American [Added]
Wee Willie Walker -¿ was -¿ American soul and blues singer

Evolving Query Decomposition 2 (Reflector Agent)
What is Ryan Adams known for besides the song When The Stars Go Blue?
Ryan Adams is known for his work as a singer-songwriter and musician, with a musical style that encompasses alternative country, rock, and folk.

Evolving Query Decomposition 3 (Reflector Agent)
What is Ryan Adams's nationality?
Ryan Adams is an American singer and musician.

Final Answer (Response Agent):
ryan adams is american.

Ground Truth:
america

Explanation:
The initial subgraph contained no explicit information about Ryan Adams' nationality.
Through evolving query decomposition, the system identified a missing link and triggered subgraph refinement.
New evidence from biographical triples (e.g., "Ryan Adams -¿ nationality -¿ American") was integrated, enabling a definitive answer.

Figure 6: Case Study: Evolving Query and Evolving Subgraph for Multi-Hop Reasoning. This example demonstrates how an initial knowledge subgraph fails to answer a complex question, prompting a decomposition into a sub-query. The Constructor Agent refines the subgraph with additional biographical facts, enabling the Response Agent to generate the correct answer. The process highlights the dynamic, iterative nature of self-evolving context retrieval.

MDP Policy as an Implicit Function of Context Recall from Section 3.3 that the Reflector Agent's policy π_{ref} maps states $s_k = (q, \mathcal{G}_k, \mathcal{H}_k)$ to actions (sub-queries or STOP). Under the implicit learning view, π_{ref} is not a fixed network but an emergent policy π_k shaped by ΔW_k. Thus, the sequence $\{\pi_k\}_{k=1}^{K}$ constitutes a trajectory of implicitly adapted policies driven by the evolving context $\mathcal{H}_k$.

Convergence via Regret Minimization We analyze convergence through the lens of episodic regret minimization in the MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r)$. Let $V_{s_k}^{\pi} = \mathbb{E}_{\pi}\left[\sum_{i=k}^{K} \gamma^{i-k} r_i \mid s_k\right]$ denote the value of policy π at state s_k, and let $V_{s_k}^{*} = \max_{\pi} V_{s_k}^{\pi}$ be the optimal value. The cumulative regret over K steps is:

$$\mathcal{R}(K) = \sum_{k=1}^{K} \left(V_{s_k}^{*} - V_{s_k}^{\pi_k}\right). \tag{8}$$

We establish sublinear regret growth $\mathcal{R}(K) = o(K)$ under the following mild assumptions:

Assumption 1 (Realizability)**.** *There exists a policy π^* such that $\texttt{Suff}(q, \mathcal{G}_q^*) = 1$, and π^* is representable by the implicit policy class induced by in-context prompts of the form $(\mathcal{H}; q)$.*

Assumption 2 (Bounded Gradient Norm)**.** *The implicit gradient direction g_k, defined as the reward-sensitive update signal from $\mathcal{H}_k$, satisfies $\|g_k\| \leq G$ for some constant $G > 0$.*

Under these assumptions, the following properties hold:

Property 1 (Smooth Policy Evolution). The value function evolves smoothly with respect to implicit updates:

$$\|V^{\pi_{k+1}} - V^{\pi_k}\|_\infty \leq L\|g_k\| + \mathcal{O}(\|g_k\|^2), \quad (9)$$

for some Lipschitz constant $L > 0$, ensuring stable policy transitions.

Property 2 (Expected Policy Improvement). Each refinement step yields non-negative expected improvement:

$$\mathbb{E}\left[V_{s_k}^{\pi_{k+1}} - V_{s_k}^{\pi_k} \mid \mathcal{H}_k\right] \geq \eta\|g_k\|^2 - \sigma_k, \quad (10)$$

where $\eta > 0$ and $\{\sigma_k\}$ is a martingale difference sequence with $\mathbb{E}[\sigma_k \mid \mathcal{H}_k] = 0$. This follows from the fact that evolving sub-queries generated by the Reflector target knowledge gaps, and the Constructor's evolving graph refinement increases the likelihood of sufficiency.

Property 3 (Vanishing Implicit Gradient). As the context becomes increasingly informative, the room for improvement diminishes:

$$\lim_{k\to\infty} \|g_k\| = 0 \quad \text{almost surely.} \quad (11)$$

This is guaranteed by Assumption 1 (Realizability) and the finite horizon K, which ensures the process either reaches a sufficient subgraph ($r_k = 1$) or exhausts its budget.

Together, these properties imply that the sequence $\{\pi_k\}$ converges to a policy $\pi^\dagger$ satisfying $V_{s_1}^{\pi^\dagger} \geq V_{s_1}^* - \epsilon$ for arbitrarily small $\epsilon > 0$ as $K \to \infty$. In practice, with a reasonable horizon (e.g., $K = 3$), MACER reliably converges to a sufficient context $\mathcal{G}_q^*$ for faithful answer synthesis.

This analysis establishes that the MACER loop performs an implicit form of policy gradient ascent on the reward landscape defined by context sufficiency, with convergence guarantees rooted in stochastic approximation theory and in-context learning dynamics, providing rigorous foundations for the empirical effectiveness of our reward-based evolving context mechanism.

M Prompt Templates

Our framework employs a multi-stage, prompt-driven reasoning pipeline that integrates structured knowledge graph (KG) extraction, community-based summarization, iterative sub-query decomposition, subgraph refinement, and faithful answer synthesis. Each stage is governed by a specialized prompt template designed to ensure modularity, interpretability, and factual consistency. The complete sequence of prompts is as follows:

1. **KG Triplets Extraction**: As shown in Figure 11, given raw textual input, this prompt instructs the model to extract structured subject-relation-object triples (e.g., `entity1 -> relation -> entity2`) to construct a fine-grained knowledge subgraph. This step transforms unstructured text into a queryable graph structure.

2. **Generate Community Summary**: As shown in Figure 12, based on densely connected subgraphs (communities), this prompt synthesizes a concise natural language summary that captures the core themes and relationships within each community, enabling high-level semantic indexing and retrieval.

3. **Keyword Expansion for Retrieval Augmentation**: As shown in Figure 13, to improve recall in the querying phase, this prompt generates a set of synonyms and related terms from the original query, considering variations in capitalization, pluralization, and common phrasings, separated by delimiter symbols.

4. **Evolving Sub-Query Decomposition**: As shown in Figure 14, for complex multi-hop questions, this prompt recursively decomposes the current query into simpler, context-answerable sub-questions, guided by previously retrieved information and reasoning traces, enabling stepwise information gathering.

5. **Evolving Subgraph Refinement**: As shown in Figure 15, this prompt cleans and enhances the retrieved or extracted subgraph by removing irrelevant triples, normalizing entity names, and optionally filling in strongly supported missing links, thereby improving the signal-to-noise ratio for downstream reasoning.

6. **Final Answer Synthesis**: As shown in Figure 16, in the final stage, the model generates a concise, context-grounded answer using *only* the refined evidence, with explicit instructions to avoid hallucination or reliance

on prior knowledge. If the answer cannot be determined, it returns “Unknown” to maintain factual integrity.

These prompts work in concert to enable structured, interpretable, and reliable reasoning over hybrid text-and-graph knowledge sources. And Figure 17 shows the LLM evaluation prompt in the broad reasoning task. Their modular design allows for independent tuning and auditing, making the overall system transparent and robust to noise and ambiguity.

Regression Metrics Case Study

Question:
How do different regression metrics (e.g., RMSE, MAE, R^2) vary by model and dataset?

Evidence (from ToG-3):
Mean Squared Error -¿ is_related_to -¿ Root Mean Squared Error...
Mean Squared Error -¿ is_similar_to -¿ Mean Absolute Error...
summary_23 -¿ SUMMARY_FOR -¿ RMSE.MSE emphasizes larger errors...
summary_287 -¿ SUMMARY_FOR -¿ MAE. ARIMA model performance evaluated using MAE...

Answer 1(from ToG-3):
Different regression metrics such as RMSE, MAE, and R^2 can vary by model and dataset...
RMSE penalizes large errors more heavily due to squared terms, while MAE is robust to outliers...
R^2 measures explained variance but can be misleading on small or overfitted datasets...

Answer 2(from GraphRAG):
The provided community summary does not detail specific variations...To evaluate how these metrics vary:
1. Train different regression models on your datasets...
2. Compute RMSE, MAE, R^2 using MLlib.RegressionMetrics...
3. Compare performance across models and datasets...

Evaluation Result (JSON):

```
{
  "Comprehensiveness": {
    "Winner": "Answer 1",
    "Explanation": "Answer 1 provides a detailed explanation of various regression metrics..."
  },
  "Diversity": {
    "Winner": "Answer 2",
   "Explanation": "Answer 2 presents a different perspective by incorporating specific tools..."
  },
  "Empowerment": {
    "Winner": "Answer 1",
    "Explanation": "Answer 1 empowers the reader by explaining what each metric means..."
  },
  "Overall Winner": {
    "Winner": "Answer 1",
   "Explanation": "Answer 1 is the overall winner as it provides a comprehensive understanding..."
  }
}
```

Figure 7: Case Study: Comparing Regression Metrics Across Models and Datasets. This example illustrates how two reasoning systems answer a technical ML question: GraphRAG emphasizes practical implementation (e.g., using Spark's MLlib), while ToG3 focuses on theoretical distinctions between RMSE, MAE, and R^2. An evaluator selects the more comprehensive and empowering answer based on evidence from the knowledge graph.

Food Access Policy Case Study

Question:
What policy changes are needed to foster equitable access to nutritious food?

Evidence (from ToG-3):
summary_216 -¿ SUMMARY_FOR -¿ Food Security. **Supermarket and Food Security**: The reliance on supermarkets for food is viewed as a threat...
Victory Garden -¿ promoted -¿ Food Security... Supermarket -¿ threatens -¿ Food Security..."
White House Childhood Obesity Task Force -¿ recommended -¿ Farm to School...
Food Works -¿ collaborates_with -¿ Sauvie Island Organics...
Michael Taylor -¿ interprets -¿ Food Modernization Act... background in representing Monsanto...

Answer 1(from ToG-3):
policy changes that could foster equitable access to nutritious food may include implementing subsidies...
increasing funding for farm to school programs in underserved communities...
establishing regulations that prioritize the availability of nutritious options in school meal programs...
and creating educational initiatives to promote awareness about nutrition and healthy eating habits...
these changes could help address disparities in access to nutritious food...

Answer 2(from GraphRAG):
1. **strengthen conflict of interest policies**: prevent individuals with significant industry ties...
2. **enhance transparency and accountability**: implement transparent processes...
3. **support diverse food producers**: promote a variety of food producers...
4. **expand nutrition education and assistance programs**: enhance nutrition education...
5. **ensure regulatory independence**: guarantee that regulatory bodies are independent...
6. **promote local food systems**: encourage local food production through grants, subsidies...
7. **support community gardens and urban agriculture**: provide funding and resources...
8. **promote sustainable farming practices**: offer incentives for sustainable farming...
These measures aim to create a more resilient and sustainable food system...

Evaluation Result (JSON):

```
{
{
  "Comprehensiveness": {
    "Winner": "Answer 2",
    "Explanation": "Answer 2 provides a more detailed and structured set of policy changes... "
    "broader scope than Answer 1..."
  },
  "Diversity": {
    "Winner": "Answer 2",
    "Explanation": "Answer 2 presents a wider variety of perspectives... "
    "including regulatory, community, and economic solutions."
  },
  "Empowerment": {
    "Winner": "Answer 1",
    "Explanation": "Answer 1 empowers readers by offering a comprehensive overview..."
    "facilitates informed judgment."
  },
  "Overall Winner": {
    "Winner": "Answer 2",
    "Explanation": "Answer 2 emerges as the overall winner due to its superior"
    "comprehensiveness, diversity, and empowerment."
  }
}
```

Figure 8: Case Study: Policy Recommendations for Equitable Food Access. This example illustrates the full reasoning pipeline: a complex policy question is answered by two different systems (GraphRAG and ToG-3), supported by retrieved knowledge snippets. An evaluator then compares both responses across multiple dimensions, selecting the more comprehensive, diverse, and empowering answer as the winner.

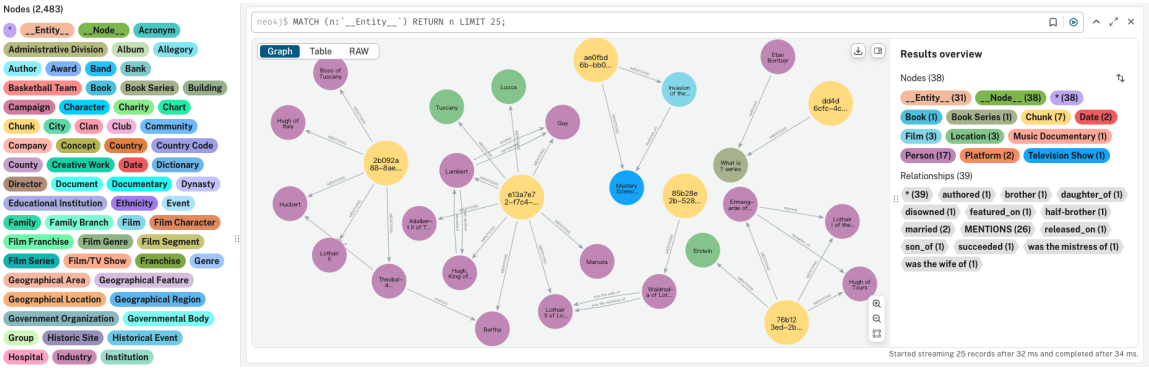

Figure 9: Structural overview of the 2WikiMultihopQA subset, exemplifying depth reasoning through multi-hop entity-relation paths (e.g., traversing "person → profession → historical event" to answer causal queries).

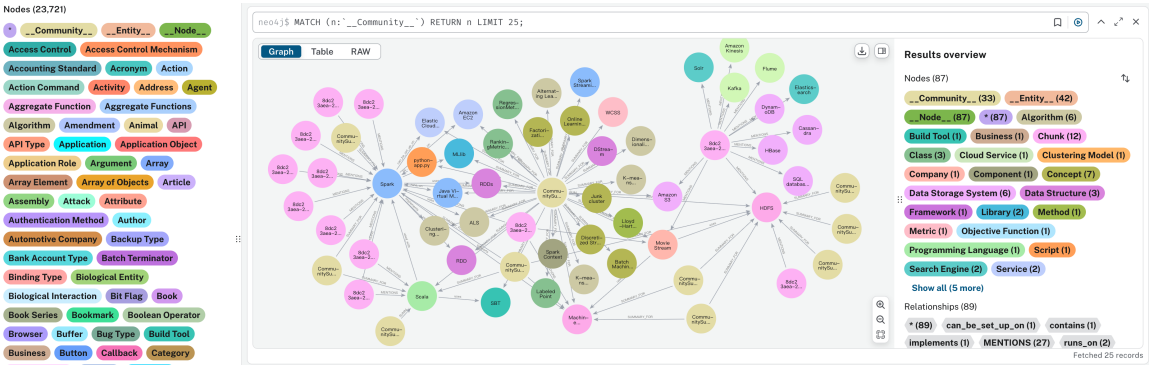

Figure 10: Visualization of the computer science domain graph in UltraDomain, showcasing breadth reasoning via diverse node types (e.g., programming languages like Scala/Spark, frameworks like HDFS/Kafka) and relationship types (e.g., `implements`, `runs_on`, `contains`).

KG Triplets Extraction Prompt Template

-Goal-
Given a text document, identify all entities and their entity types from the text and all relationships among the identified entities.
Given the text, extract up to **{max_knowledge_triplets}** entity-relation triplets.

-Steps-
1. Identify all entities. For each, extract:
entity_name — entity_type — entity_description

2. Identify all related (source, target) pairs. For each, extract:
source_entity — target_entity — relation — relationship_description

3. Output valid JSON only:
{ "entities": [...], "relationships": [...] }

-An Output Example-
{
"entities": [
{ "entity_name": "Albert Einstein", "entity_type": "Person", "entity_description": "..." },
{ "entity_name": "Theory of Relativity", "entity_type": "Scientific Theory", "entity_description": "..." },
{ "entity_name": "Nobel Prize in Physics", "entity_type": "Award", "entity_description": "..." }
],
"relationships": [
{ "source_entity": "Albert Einstein", "target_entity": "Theory of Relativity", "relation": "developed", "relationship_description": "..." },
{ "source_entity": "Albert Einstein", "target_entity": "Nobel Prize in Physics", "relation": "won", "relationship_description": "..." }
]
}

-Real Data-
##################
text: **{text}**
##################
output:

Figure 11: KG Triplets Extraction Prompt Template. The template provides structured instructions for extracting entities and relationships from text, with clear formatting for both input requirements and JSON output format.

Community Summary Prompt Template

role="system"
You are provided with a set of relationships from a knowledge graph, each represented as
`entity1 -> entity2 -> relation -> relationship_description`.
Your task is to create a summary of these relationships. The summary should include: Names of the entities involved, A concise synthesis of the relationship descriptions. The goal is to capture the most critical and relevant details that highlight the nature
and significance of each relationship. Ensure the summary is coherent and integrates
information to emphasize key aspects. Avoid redundancy and maintain clarity.

role="user"
##################
`text: {community_info}`
##################

assistant:
`% Generated summary based on {community_info} will appear here.`

Figure 12: Community Summary Template. This template provides structured instructions for extracting entities and relationships from text, with clear formatting for input specifications and expected JSON-like output format.

Keyword Expansion Prompt Template

role="system"
Given some initial query, generate synonyms or related keywords up to {max_keywords} in total,
considering possible cases of capitalization, pluralization, common expressions, etc.
Provide all synonyms/keywords separated by '^' symbols: 'keyword1^keyword2^...'.
Note: result should be in one line, separated by '^' symbols.

role="user"

QUERY: {query_str}
---

assistant:
% Example: KEYWORDS: machine learning^ML learning machines^AI models^neural networks^deep learning ...

Figure 13: Keyword Expansion Prompt Template. This template instructs the model to generate up to {max_keywords} synonyms or related terms for a given query, formatted as a single line separated by '^' symbols.

Step-wise Query Evolution and Decomposition Prompt Template

role="system"
The original question is as follows: {query_str}
We have an opportunity to answer some, or all of the question from a knowledge source.
Context information for the knowledge source is provided below, as well as previous reasoning steps.
Given the context and previous reasoning, return a question that can be answered from the context.
This question can be the same as the original question, or represent a subcomponent.
It should not be irrelevant to the original question.
If no further information can be extracted, return 'None'.

Examples:

Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
Knowledge source context: Provides names of the winners of the 2020 Australian Open
Previous reasoning: None
Next question: Who was the winner of the 2020 Australian Open?

Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
Knowledge source context: Includes biographical info for each winner
Previous reasoning:
- Who was the winner of the 2020 Australian Open?
- The winner was Novak Djokovic.
Next question: How many Grand Slam titles does Novak Djokovic have?

Current Input:

Question: {query_str}
Knowledge source context: {context_str}
Previous reasoning: {prev_reasoning}

assistant:
% Output: <decomposed sub-question> OR 'None'

Figure 14: Step-wise Query Evolution and Decomposition Prompt Template. This template guides the model to recursively break down a complex question into answerable sub-questions based on available context and prior reasoning, enabling multi-hop reasoning over knowledge sources.

Subgraph Evolution and Refinement Prompt Template

role="system"
You are given a subgraph extracted from a knowledge graph, represented as a list of triples:
`entity1 -> relation -> entity2`.
This subgraph may contain irrelevant, redundant, or incomplete information.
Your task is to **refine** the subgraph by:
Removing irrelevant or noisy triples not related to the query, Filling in missing but inferable relationships (if strongly supported),
Ensuring entity names are normalized (e.g., consistent capitalization, singular/plural).
Return the refined subgraph in the same triple format, one per line.
If no refinement is needed, return the original subgraph.
If all triples are irrelevant, return 'None'.

Example Input:

Query: What are the major achievements of Marie Curie?
Sub-graph:
Marie Curie -¿ won -¿ Nobel Prize in Physics
Marie Curie -¿ born in -¿ Warsaw
Marie Curie -¿ spouse -¿ Pierre Curie
Apple Inc. -¿ founded by -¿ Steve Jobs

Refined Output:

```
Marie Curie -> won -> Nobel Prize in Physics
Marie Curie -> won -> Nobel Prize in Chemistry
Marie Curie -> spouse -> Pierre Curie
```
(Note: Added Chemistry prize based on strong prior knowledge; removed birthplace and unrelated Apple fact)

Current Input:

Query: {query_str}
Sub-graph:
{subgraph_triples}

assistant:

Figure 15: Subgraph Evolution and Refinement Prompt Template. This template guides the model to clean, complete, and normalize a noisy or incomplete knowledge subgraph in response to a given query, improving its relevance and coherence for downstream reasoning.

Final Answer Synthesis Prompt Template

role="system"
Context information is provided below.
You must answer the query **using only this context**, and **not** any prior knowledge.
Do not make assumptions or add information not present in the context.
If the answer cannot be determined from the context, respond with 'Unknown'.

{context_str}

Query: {query_str}

Instructions:
Extract or synthesize the answer strictly from the provided context.
Keep the answer concise and factual.
Avoid phrases like "The context states that..." — just give the answer.

assistant:
`% Final answer derived solely from context.`

Figure 16: Final Answer Synthesis Prompt Template. This template enforces faithful response generation based exclusively on retrieved context, a core principle in Retrieval-Augmented Generation (RAG) systems. It suppresses model hallucination by explicitly forbidding the use of prior knowledge.

Answer Evaluator Prompt Template

role="system"
You are an expert tasked with evaluating two answers to the same question based on three criteria: **Comprehensiveness**, **Diversity**, and **Empowerment**.

Evaluation Criteria:

• Comprehensiveness:
How much detail does the answer provide to cover all aspects and sub-questions implied by the original query?

• Diversity:
How varied and rich is the answer in providing different perspectives, evidence sources, or reasoning paths?

• Empowerment:
How well does the answer help the reader understand the topic and make informed judgments or decisions?

Instructions:
Compare Answer 1 and Answer 2 for each criterion.
Choose the better answer and explain why.
Select an overall winner based on balance across all three.

Input:

Question: {query}
Answer 1: {answer1}
Answer 2: {answer2}

Output Format (JSON):

```
{
  "Comprehensiveness": {
    "Winner": "Answer 1 or Answer 2",
    "Explanation": "..."
  },
  "Diversity": {
    "Winner": "Answer 1 or Answer 2",
    "Explanation": "..."
  },
  "Empowerment": {
    "Winner": "Answer 1 or Answer 2",
    "Explanation": "..."
  },
  "Overall Winner": {
    "Winner": "Answer 1 or Answer 2",
    "Explanation": "..."
  }
}
```

Figure 17: Answer Evaluator Prompt Template. This template guides a dedicated agent to compare two candidate responses along three dimensions: comprehensiveness, diversity, and empowerment, promoting high-quality, informative, and user-centered answer selection in multi-agent systems.