

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 THINK-ON-GRAPH 3.0: EFFICIENT AND ADAPTIVE LLM REASONING ON HETEROGENEOUS GRAPHS VIA MULTI-AGENT DUAL-EVOLVING CONTEXT RE- TRIEVAL

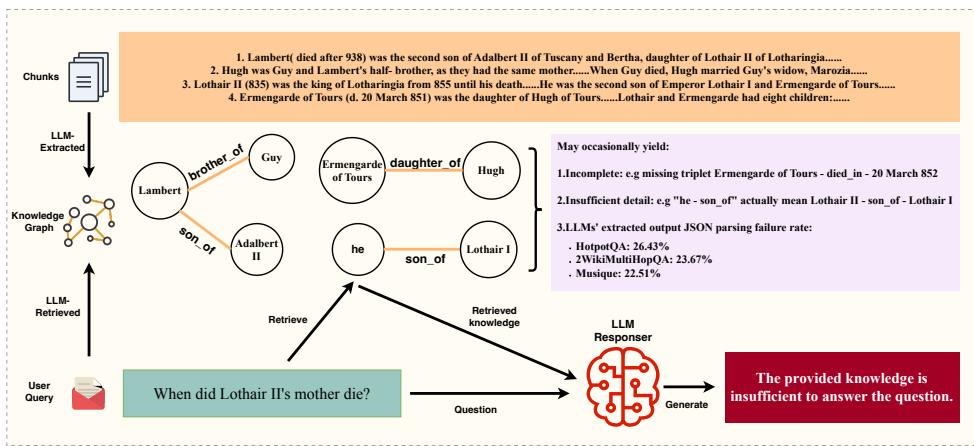
## ABSTRACT

Retrieval-Augmented Generation (RAG) and Graph-based RAG has become the important paradigm for enhancing Large Language Models (LLMs) with external knowledge. However, existing approaches face a fundamental trade-off. While graph-based methods are inherently dependent on high-quality graph structures, they face significant practical constraints: manually constructed knowledge graphs are prohibitively expensive to scale, while automatically extracted graphs from corpora are limited by the performance of the underlying LLM extractors, especially when using smaller, local-deployed models. This paper presents Think-on-Graph 3.0 (ToG-3), a novel framework that introduces Multi-Agent Context Evolution and Retrieval (MACER) mechanism to overcome these limitations. Our core innovation is the dynamic construction and refinement of a Chunk-Triplets-Community heterogeneous graph index, which pioneeringly incorporates a dual-evolution mechanism of Evolving Query and Evolving Sub-Graph for precise evidence retrieval. This approach addresses a critical limitation of prior Graph-based RAG methods, which typically construct a static graph index in a single pass without adapting to the actual query. A multi-agent system, comprising Constructor, Retriever, Reflector, and Responser agents, collaboratively engages in an iterative process of evidence retrieval, answer generation, sufficiency reflection, and, crucially, evolving query and subgraph. This dual-evolving multi-agent system allows ToG-3 to adaptively build a targeted graph index during reasoning, mitigating the inherent drawbacks of static, one-time graph construction and enabling deep, precise 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 components of 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. Notably, open-source models play a crucial role in enabling AI applications in offline environments. However, current LLMs still face notable limitations, including issues with factual hallucinations and inadequate performance in complex reasoning tasks. 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. While traditional RAG systems rely on vector similarity to retrieve relevant text chunks, they often struggle with complex reasoning tasks that require integrating information across multiple documents or understanding structural relationships between entities. To address the above limitations, recent advancements have explored using Knowledge Graphs (KGs) or extracted Graph using LLMs to represent and retrieve structured information. 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, as shown in Figure 1, the quality of the generated graph is highly dependent on the LLM’s ability to accurately extract entities and

054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067



069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107

**Figure 1:** Performance Limitations of Graph-Based RAG systems under Resource-Constrained and Locally-Deployed Scenarios. In such scenarios, developers typically adopt open-source models such as Llama or Qwen as the backbone LLMs. Limitations like incomplete extracted triplets, insufficient extraction details and parsing failure may lead to insufficient knowledge provision, ultimately resulting in failure to adequately answer the query.

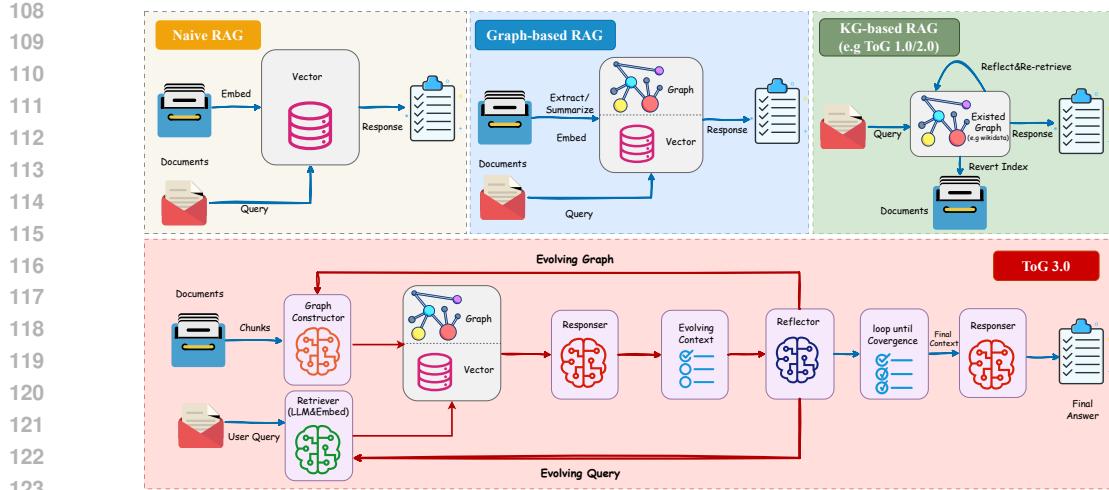
relationships, which can be a bottleneck for lightweight models like Qwen2.5-7B~72B (Yang et al., 2024), which is broadly deployed in private and offline environments. Moreover, these methods often separate the handling of local and global questions.

To overcome these limitations, we introduce **Think-on-Graph 3.0** (ToG-3), a new RAG framework that integrates the strengths of both paradigms. Our core contribution lies in the introduction of a novel Chunk-Triplets-Community heterogeneous graph architecture and a novel MACER (Multi-Agent Context Evolution and Retrieval) mechanism, which pioneeringly incorporates a dual-evolution mechanism of **Evolving Query** and **Evolving Sub-Graph** for precise evidence retrieval. Figure 2 illustrates the key distinctions between ToG-3 and classical RAG paradigms such as NaiveRAG and GraphRAG. ToG-3 introduces a novel dual-evolution mechanism—comprising Evolving Query and Evolving Subgraph—that dynamically refines both the query representation and the graph structure in an iterative manner. This approach addresses a critical limitation of prior RAG methods, which typically construct a static graph index in a single pass without adapting to the actual query. The framework is particularly suited for resource-constrained and on-premises deployment scenarios, where lightweight open-source LLMs (e.g., Llama or Qwen) are often employed as the backbone of the RAG system.

Extensive experiments on complex multi-hop reasoning benchmarks demonstrate that our method achieves the highest average Exact Match and F1 scores on HotpotQA, 2WikiMultihopQA, and Musique. For broad reasoning tasks, ToG-3 also achieves remarkable win rates over baselines across comprehensiveness, diversity, and empowerment dimensions.

Our key contributions are summarized as follows:

1. We propose **MACER** (Multi-Agent Context Evolution and Retrieval), a novel multi-agent framework that introduces a dual-evolution mechanism integrating Evolving Query and Evolving Sub-Graph within graph-based RAG. This design significantly enhances retrieval performance and complex reasoning capabilities, especially when using lightweight open-source LLMs as the backbone of the RAG system.
2. We present **ToG-3**, a unified reasoning system that effectively combines the complementary advantages of prior graph-based and ToG methods through a Chunk-Triplet-Community Heterogeneous Graph Index and a Dual-Evolving Context Retrieval Loop Process.
3. We conduct extensive experiments on both **Deep and Broad Reasoning Tasks**, demonstrating that our approach consistently supports multi-hop inference and large-scale contextual integration, achieving competitive results across diverse benchmarks.



**Figure 2: 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 *four-agent* loop—Retriever, Constructor, Reflector, Response—where the graph and the query sub-tasks *co-evolve* at runtime, yielding dynamic, query-adaptive context that converges to a minimal, sufficient subgraph.

## 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), Chen et al. (2023) enhance sense-making but often introduce noise through uncontrolled summarization or lack explicit support for multi-hop reasoning.

### 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-RETEGEN (Shao et al., 2023) introduces an iterative loop that alternates between retrieval and generation, using generated hypotheses to guide further search. Trivedi et al. (2023) combine Chain-of-Thought (CoT) with retrieval, interleaving reasoning steps with evidence gathering, significantly

162 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.

### 171 3 METHODOLOGY

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

#### 176 3.1 PROBLEM FORMULATION

178 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  
 179 is both accurate and *faithful* to the source corpus, derived from a *minimal, sufficient subgraph*  $\mathcal{G}_q^*$  of  
 180 a heterogeneous graph  $\mathcal{G}$  constructed from  $\mathcal{D}$ :

$$181 \quad \mathcal{G}_q^* = \arg \min_{\mathcal{G}' \subseteq \mathcal{G}} |\mathcal{G}'| \quad \text{subject to} \quad \text{Suff}(q, \mathcal{G}') = 1, \quad (1)$$

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

185 Existing methods face a critical dilemma: **(1)** Systems like ToG-1 or 2 rely on high-quality, pre-  
 186 constructed KGs, limiting their applicability to private or specialized domains. **(2)** Corpus-based  
 187 GraphRAG methods (e.g., GraphRAG, LightRAG) build a static graph from  $\mathcal{D}$  in one go. Their  
 188 performance is bottlenecked by the quality of this initial graph, which in turn depends heavily on  
 189 the capability of the LLM used for information extraction.

#### 191 3.2 HETEROGENEOUS GRAPH INDEX: SCHEMA AND CONSTRUCTION

##### 192 3.2.1 NODE AND EDGE SCHEMA

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

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

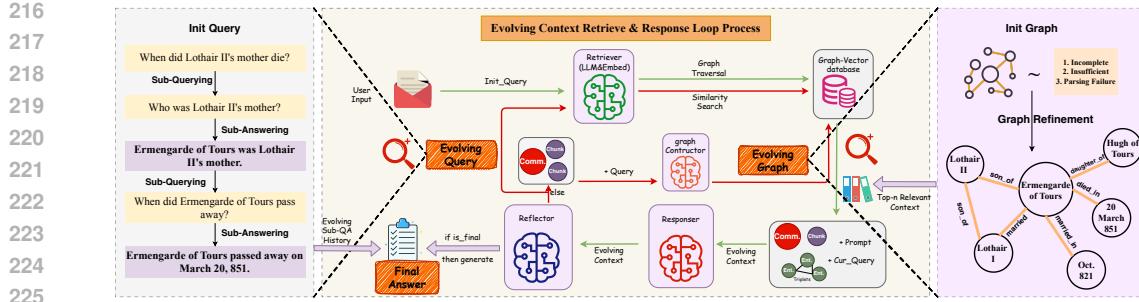
202 Edges are defined by three type relations:

- 203 • **OPENREL**( $s, p, o$ ): Connects entities  $s$  and  $o$  via predicate  $p$  extracted by the LLM, forming  
 204 an open-domain semantic triple.
- 205 • **MENTIONEDIN**( $t, c$ ): Connects a triplet  $t$  to the chunk  $c$  from which it was extracted.
- 206 • **SUMMARYFOR**( $m, e$ ): Connects a community summary node  $m$  to an entity  $e$  that be-  
 207 longs to that community.

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

##### 213 3.2.2 OFFLINE INDEX CONSTRUCTION

215 Algorithm 1 in Appendix B details the one-time construction of the universal index  $\mathcal{G}$ . A key  
 design choice is the use of a single frozen encoder  $E_\theta$  (e.g., jina-mbedding-v3 (Sturua et al., 2024))



**Figure 3: Multi-Agent Dual-Evolving Context Retrieval-Response Loop.** The Retriever fetches an initial chunk–triplet–community subgraph. 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 (sub-graph 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.

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  $\pi_{\text{ret}}$ , 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 sub-graphs.

**Action Space ( $\mathcal{A}$ )** : The Reflector Agent  $\pi_{\text{ref}}$  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  $r_k$ :

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

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  $\pi_{\text{const}}$  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 including iterative sequence of *evolving queries* and *evolving sub-graphs* reflects the structural evolution of the graph based on the agent’s reasoning action, formally defined by the recurrence:

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

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

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, 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.

270

271 
$$\mathcal{H}_{k+1} = \mathcal{H}_k \cup (q'_k, a_k, r_k, \mathcal{G}_{k+1}) \quad (5)$$

272 
$$a^* \leftarrow \pi_{\text{resp}}^{\text{final}}(q, \mathcal{H}_k) \quad (6)$$

273

274 The complete MACER process, now cast as an MDP, is summarized in Algorithm 2. The loop  
 275 continues until the Reflector’s policy  $\pi_{\text{ref}}$  outputs the STOP action (via  $r_k = 1$ ) or a maximum  
 276 horizon  $K$  is reached. The final answer  $a^*$  is synthesized from the full trajectory  $\mathcal{H}_k$  of states and  
 277 actions, ensuring faithfulness to the evolved evidence. This MDP formulation provides the formal  
 278 foundation for establishing the convergence of the MACER process under mild assumptions, as  
 279 detailed in Appendix. I. This iterative refinement allows ToG-3 to start from a potentially weak  
 280 initial graph but *specialize* it towards the reasoning path of the specific query, converging on a  
 281 high-quality evidence subgraph  $\mathcal{G}_q^*$ . This evolving and refinement mechanism alleviate the three  
 282 fundamental weaknesses of small LMs in static GraphRAG, including incomplete triplet recall,  
 283 insufficient knowledge details and high parsing failure of LLMs’ output, as mentioned in Section 1.

284

285 

## 4 EXPERIMENT

286

287 

### 4.1 EXPERIMENTAL SETUP

288

**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 4 and Appendix. C.

289

**Evaluation Metrics** For **Deep Reasoning Tasks**, we follow standard QA evaluation practices with **Exact Match (EM)** 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). Metrics detail are provide Appendix.E.

290

**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.D. 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 provide Appendix.A.

291

292 

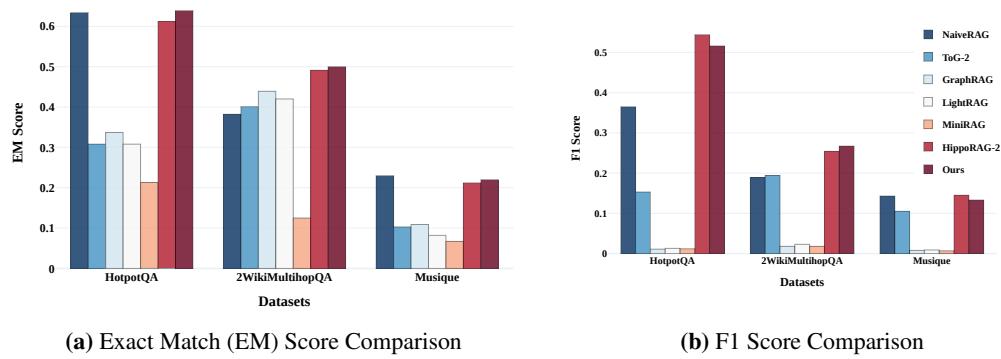
### 4.2 RESULT OF DEEP REASONGING BENCHMARK

293

**Result Analysis from a Method Perspective.** Results shown in Table 1 represent the average of three independent reasoning experiments. Previsou Graph-based methods like GraphRAG that rely on LLM-based graph construction show limited performance. Their performance is the lowest, particularly in terms of F1 scores as shown in Figure 4b, 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. F.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 (0.453) and F1 (0.312) 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.

324 **Table 1:** Exact Match (EM) and F1 scores on Deep Reasoning datasets. We highlight the **best**, **second-best**,  
 325 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	0.639	0.516	0.500	0.267	0.221	0.153	0.453	0.312

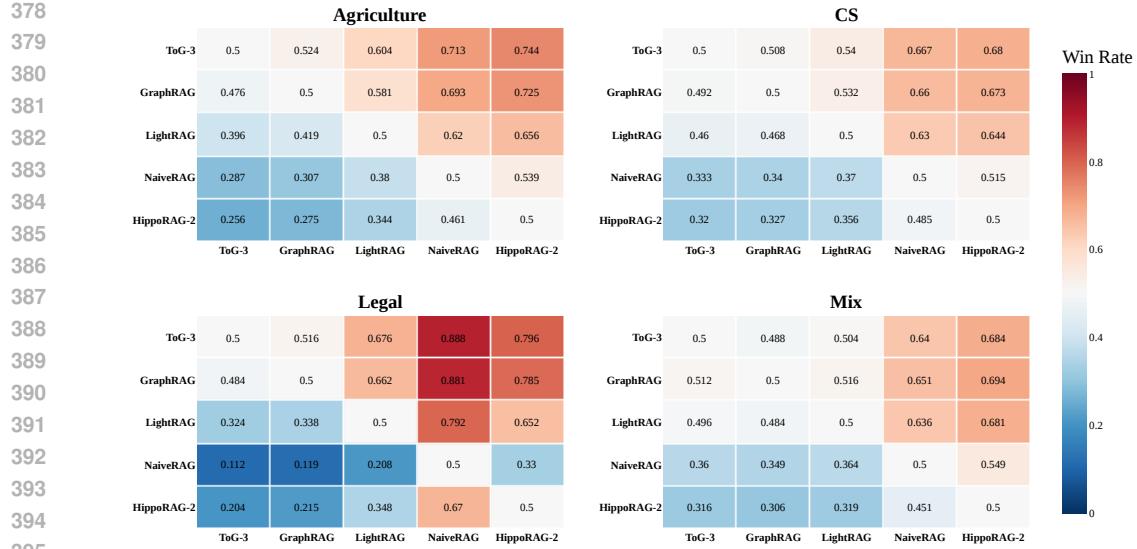


352 **Figure 4:** Performance comparison of different RAG methods on multi-hop QA datasets. (a) Exact Match  
 353 scores measure the percentage of questions where the model’s answer exactly matches the ground truth. (b) F1  
 354 scores provide a harmonic mean of precision and recall for token-level answer matching.

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

### 4.3 RESULT OF BROAD REASONING TASKS

369 As shown in Figure 5, The four heatmaps clearly demonstrate that the five methods can be distinctly  
 370 divided into two clusters: the upper-right region (predominantly red, indicating superior perfor-  
 371 mance) and the lower-left region (predominantly blue, indicating inferior performance). Specifically,  
 372 ToG-3, GraphRAG, and LightRAG exhibit significantly higher win rates compared to NaiveRAG  
 373 and HippoRAG-2. Detailed win rates (%) of baselines v.s. ToG-3 across four datasets are pro-  
 374 vided in Table 6 of Appendix. F. Our framework outperforms NaiveRAG by substantial margins  
 375 (up to 88.8% overall win rate on Legal dataset and 72.9% average win rate on all four datasets),  
 376 highlighting the limitations of chunk-based retrieval for complex queries. While GraphRAG shows  
 377 competitive performance in comprehensiveness due to its extensive community summarization and  
 378 retrieval, ToG-3 achieves better balance across all metrics, particularly excelling in diversity and em-  
 379 powerment through its heterogeneous graph architecture that integrates chunk-level, triplet-level,



**Figure 5: 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:** 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 Time (h)	Inference Time (s/q)	Avg. EM
		Entities	Relations	Communities			
HotpotQA	ToG-3	37,358	30,987	5,041	12.5	16.82	<b>0.639</b>
	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	14.95	<b>0.500</b>
	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	13.24	<b>0.221</b>
	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	15.00	<b>0.453</b>
	GraphRAG	83,658	64,748	8,883	13.10	8.58	0.295
	LightRAG	79,792	63,358	-	10.06	6.30	0.270

and community-level information. Detailed ELO rating calculation for broad reasoning tasks can be found in Appendix. F.3. The multi-agent dual-evolving context retrieval mechanism enables both deep knowledge reasoning through entity-relation exploration and broad community reasoning. This balanced architectural approach makes ToG-3 particularly effective for real-world applications requiring both comprehensive coverage and precise, actionable insights. Our analysis reveals that, on average, 20% of the samples require one evolving-context iteration, 32% require two iterations, and 48% require three iterations. Case studies of ToG-3 retrieval and response output are provided in Appendix. G.

#### 4.4 ANALYSIS OF COMPUTATION COST

The Table 2 reveal 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

432 **Table 3:** Ablation studies of MACER components and foundation model scaling. Standard ToG-3 settings  
 433 incorporates all MACER components, employs the Qwen2.5-32B-instruct as the backbone LLM, and utilizes  
 434 the Jina-v3-embedding model for representation encoding.

Ablation Setting	HotpotQA		2WikiMultihopQA		Musique		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
Standard ToG-3 settings	0.639	0.516	0.500	0.267	0.221	0.153	0.453	0.312
<b>MACER Components Ablation</b>								
w/o Evolving Query	0.598	0.443	0.412	0.203	0.178	0.121	0.396	0.256
w/o Evolving Sub-Graph	0.613	0.472	0.458	0.234	0.203	0.138	0.425	0.281
w/o Community Node	0.641	0.519	0.487	0.259	0.216	0.148	0.448	0.309
<b>Foundation Model Scaling Abalation</b>								
<b>LLM Model</b>								
Qwen2.5-14B	0.573	0.469	0.453	0.231	0.198	0.134	0.408	0.278
Qwen2.5-72B	0.668	0.538	0.523	0.281	0.235	0.162	0.475	0.327
<b>Embedding Model</b>								
Qwen3-Embed-0.6B	0.638	0.517	0.505	0.269	0.224	0.155	0.456	0.314
Qwen3-Embed-4B	0.643	0.523	0.508	0.271	0.227	0.158	0.459	0.317

451 constructs the graph more efficiently by extracting fewer relational structures during graph initial-  
 452 ization compared to both LightRAG and GraphRAG. While LightRAG achieve faster inference  
 453 times, they suffer from lower accuracy due to redundant graph elements or simpler retrieval mecha-  
 454 nisms. GraphRAG’s expensive two-stage indexing yields suboptimal results despite longer process-  
 455 ing times. ToG-3 demonstrates an effective balance: its efficient heterogeneous graph construction  
 456 produces refined knowledge bases across all datasets, and while its multi-agent reasoning requires  
 457 higher inference time, this cost is directly justified by its best performance on all benchmarks, mak-  
 458 ing it ideal for quality-sensitive applications requiring reliable reasoning capabilities. Note that  
 459 HippoRAG-2 was excluded from the comparative analysis due to its reliance on OpenIE-based ex-  
 460 traction rather than pure LLM extraction, which differs fundamentally from the approaches under  
 461 investigation.

## 463 4.5 ABALATION STUDY

464 **Abalation Study of MACER component** Our ablation study reveals the relative importance of  
 465 each MACER component for deep reasoning performance. The most significant performance degra-  
 466 dation occurs when removing the evolving query mechanism (average performance drop of 12.6% in  
 467 EM and 17.9% in F1), underscoring its critical role in complex question answering, especially when  
 468 using smaller LLMs. The iterative query decomposition enables the framework to break down mul-  
 469 tifaceted questions into tractable sub-problems, which is essential for navigating the heterogeneous  
 470 graph structure. Removing subgraph refinement causes a moderate performance decrease (average  
 471 drop of 6.2% in EM and 10.0% in F1), indicating its importance in adapting the knowledge structure  
 472 to the specific reasoning context. Interestingly, community nodes show the smallest impact on deep  
 473 reasoning tasks (a slight drop in the average EM and F1 scores), suggesting that while they con-  
 474 tribute to performance, the chunk and triplet representations carry most of the relevant information  
 475 for precise answer generation. However, in broad reasoning tasks, community nodes are essential  
 476 for comprehensive coverage and diversity, highlighting the complementary roles of different node  
 477 types in our heterogeneous graph architecture.

478 **Abalation Study of used foundation model** The foundation model scaling analysis reveals sev-  
 479 eral important patterns. First, LLM capacity has a substantially greater impact on performance than  
 480 embedding model size. Scaling from Qwen2.5-14B to Qwen2.5-72B yields a 16.4% average im-  
 481 provement in EM scores, highlighting the critical role of reasoning capability in complex QA tasks.  
 482 Second, larger embedding models provide consistent but more modest improvements. Qwen3-  
 483 Embed-0.6B shows a slight average EM improvement over jina-embeddings-v3, while Qwen3-  
 484 Embed-4B provides a 1.6% improvement. This suggests that while retrieval quality matters and  
 485 larger embedding models contribute to better performance, the LLM’s reasoning capacity remains

486 the primary bottleneck for complex reasoning tasks. These findings provide practical guidance for  
 487 resource allocation in real-world deployments.  
 488

489 **5 CONCLUSION**  
 490

491 In this work, we introduced Think-on-Graph 3.0, a novel framework that fundamentally rethinks  
 492 the paradigm of RAG for complex reasoning. By proposing the Multi-Agent Context Evolution  
 493 Retrieval (MACER) mechanism and a dynamic Chunk-Triplets-Community heterogeneous  
 494 graph architecture, we address critical limitations in both existing graph-based RAG methods and  
 495 knowledge-graph-dependent approaches. Our comprehensive experimental evaluation demonstrates  
 496 that ToG-3 achieves state-of-the-art performance across multiple challenging benchmarks. The  
 497 framework’s core innovation is its dual-evolving mechanism—comprising Evolving Query and  
 498 Evolving Subgraph—which dynamically refines both the query representation and the underlying  
 499 graph structure throughout the reasoning process. This iterative co-evolution enables deep, multi-  
 500 hop inference while preserving broad coverage across complex queries. This adaptive capability  
 501 proves particularly valuable for overcoming the quality constraints of static graph construction and  
 502 the domain limitations of pre-existing knowledge bases. The framework’s ability to work with light  
 503 LLMs also opens possibilities for more efficient and deployable AI systems. We believe the prin-  
 504 ciples established in ToG-3—dynamic graph evolution, multi-agent collaboration, and heterogeneous  
 505 knowledge integration—provide a foundation for the next generation of RAG systems.  
 506

507 **6 LIMITATION AND FUTURE DIRECTIONS**  
 508

509 Of course our work has several limitations. First, constrained by GPU resources, our experiments  
 510 are primarily conducted with LLMs up to 72B parameters and embedding models up to 4B pa-  
 511 rameters—though these sizes are practical for most developers and small-to-medium enterprises  
 512 for local deployment. Second, the evolving query and sub-graph refinement components increase  
 513 inference latency, typically 2–3x slower than baseline methods, making our approach more suit-  
 514 able for accuracy-critical applications where sacrificing speed for improved knowledge fidelity is  
 515 acceptable. Third, the same mechanisms result in longer context inputs, which demand larger GPU  
 516 memory capacity for efficient processing. These limitations could be mitigated through model dis-  
 517 tillation, optimized graph traversal algorithms, and dynamic context pruning techniques in future  
 518 improvement.

519 Future work will explore three promising directions for further advancement. First, we plan to  
 520 extend our multi-agent evolving framework to support larger-scale and more complex knowledge-  
 521 intensive tasks, such as programming assistance, financial analysis, and clinical decision-making.  
 522 Second, we aim to generalize our method from text to multimodal reasoning, integrating audio,  
 523 image, and video modalities to construct a world model that bridges textual knowledge with per-  
 524 ceptual grounding. Third, inspired by human cognitive science and brain science, we will explore  
 525 novel architectures that combine parametric memory with large language models to unify memory  
 526 and reasoning in a seamless framework, enabling more efficient knowledge retention and tool-use  
 527 capabilities akin to human intelligence.  
 528  
 529  
 530  
 531  
 532  
 533  
 534  
 535  
 536  
 537  
 538  
 539

540 REFERENCES  
541

542 Anthropic AI. Introducing claude 4. 2025a. URL <https://www.anthropic.com/news/clause-4>.

543 Meta AI. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. 2025b.  
544 URL <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>.

545 Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning  
546 to retrieve, generate, and critique through self-reflection, 2023. URL <https://arxiv.org/abs/2310.11511>.

547 Howard Chen, Ramakanth Pasunuru, Jason Weston, and Asli Celikyilmaz. Walking down the mem-  
548 ory maze: Beyond context limit through interactive reading, 2023. URL <https://arxiv.org/abs/2310.05029>.

549 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit  
550 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the  
551 frontier with advanced reasoning, multimodality, long context, and next generation agentic capa-  
552 bilities. *arXiv preprint arXiv:2507.06261*, 2025.

553 Benoit Dherin, Michael Munn, Hanna Mazzawi, Michael Wunder, and Javier Gonzalvo. Learning  
554 without training: The implicit dynamics of in-context learning. *arXiv preprint arXiv:2507.16003*,  
555 2025.

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

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

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

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

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

569 Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. Hipporag: Neuro-  
570 biologically inspired long-term memory for large language models. In *The Thirty-eighth Annual  
571 Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=hkujvAPVsg>.

572 Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. From rag to memory:  
573 Non-parametric continual learning for large language models, 2025. URL <https://arxiv.org/abs/2502.14802>.

574 Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing A multi-  
575 hop QA dataset for comprehensive evaluation of reasoning steps. In Donia Scott, Núria Bel,  
576 and Chengqing Zong (eds.), *Proceedings of the 28th International Conference on Compu-  
577 tational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pp. 6609–  
578 6625. International Committee on Computational Linguistics, 2020. doi: 10.18653/V1/2020.  
579 COLING-MAIN.580. URL <https://doi.org/10.18653/v1/2020.coling-main.580>.

580 Jianheng Huang, Leyang Cui, Ante Wang, Chengyi Yang, Xinting Liao, Linfeng Song, Junfeng Yao,  
581 and Jinsong Su. Mitigating catastrophic forgetting in large language models with self-synthesized  
582 rehearsal. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd  
583 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.  
584 1416–1428, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:  
585 10.18653/v1/2024.acl-long.77. URL <https://aclanthology.org/2024.acl-long.77/>.

594 Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, and  
 595 Xiang Ren. Lifelong pretraining: Continually adapting language models to emerging corpora. In  
 596 Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (eds.), *Proceedings*  
 597 *of the 2022 Conference of the North American Chapter of the Association for Computational*  
 598 *Linguistics: Human Language Technologies*, pp. 4764–4780, Seattle, United States, July 2022.  
 599 Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.351. URL <https://aclanthology.org/2022.naacl-main.351/>.

600

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

605

606 Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Li-  
 607 dong Bing. Chain-of-knowledge: Grounding large language models via dynamic knowledge  
 608 adapting over heterogeneous sources. In *International Conference on Learning Representations*,  
 609 2024. URL <https://openreview.net/forum?id=cPgh4gWZlz>.

610

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

614

615 Shengjie Ma, Chengjin Xu, Xuhui Jiang, Muzhi Li, Huaren Qu, Cehao Yang, Jiaxin Mao, and Jian  
 616 Guo. Think-on-graph 2.0: Deep and faithful large language model reasoning with knowledge-  
 617 guided retrieval augmented generation, 2024. URL <https://arxiv.org/abs/2407.10805>.

618

619 OpenAI. Introducing gpt-5. 2025. URL <https://openai.com/index/introducing-gpt-5/>.

620

621 Hongjin Qian, Zheng Liu, Peitian Zhang, Kelong Mao, Defu Lian, Zhicheng Dou, and Tiejun  
 622 Huang. Memorag: Boosting long context processing with global memory-enhanced retrieval  
 623 augmentation. In *Proceedings of the ACM Web Conference 2025 (TheWebConf 2025)*, Sydney,  
 624 Australia, 2025. ACM. URL <https://arxiv.org/abs/2409.05591>. arXiv:2409.05591.

625

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

629

630 Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Man-  
 631 ning. RAPTOR: recursive abstractive processing for tree-organized retrieval. In *The Twelfth*  
 632 *International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11,*  
 633 2024. OpenReview.net, 2024. URL <https://openreview.net/forum?id=GN921JHCRw>.

634

635 Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. Enhancing  
 636 retrieval-augmented large language models with iterative retrieval-generation synergy, 2023. URL  
 637 <https://arxiv.org/abs/2305.15294>.

638

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

642

643 Saba Sturua, Isabelle Mohr, Mohammad Kalim Akram, Michael Günther, Bo Wang, Markus Krim-  
 644 mel, Feng Wang, Georgios Mastrapas, Andreas Koukounas, Andreas Koukounas, Nan Wang,  
 645 and Han Xiao. jina-embeddings-v3: Multilingual embeddings with task lora, 2024. URL  
 646 <https://arxiv.org/abs/2409.10173>.

647

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

651

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

648 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. MuSiQue: Multi-  
 649 hop questions via single-hop question composition. *Transactions of the Association for Compu-  
 650 tational Linguistics*, 10:539–554, 2022. doi: 10.1162/tacl\_a\_00475. URL <https://aclanthology.org/2022.tacl-1.31/>.  
 651

652 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Interleaving re-  
 653 trieval with chain-of-thought reasoning for knowledge-intensive multi-step questions, 2023. URL  
 654 <https://arxiv.org/abs/2212.10509>.  
 655

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

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

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

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

673 Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen,  
 674 and Ningyu Zhang. Editing large language models: Problems, methods, and opportunities.  
 675 In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference  
 676 on Empirical Methods in Natural Language Processing*, pp. 10222–10240, Singapore, Decem-  
 677 ber 2023b. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.632.  
 678 URL <https://aclanthology.org/2023.emnlp-main.632/>.  
 679

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

683 Zihan Zhang, Meng Fang, Ling Chen, and Mohammad-Reza Namazi-Rad. CITB: A benchmark for  
 684 continual instruction tuning. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the  
 685 Association for Computational Linguistics: EMNLP 2023*, pp. 9443–9455, Singapore, December  
 686 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.633.  
 687 URL <https://aclanthology.org/2023.findings-emnlp.633/>.  
 688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 APPENDICES  
703704 Within this supplementary material, we elaborate on the following aspects:  
705

- 706 • Appendix A: Implementation Details and Hyperparameters
- 707 • Appendix B: Detailed ToG-3 Algorithms
- 708 • Appendix C: Datasets Statistics and Details
- 709 • Appendix D: Baselines Details
- 710 • Appendix E: Evaluation Metrics
- 711 • Appendix F: More Experiment Results and Details
- 712 • Appendix G: Case Study for ToG-3
- 713 • Appendix H: Graph Visualization Examples
- 714 • Appendix I: Theoretical Support for ToG-3
- 715 • Appendix J: LLM Prompts

719 A IMPLEMENTATION DETAILS  
720

721 We implement ToG-3 experiments with the following configuration: **Data Processing**: Chunk size  
722 is set to 1024 tokens with 20-token overlap between consecutive chunks to maintain contextual  
723 continuity. **Multi-Agent hyperparameter**: Constructor Agent extracts a maximum of 2 knowledge  
724 triplets per chunk and employs hierarchical Leiden clustering (Traag et al., 2019) with maximum  
725 cluster size of 5 for community detection. Retriever Agent retrieves top-5 most relevant nodes using  
726 hybrid vector-graph similarity matching. Reflector/Responser Agent utilizes the top-5 retrieved  
727 passages as context for answer generation. **Backend Infrastructure**: LLM service is based on  
728 Qwen2.5-32B-Instruct (Yang et al., 2024) deployed with vLLM (Kwon et al., 2023) engine using  
729 bfloat16 precision and prefix caching enabled and greedy-search generation method, which is more  
730 stable than the Qwen3 model in mixed reasoning mode in our task; embeddings are generated using  
731 Jina-embeddings-v3 (1024-dimensional) (Sturua et al., 2024); Our server is equipped with 8 A100  
732 40GB cards, AMD EPYC 256-core Processor, 2TB memory, and Ubuntu 20.04.1 system. and the  
733 hybrid vector-graph storage is implemented using Neo4j community edition<sup>1</sup> for efficient knowl-  
734 edge representation and retrieval, see Appendix.H for visualized graph example.

735 B TOG-3 ALGORITHMS  
736

737 Algorithms 1 and 2 present the two-stage pipeline of ToG-3. The first stage constructs a heteroge-  
738 neous graph index comprising chunks, triplets, and communities, while the second stage implements  
739 a Multi-Agent Context Evolution and Retrieval (MACER) loop featuring a novel dual-evolution  
740 mechanism—Evolving Query and Evolving Subgraph—that dynamically refines both the query rep-  
741 resentation and the graph structure through iterative interaction.

743 C DATASET DETAIL  
744

745 This section presents a comprehensive statistical overview of the **Deep and Broad datasets** we  
746 use in this paper, including detailed statistics metadata and licensing information, as summarized in  
747 Table 4. Additionally, we provide individual descriptions of each dataset to elucidate their respective  
748 characteristics and intended use cases.

750 C.1 DEEP REASONING DATASETS  
751

- 752 • **HotpotQA** (Yang et al., 2018): A crowdsourced question answering dataset built on En-  
753 glish Wikipedia, comprising approximately 113K questions. Each question is constructed  
754 to require the combination of information from the introductory sections of two Wikipedia

755 <sup>1</sup><https://neo4j.com/product/community-edition>

---

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

---

757 **Require:** Corpus  $\mathcal{D} = \{d_i\}_{i=1}^N$ , lightweight LM  $\mathcal{L}_{\text{light}}$ , encoder  $E_\theta$

758 **Ensure:** Heterogeneous graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

759 1:  $\mathcal{V} \leftarrow \emptyset, \mathcal{E} \leftarrow \emptyset$

760 2:  $\mathcal{C} \leftarrow \text{SplitIntoChunks}(\mathcal{D})$  ▷ Sentence-level segmentation

761 3:  $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{C}$

762 4: **for** each chunk  $c \in \mathcal{C}$  **do**

763 5:    $\mathcal{T}_c \leftarrow \mathcal{L}_{\text{light}}(c)$  ▷ Extract semantic triplets  $(s, p, o, \text{type}_s, \text{type}_p, \text{type}_o)$

764 6:    $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{T}_c$

765 7:   **for** each triplet  $t \in \mathcal{T}_c$  **do**

766 8:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{\text{MENTIONEDIN}(t, c)\}$

767 9:   **end for**

768 10: **end for**

769 11:  $G_e \leftarrow \text{BuildEntityCoOccurrenceGraph}(\mathcal{T})$  ▷  $\mathcal{T}$  is all triplets

770 12:  $\{M_\ell\}_\ell \leftarrow \text{LeidenClustering}(G_e)$

771 13: **for** each community  $M_\ell$  **do**

772 14:    $m_\ell \leftarrow \mathcal{L}_{\text{light}}(M_\ell)$  ▷ Generate community summary

773 15:    $\mathcal{V} \leftarrow \mathcal{V} \cup \{m_\ell\}$

774 16:   **for** each entity  $e \in M_\ell$  **do**

775 17:      $\mathcal{E} \leftarrow \mathcal{E} \cup \{\text{SUMMARYFOR}(m_\ell, e)\}$

776 18:   **end for**

777 19: **end for**

778 20: **Encode** every node  $v \in \mathcal{V}$  using  $E_\theta$  ▷ Unified dense encoding

779 21: **return**  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

---

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

---

781 **Require:** Query  $q$ , heterogeneous graph  $\mathcal{G}$ , LLM  $\mathcal{L}$ , max rounds  $K$

782 **Ensure:** Final answer  $a^*$

783 1:  $k \leftarrow 0, \mathcal{G}_0 \leftarrow \text{Retriever}(q, \mathcal{G})$  ▷ Initial retrieval

784 2:  $\mathcal{H}_0 \leftarrow \{(q, \mathcal{G}_0, \text{init})\}$  ▷ Initialize trajectory history

785 3: **repeat**

786 4:    $a_k \leftarrow \pi_{\text{resp}}(q, \mathcal{G}_k, \mathcal{H}_k)$  ▷ Response Agent generates answer

787 5:    $r_k \leftarrow \pi_{\text{ref}}^{\text{suff}}(q, \mathcal{G}_k, a_k)$  ▷ Reflector judges sufficiency

788 6:   **if**  $r_k = 1$  **then break**

789 7:   **end if**

790 8:    $q'_k \leftarrow \pi_{\text{ref}}^{\text{evolve}}(q, \mathcal{G}_k)$  ▷ Reflector evolves query

791 9:    $\mathcal{G}_{k+1} \leftarrow \pi_{\text{const}}^{\text{evolve}}(q'_k, \mathcal{G}_k)$  ▷ Constructor evolves subgraph

792 10:    $\mathcal{H}_{k+1} \leftarrow \mathcal{H}_k \cup \{(q'_k, a_k, r_k, \mathcal{G}_{k+1})\}$

793 11:    $k \leftarrow k + 1$

794 12: **until**  $k = K$

795 13:  $a^* \leftarrow \pi_{\text{resp}}^{\text{final}}(q, \mathcal{H}_k)$  ▷ Synthesize answer from full trajectory

796 14: **return**  $a^*$

---

797  
798  
799 articles for answering. The dataset provides two gold paragraphs per question, along with a  
800 list of sentences identified as supporting facts necessary to answer the question. HotpotQA  
801 includes various reasoning strategies such as bridge questions (involving missing entities),  
802 intersection questions (e.g., “what satisfies both property A and property B?”), and com-  
803 parison questions (comparing two entities through a common attribute). It is available in  
804 two settings: a *few-shot distractor setting* where models are provided with 10 paragraphs  
805 including the gold ones, and an *open-domain full-wiki setting* where models must retrieve  
806 relevant passages from the entire Wikipedia corpus given only the question.

807 • **2WikiMultihopQA** (Ho et al., 2020): A multi-hop question answering dataset that  
808 contains complex questions requiring reasoning over multiple Wikipedia paragraphs. Each  
809 question is designed to necessitate logical connections across different pieces of informa-  
tion to arrive at the correct answer.

810 **Table 4:** Statistics of Deep Reasoning and Broad Reasoning Datasets. Metrics abbreviations: Comp. (Com-  
 811 prehensiveness), Div. (Diversity), Emp. (Empowerment).

813 Dataset	813 Corpus Size	813 Chunks	813 Entities/Relations	813 Communities	813 Metrics	813 License
<b>Deep Reasoning Tasks</b>						
815 HotpotQA	9,809	9,812	37,358/30,987	5,041	EM, F1	Apache-2.0
816 2WikiMultihopQA	6,119	6122	19,311/21,077	3,417	EM, F1	Apache-2.0
817 Musique	11,254	11,300	32,842/39,134	6,258	EM, F1	CC-BY-4.0
<b>Broad Reasoning Tasks</b>						
818 CS	10	2,134	3,530/33,507	1,166		
819 Agriculture	12	2,025	6,043/12,571	1,039	Comp., Div., Emp.	Apache-2.0
820 Legal	94	5,900	26,180/44,334	1,359		
Mix	61	658	2,784/5,089	425		

821

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

828

## 829 C.2 BROAD REASONING DATASETS

830

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

835

836 • **CS**: Computer science domain focusing on data science, software engineering, and pro-  
 837 gramming topics, requiring technical comprehension and analytical reasoning.

838 • **Agriculture**: Covers agricultural practices including beekeeping, crop production, and dis-  
 839 ease prevention, demanding domain-specific knowledge integration.

840 • **Legal**: Derived from legal contracts and documents, focusing on corporate legal practices,  
 841 regulatory compliance, and governance, requiring precise interpretation of nuanced legal  
 842 language.

843 • **Mix**: Contains diverse contexts from college textbooks spanning natural sciences, humani-  
 844 ties, and social sciences, testing generalization capabilities across interdisciplinary topics.

845

## 846 D BASELINES

847

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

851

852 • **NaiveRAG** (Gao et al., 2023): A standard chunk-based retrieval baseline that segments raw  
 853 texts into chunks and stores them in a vector database using text embeddings. For queries,  
 854 it generates vectorized representations to directly retrieve text chunks based on semantic  
 855 similarity.

856 • **GraphRAG** (Edge et al., 2024): A graph-enhanced RAG system that utilizes an LLM  
 857 to extract entities and relationships from text, representing them as nodes and edges. It  
 858 generates community summaries through graph clustering and employs both local (entity-  
 859 based) and global (community-based) retrieval strategies for comprehensive information  
 860 access.

861 • **LightRAG** (Guo et al., 2024): A graph-structured RAG framework that employs a dual-  
 862 level retrieval system combining low-level entity retrieval with high-level knowledge dis-  
 863 covery. It integrates graph structures with vector representations for efficient retrieval of  
 864 related entities and their relationships.

- 864 • **ToG-2** (Ma et al., 2024): A knowledge graph-based framework implements a tight-  
865 coupling hybrid RAG paradigm that iteratively retrieves information from both unstruc-  
866 tured texts and structured knowledge sources. It alternates between graph retrieval and  
867 context retrieval for in-depth knowledge exploration.
- 868 • **HippoRAG-2** (Gutiérrez et al., 2025): A non-parametric continual learning framework that  
869 leverages Personalized PageRank algorithm over an open knowledge graph constructed us-  
870 ing LLM-extracted triples. It enhances multi-hop reasoning capabilities through sophisti-  
871 cated graph traversal and passage integration mechanisms.

## 873 E METRICS

875 We employ different evaluation protocols for the two task categories:

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

- 879 • **Exact Match (EM)**: Measures the percentage of predictions that exactly match the ground  
880 truth answer. Specifically, we use Sub-string Exact Match version (Sun et al., 2023; Ma  
881 et al., 2024) which goes through the whole response to check whether the answer is in.
- 882 • **F1 Score**: Computes word-level overlap between predictions and ground truth answers.

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

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

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

## 895 F MORE EXPERIMENT RESULTS AND DETAILS

897 This section presents extended experimental results, including detailed precision and recall metrics  
898 on Deep Reasoning tasks, as well as one-to-one win rates from Broad Reasoning tasks. The pairwise  
899 win rates are converted into a unified ELO rating system, with the resulting ratings visualized in the  
900 heatmap shown in Figure 5.

### 902 F.1 PRECISION AND RECALL RATE RESULTS

904 Table 5 reveals the underlying reason for the relatively low F1 scores of GraphRAG and LightRAG:  
905 these methods are not specifically designed for deep reasoning tasks. By examining both preci-  
906 sion/recall metrics and output cases, we observe that excessively long or unfocused responses tend  
907 to substantially reduce recall, thereby diminishing overall F1 performance.

### 908 F.2 RESULT DETAIL IN BRAOD REASONING TASKS

910 Table 6 presents the pairwise win rates (%) of baseline methods against ToG-3 across four datasets  
911 and four evaluation dimensions. The results demonstrate that ToG-3 consistently outperforms all  
912 compared baselines.

### 914 F.3 ELO RATING CALCULATION FOR BROAD REASONING TASKS

916 This appendix details the mathematical framework and computational process for deriving ELO  
917 ratings from pairwise comparison data across four benchmark datasets. The ELO rating system pro-  
918 vides a mathematically consistent approach to quantify relative performance differences between

918 **Table 5:** Comprehensive Evaluation Metrics of five RAG methods across three deep reasoning datasets. The  
 919 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	<b>0.516</b>	<b>0.595</b>	<b>0.454</b>	<b>0.267</b>	<b>0.485</b>	<b>0.312</b>	<b>0.153</b>	<b>0.286</b>	<b>0.132</b>

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

930 **Table 6:** Win rates (%) of baselines v.s. ToG-3 across four datasets and four evaluation dimensions. The better  
 931 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	28.4%	<b>71.6%</b>	32.4%	<b>67.6%</b>	12.4%	<b>87.6%</b>	34.8%	<b>65.2%</b>
	19.6%	<b>80.4%</b>	32.0%	<b>68.0%</b>	9.6%	<b>90.4%</b>	28.4%	<b>71.6%</b>
	29.4%	<b>70.6%</b>	32.8%	<b>67.2%</b>	12.4%	<b>87.6%</b>	38.8%	<b>61.2%</b>
	28.7%	<b>71.3%</b>	33.3%	<b>66.7%</b>	11.2%	<b>88.8%</b>	36.0%	<b>64.0%</b>
Diversity	GraphRAG	ToG-3	GraphRAG	ToG-3	GraphRAG	ToG-3	GraphRAG	ToG-3
	46.8%	<b>53.2%</b>	49.6%	<b>50.4%</b>	49.6%	<b>50.4%</b>	<b>51.6%</b>	48.4%
	44.4%	<b>55.6%</b>	48.4%	<b>51.6%</b>	46.8%	<b>53.2%</b>	<b>52.0%</b>	48.0%
	25.2%	<b>74.8%</b>	43.2%	<b>56.8%</b>	29.6%	<b>70.4%</b>	38.4%	<b>61.6%</b>
Empowerment	Overall	47.6%	<b>52.4%</b>	49.2%	<b>50.8%</b>	48.4%	<b>51.6%</b>	48.8%
	GraphRAG	ToG-3	GraphRAG	ToG-3	GraphRAG	ToG-3	GraphRAG	ToG-3
	38.9%	<b>61.1%</b>	45.6%	<b>54.4%</b>	33.6%	<b>66.4%</b>	47.6%	<b>52.4%</b>
	Diversity	32.0%	<b>68.0%</b>	42.0%	<b>58.0%</b>	28.0%	<b>72.0%</b>	39.2%
Overall	Empowerment	40.5%	<b>59.5%</b>	46.0%	<b>54.0%</b>	33.6%	<b>66.4%</b>	52.0%
	Comprehensiveness	39.6%	<b>60.4%</b>	46.0%	<b>54.0%</b>	32.4%	<b>67.6%</b>	49.6%
	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3
	24.5%	<b>75.5%</b>	31.6%	<b>68.4%</b>	21.6%	<b>78.4%</b>	29.6%	<b>70.4%</b>
Comprehensiveness	Diversity	18.8%	<b>81.2%</b>	28.0%	<b>72.0%</b>	17.2%	<b>82.8%</b>	23.6%
	Empowerment	27.8%	<b>72.2%</b>	32.9%	<b>67.1%</b>	21.6%	<b>78.4%</b>	34.0%
	Overall	25.6%	<b>74.4%</b>	32.0%	<b>68.0%</b>	20.4%	<b>79.6%</b>	31.6%
	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3	HippoRAG-2	ToG-3

951 retrieval-augmented generation methods. The ELO rating system transforms raw win rates into a  
 952 logarithmic scale that ensures transitive consistency in performance rankings. The core transformation  
 953 is defined as follows:

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

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

956 The absolute ELO rating for method  $i$  is then:

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

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

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

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

## G CASE STUDY FOR TOG-3

959 This section provides a detailed case study of ToG-3 in deep reasoning task (Figure 6) and broad  
 960 reasoning task (Figure 7 and Figure 8), offering an intuitive demonstration of the execution dynam-

```

972
973 Question:
974   What nationality is the performer of the song When The Stars Go Blue?
975 Initial Evidence (Sub-Graph):
976   When The Stars Go Blue -> performed_by -> Ryan Adams
977   Ryan Adams -> known_for -> alternative country, rock, folk
978   Ryan Adams -> also_includes_genre -> indie rock, Americana
979   ... (no nationality information present)
980 Evolving Query Decomposition 1 (Reflector Agent)
981   Who is the performer of the song When The Stars Go Blue?
982   Ryan Adams.
983 Evolving Sub-Graph Refinement (Constructor Agent):
984   Sammy Hagar -> is -> American rock vocalist
985   Chuck Berry -> is -> American singer and songwriter
986   Elvis Presley -> pioneered -> Rock and Roll (American origin)
987   Ryan Adams -> nationality -> American [Added]
988   Wee Willie Walker -> was -> American soul and blues singer
989 Evolving Query Decomposition 2 (Reflector Agent)
990   What is Ryan Adams known for besides the song When The Stars Go Blue?
991   Ryan Adams is known for his work as a singer-songwriter and musician, with a musical style that
992   encompasses alternative country, rock, and folk.
993 Evolving Query Decomposition 3 (Reflector Agent)
994   What is Ryan Adams's nationality?
995   Ryan Adams is an American singer and musician.
996 Final Answer (Response Agent):
997   ryan adams is american.
998 Ground Truth:
999   america
1000 Explanation:
1001   The initial sub-graph contained no explicit information about Ryan Adams' nationality.
1002   Through evolving query decomposition, the system identified a missing link and triggered sub-graph
1003   refinement.
1004   New evidence from biographical triples (e.g., "Ryan Adams -> nationality -> American") was
1005   integrated, enabling a definitive answer.
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

**Figure 6:** Case Study: Evolving Query and Evolving Sub-Graph for Multi-Hop Reasoning. This example demonstrates how an initial knowledge sub-graph fails to answer a complex question, prompting a decomposition into a sub-query. The Constructor Agent refines the sub-graph 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.

ics of its dual-evolution mechanism—comprising Evolving Query and Evolving Subgraph—across multi-step reasoning processes.

## H 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 <sup>2</sup>.

**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 graphs emphasize precision in step-by-step reasoning over

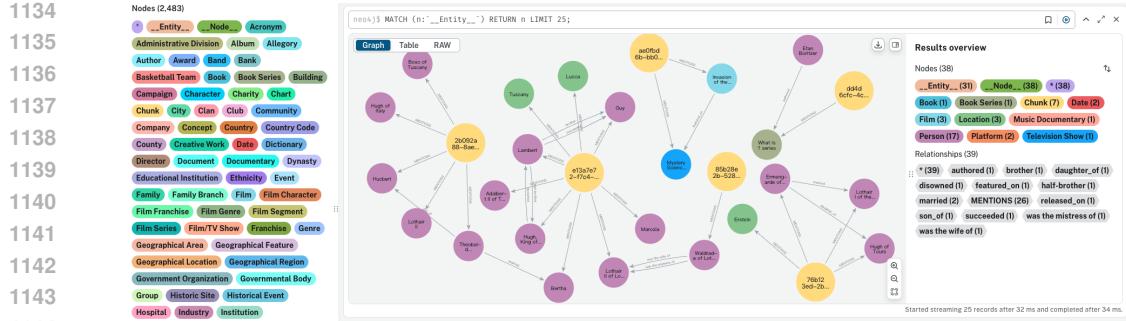
<sup>2</sup><https://neo4j.com/product/community-edition>

```

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035 Question:
1036 How do different regression metrics (e.g., RMSE, MAE,  $R^2$ ) vary by model and dataset?
1037 Evidence (from ToG-3):
1038 Mean Squared Error -> is_related_to -> Root Mean Squared Error...
1039 Mean Squared Error -> is_similar_to -> Mean Absolute Error...
1040 summary_23 -> SUMMARY_FOR -> RMSE.MSE emphasizes larger errors...
1041 summary_287 -> SUMMARY_FOR -> MAE. ARIMA model performance evaluated using MAE...
1042 Answer 1 (from ToG-3):
1043 Different regression metrics such as RMSE, MAE, and  $R^2$  can vary by model and dataset...
1044 RMSE penalizes large errors more heavily due to squared terms, while MAE is robust to outliers...
1045  $R^2$  measures explained variance but can be misleading on small or overfitted datasets...
1046 Answer 2 (from GraphRAG):
1047 The provided community summary does not detail specific variations...To evaluate how these
1048 metrics vary:
1. Train different regression models on your datasets...
1049 2. Compute RMSE, MAE,  $R^2$  using MLlib.RegressionMetrics...
1050 3. Compare performance across models and datasets...
1051
1052 Evaluation Result (JSON):
1053 {
1054     "Comprehensiveness": {
1055         "Winner": "Answer 1",
1056         "Explanation": "Answer 1 provides a detailed explanation of various regression metrics..."
1057     },
1058     "Diversity": {
1059         "Winner": "Answer 2",
1060         "Explanation": "Answer 2 presents a different perspective by incorporating specific tools..."
1061     },
1062     "Empowerment": {
1063         "Winner": "Answer 1",
1064         "Explanation": "Answer 1 empowers the reader by explaining what each metric means..."
1065     },
1066     "Overall Winner": {
1067         "Winner": "Answer 1",
1068         "Explanation": "Answer 1 is the overall winner as it provides a comprehensive understanding..."
1069     }
1070 }
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

1080  
 1081  
 1082  
 1083  
 1084 **Question:**  
 1085 What policy changes are needed to foster equitable access to nutritious food?  
 1086 **Evidence (from ToG-3):**  
 1087 summary\_216 -> SUMMARY\_FOR -> Food Security. \*\*Supermarket and Food Security\*\*: The reliance on  
 1088 supermarkets for food is viewed as a threat...  
 1089 Victory Garden -> promoted -> Food Security... Supermarket -> threatens -> Food Security..."  
 1090 White House Childhood Obesity Task Force -> recommended -> Farm to School...  
 1091 Food Works -> collaborates\_with -> Sauvie Island Organics...  
 1092 Michael Taylor -> interprets -> Food Modernization Act... background in representing Monsanto...  
 1093  
 1094 **Answer 1(from ToG-3):**  
 1095 policy changes that could foster equitable access to nutritious food may include implementing  
 1096 subsidies...  
 1097 increasing funding for farm to school programs in underserved communities...  
 1098 establishing regulations that prioritize the availability of nutritious options in school meal  
 1099 programs...  
 1100 and creating educational initiatives to promote awareness about nutrition and healthy eating  
 1101 habits...  
 1102 these changes could help address disparities in access to nutritious food...  
 1103  
 1104 **Answer 2(from GraphRAG):**  
 1105 1. \*\*strengthen conflict of interest policies\*\*: prevent individuals with significant industry  
 1106 ties...  
 1107 2. \*\*enhance transparency and accountability\*\*: implement transparent processes...  
 1108 3. \*\*support diverse food producers\*\*: promote a variety of food producers...  
 1109 4. \*\*expand nutrition education and assistance programs\*\*: enhance nutrition education...  
 1110 5. \*\*ensure regulatory independence\*\*: guarantee that regulatory bodies are independent...  
 1111 6. \*\*promote local food systems\*\*: encourage local food production through grants, subsidies...  
 1112 7. \*\*support community gardens and urban agriculture\*\*: provide funding and resources...  
 1113 8. \*\*promote sustainable farming practices\*\*: offer incentives for sustainable farming...  
 1114 These measures aim to create a more resilient and sustainable food system...  
 1115  
 1116 **Evaluation Result (JSON):**  
 1117 {  
 1118 {  
 1119 "Comprehensiveness": {  
 1120 "Winner": "Answer 2",  
 1121 "Explanation": "Answer 2 provides a more detailed and structured set of policy changes... "  
 1122 "broader scope than Answer 1..."  
 1123 },  
 1124 "Diversity": {  
 1125 "Winner": "Answer 2",  
 1126 "Explanation": "Answer 2 presents a wider variety of perspectives... "  
 1127 "including regulatory, community, and economic solutions."  
 1128 },  
 1129 "Empowerment": {  
 1130 "Winner": "Answer 1",  
 1131 "Explanation": "Answer 1 empowers readers by offering a comprehensive overview..."  
 1132 "facilitates informed judgment."  
 1133 },  
 1134 "Overall Winner": {  
 1135 "Winner": "Answer 2",  
 1136 "Explanation": "Answer 2 emerges as the overall winner due to its superior"  
 1137 "comprehensiveness, diversity, and empowerment."  
 1138 }  
 1139 }  
 1140  
 1141  
 1142 **Figure 8:** Case Study: Policy Recommendations for Equitable Food Access. This example illustrates the full  
 1143 reasoning pipeline: a complex policy question is answered by two different systems (GraphRAG and ToG-  
 1144 3), supported by retrieved knowledge snippets. An evaluator then compares both responses across multiple  
 1145 dimensions, selecting the more comprehensive, diverse, and empowering answer as the winner.  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153



1188 for a transformer module with MLP layer weights  $W$ , the context  $\mathcal{H}_k$  generates an implicit weight  
 1189 update  $\Delta W_k$  through the attention mechanism:  
 1190

$$1191 \quad \Delta W_k = \frac{(W \Delta A_k) A(q)^\top}{\|A(q)\|^2}, \quad \text{where } \Delta A_k = A(\mathcal{H}_k, q) - A(q). \quad (7)$$

1195 Here,  $A(\cdot)$  denotes the activation pattern from the attention layer,  $A(q)$  represents the baseline activation  
 1196 without context, and  $A(\mathcal{H}_k, q)$  captures the contextualized activation with the full reasoning  
 1197 history. The term  $\Delta A_k$  quantifies the information injected by the evolving context  $\mathcal{H}_k$ . The low-  
 1198 rank nature of  $\Delta W_k$  ensures efficient and targeted parameter updates without catastrophic forgetting  
 1199 of pre-trained knowledge.

1200 **MDP Policy as an Implicit Function of Context** Recall from Section 3.3 that the Reflector  
 1201 Agent’s policy  $\pi_{\text{ref}}$  maps states  $s_k = (q, \mathcal{G}_k, \mathcal{H}_k)$  to actions (sub-queries or STOP). Under the im-  
 1202 plicit learning view,  $\pi_{\text{ref}}$  is not a fixed network but an emergent policy  $\pi_k$  shaped by  $\Delta W_k$ . Thus,  
 1203 the sequence  $\{\pi_k\}_{k=1}^K$  constitutes a trajectory of implicitly adapted policies driven by the evolving  
 1204 context  $\mathcal{H}_k$ .

1206 **Convergence via Regret Minimization** We analyze convergence through the lens of episodic  
 1207 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  
 1208 value of policy  $\pi$  at state  $s_k$ , and let  $V_{s_k}^* = \max_\pi V_{s_k}^\pi$  be the optimal value. The cumulative regret  
 1209 over  $K$  steps is:

$$1211 \quad \mathcal{R}(K) = \sum_{k=1}^K (V_{s_k}^* - V_{s_k}^{\pi_k}). \quad (8)$$

1214 We establish sublinear regret growth  $\mathcal{R}(K) = o(K)$  under the following mild assumptions:

1215 **Assumption 1 (Realizability).** *There exists a policy  $\pi^*$  such that  $\text{Suff}(q, \mathcal{G}_q^*) = 1$ , and  $\pi^*$  is repre-  
 1216 sentable by the implicit policy class induced by in-context prompts of the form  $(\mathcal{H}; q)$ .*

1217 **Assumption 2 (Bounded Gradient Norm).** *The implicit gradient direction  $g_k$ , defined as the reward-  
 1218 sensitive update signal from  $\mathcal{H}_k$ , satisfies  $\|g_k\| \leq G$  for some constant  $G > 0$ .*

1220 Under these assumptions, the following properties hold:

1221 **Property 1 (Smooth Policy Evolution).** The value function evolves smoothly with respect to im-  
 1222 plicit updates:

$$1224 \quad \|V^{\pi_{k+1}} - V^{\pi_k}\|_\infty \leq L\|g_k\| + \mathcal{O}(\|g_k\|^2), \quad (9)$$

1225 for some Lipschitz constant  $L > 0$ , ensuring stable policy transitions.

1226 **Property 2 (Expected Policy Improvement).** Each refinement step yields non-negative expected  
 1227 improvement:

$$1229 \quad \mathbb{E} [V_{s_k}^{\pi_{k+1}} - V_{s_k}^{\pi_k} \mid \mathcal{H}_k] \geq \eta\|g_k\|^2 - \sigma_k, \quad (10)$$

1230 where  $\eta > 0$  and  $\{\sigma_k\}$  is a martingale difference sequence with  $\mathbb{E}[\sigma_k \mid \mathcal{H}_k] = 0$ . This follows  
 1231 from the fact that evolving sub-queries generated by the Reflector target knowledge gaps, and the  
 1232 Constructor’s evolving graph refinement increases the likelihood of sufficiency.

1233 **Property 3 (Vanishing Implicit Gradient).** As the context becomes increasingly informative, the  
 1234 room for improvement diminishes:

$$1236 \quad \lim_{k \rightarrow \infty} \|g_k\| = 0 \quad \text{almost surely.} \quad (11)$$

1237 This is guaranteed by Assumption 1 (Realizability) and the finite horizon  $K$ , which ensures the  
 1238 process either reaches a sufficient subgraph ( $r_k = 1$ ) or exhausts its budget.

1239 Together, these properties imply that the sequence  $\{\pi_k\}$  converges to a policy  $\pi^\dagger$  satisfying  $V_{s_1}^{\pi^\dagger} \geq$   
 1240  $V_{s_1}^* - \epsilon$  for arbitrarily small  $\epsilon > 0$  as  $K \rightarrow \infty$ . In practice, with a reasonable horizon (e.g.,  $K = 3$ ),  
 1241 MACER reliably converges to a sufficient context  $\mathcal{G}_q^*$  for faithful answer synthesis.

1242 This analysis establishes that the MACER loop performs an implicit form of policy gradient as-  
 1243 cent on the reward landscape defined by context sufficiency, with convergence guarantees rooted in  
 1244 stochastic approximation theory and in-context learning dynamics, providing rigorous foundations  
 1245 for the empirical effectiveness of our reward-based evolving context mechanism.

## 1247 J PROMPT TEMPLATES

1249 Our framework employs a multi-stage, prompt-driven reasoning pipeline that integrates structured  
 1250 knowledge graph (KG) extraction, community-based summarization, iterative sub-query decom-  
 1251 position, sub-graph refinement, and faithful answer synthesis. Each stage is governed by a specialized  
 1252 prompt template designed to ensure modularity, interpretability, and factual consistency. The com-  
 1253 plete sequence of prompts is as follows:

- 1255 **1. KG Triplets Extraction:** As shown in Figure 11, given raw textual input, this prompt  
 1256 instructs the model to extract structured subject-relation-object triples (e.g., entity1 ->  
 1257 relation -> entity2) to construct a fine-grained knowledge sub-graph. This step trans-  
 1258 forms unstructured text into a queryable graph structure.
- 1259 **2. Generate Community Summary:** As shown in Figure 12, based on densely connected  
 1260 sub-graphs (communities), this prompt synthesizes a concise natural language summary  
 1261 that captures the core themes and relationships within each community, enabling high-level  
 1262 semantic indexing and retrieval.
- 1263 **3. Keyword Expansion for Retrieval Augmentation:** As shown in Figure 13, to improve  
 1264 recall in the querying phase, this prompt generates a set of synonyms and related terms  
 1265 from the original query, considering variations in capitalization, pluralization, and common  
 1266 phrasings, separated by delimiter symbols.
- 1267 **4. Evolving Sub-Query Decomposition:** As shown in Figure 14, for complex multi-hop  
 1268 questions, this prompt recursively decomposes the current query into simpler, context-  
 1269 answerable sub-questions, guided by previously retrieved information and reasoning traces,  
 1270 enabling stepwise information gathering.
- 1271 **5. Evolving Sub-Graph Refinement:** As shown in Figure 15, this prompt cleans and en-  
 1272 hances the retrieved or extracted sub-graph by removing irrelevant triples, normalizing  
 1273 entity names, and optionally filling in strongly supported missing links, thereby improving  
 1274 the signal-to-noise ratio for downstream reasoning.
- 1275 **6. Final Answer Synthesis:** As shown in Figure 16, in the final stage, the model generates a  
 1276 concise, context-grounded answer using *only* the refined evidence, with explicit instructions  
 1277 to avoid hallucination or reliance on prior knowledge. If the answer cannot be determined,  
 1278 it returns “Unknown” to maintain factual integrity.

1279 These prompts work in concert to enable structured, interpretable, and reliable reasoning over hybrid  
 1280 text-and-graph knowledge sources. And Figure 17 shows the LLM evaluation prompt in the broad  
 1281 reasoning task. Their modular design allows for independent tuning and auditing, making the overall  
 1282 system transparent and robust to noise and ambiguity.

1283  
 1284  
 1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295

```

1296
1297 -Goal-
1298 Given a text document, identify all entities and their entity types from the text and all
1299 relationships among the identified entities.
1300 Given the text, extract up to {max_knowledge_triplets} entity-relation triplets.
1301
1302 -Steps-
1303 1. Identify all entities. For each, extract:
1304 entity_name | entity_type | entity_description
1305
1306 2. Identify all related (source, target) pairs. For each, extract:
1307 source_entity | target_entity | relation | relationship_description
1308
1309 3. Output valid JSON only:
1310 { "entities": [...], "relationships": [...] }
1311
1312 -An Output Example-
1313 {
1314   "entities": [
1315     { "entity_name": "Albert Einstein", "entity_type": "Person", "entity_description": "..." },
1316     { "entity_name": "Theory of Relativity", "entity_type": "Scientific Theory",
1317       "entity_description": "..." },
1318     { "entity_name": "Nobel Prize in Physics", "entity_type": "Award", "entity_description": "..." }
1319   ],
1320   "relationships": [
1321     { "source_entity": "Albert Einstein", "target_entity": "Theory of Relativity", "relation": "developed",
1322       "relationship_description": "..." },
1323     { "source_entity": "Albert Einstein", "target_entity": "Nobel Prize in Physics", "relation": "won",
1324       "relationship_description": "..." }
1325   ]
1326 }
1327
1328 -Real Data-
1329 #####
1330 text: {text}
1331 #####
1332 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.

```

1333
1334 role="system"
1335 You are provided with a set of relationships from a knowledge graph, each represented as
1336 entity1 -> entity2 -> relation -> relationship_description.
1337 Your task is to create a summary of these relationships. The summary should include: Names
1338 of the entities involved, A concise synthesis of the relationship descriptions. The goal is
1339 to capture the most critical and relevant details that highlight the nature
1340 and significance of each relationship. Ensure the summary is coherent and integrates
1341 information to emphasize key aspects. Avoid redundancy and maintain clarity.
1342
1343 role="user"
1344 #####
1345 text: {community.info}
1346 #####
1347
1348 assistant:
1349 % 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.

```

1350
1351
1352
1353 role="system"
1354 Given some initial query, generate synonyms or related keywords up to {max_keywords} in
1355 total,
1356 considering possible cases of capitalization, pluralization, common expressions, etc.
1357 Provide all synonyms/keywords separated by '^' symbols: 'keyword1^keyword2^...'.
1358 Note: result should be in one line, separated by '^' symbols.
1359
1360 role="user"
1361 -----
1362 assistant:
1363 % Example: KEYWORDS: machine learning^ML learning machines^AI models^neural networks^deep
1364 learning ...
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375 role="system"
1376 The original question is as follows: {query_str}
1377 We have an opportunity to answer some, or all of the question from a knowledge source.
1378 Context information for the knowledge source is provided below, as well as previous reasoning
1379 steps.
1380 Given the context and previous reasoning, return a question that can be answered from the
1381 context.
1382 This question can be the same as the original question, or represent a subcomponent.
1383 It should not be irrelevant to the original question.
1384 If no further information can be extracted, return 'None'.
1385
1386 Examples:
1387 Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
1388 Knowledge source context: Provides names of the winners of the 2020 Australian Open
1389 Previous reasoning: None
1390 Next question: Who was the winner of the 2020 Australian Open?
1391 Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
1392 Knowledge source context: Includes biographical info for each winner
1393 Previous reasoning:
1394 - Who was the winner of the 2020 Australian Open?
1395 - The winner was Novak Djokovic.
1396 Next question: How many Grand Slam titles does Novak Djokovic have?
1397
1398 Current Input:
1399 Question: {query_str}
1400 Knowledge source context: {context_str}
1401 Previous reasoning: {prev_reasoning}
1402
1403 assistant:
1404 % Output: <decomposed sub-question> OR 'None'

```

**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.

```

1374
1375 role="system"
1376 The original question is as follows: {query_str}
1377 We have an opportunity to answer some, or all of the question from a knowledge source.
1378 Context information for the knowledge source is provided below, as well as previous reasoning
1379 steps.
1380 Given the context and previous reasoning, return a question that can be answered from the
1381 context.
1382 This question can be the same as the original question, or represent a subcomponent.
1383 It should not be irrelevant to the original question.
1384 If no further information can be extracted, return 'None'.
1385
1386 Examples:
1387 Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
1388 Knowledge source context: Provides names of the winners of the 2020 Australian Open
1389 Previous reasoning: None
1390 Next question: Who was the winner of the 2020 Australian Open?
1391 Question: How many Grand Slam titles does the winner of the 2020 Australian Open have?
1392 Knowledge source context: Includes biographical info for each winner
1393 Previous reasoning:
1394 - Who was the winner of the 2020 Australian Open?
1395 - The winner was Novak Djokovic.
1396 Next question: How many Grand Slam titles does Novak Djokovic have?
1397
1398 Current Input:
1399 Question: {query_str}
1400 Knowledge source context: {context_str}
1401 Previous reasoning: {prev_reasoning}
1402
1403 assistant:
1404 % 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.

```

1404
1405
1406 role="system"
1407 You are given a sub-graph extracted from a knowledge graph, represented as a list of triples:
1408 entity1 -> relation -> entity2.
1409 This sub-graph may contain irrelevant, redundant, or incomplete information.
1410 Your task is to refine the sub-graph by:
1411 Removing irrelevant or noisy triples not related to the query, Filling in missing but inferable
1412 relationships (if strongly supported),
1413 Ensuring entity names are normalized (e.g., consistent capitalization, singular/plural).
1414 Return the refined sub-graph in the same triple format, one per line.
1415 If no refinement is needed, return the original sub-graph.
1416 If all triples are irrelevant, return 'None'.
1417
1418 Example Input:
1419 Query: What are the major achievements of Marie Curie?
1420 Sub-graph:
1421 Marie Curie -> won -> Nobel Prize in Physics
1422 Marie Curie -> born in -> Warsaw
1423 Marie Curie -> spouse -> Pierre Curie
1424 Apple Inc. -> founded by -> Steve Jobs
1425
1426 Refined Output:
1427 Marie Curie -> won -> Nobel Prize in Physics
1428 Marie Curie -> won -> Nobel Prize in Chemistry
1429 Marie Curie -> spouse -> Pierre Curie
1430 (Note: Added Chemistry prize based on strong prior knowledge; removed birthplace and unrelated
1431 Apple fact)
1432
1433 Current Input:
1434 Query: {query_str}
1435 Sub-graph:
1436 {subgraph_triples}
1437
1438 assistant:
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

```

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

```

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.

```

1458
1459
1460
1461
1462
1463
1464
1465
1466 role="system"
1467 You are an expert tasked with evaluating two answers to the same question
1468 based on three criteria: Comprehensiveness, Diversity, and Empowerment.
1469
1470 Evaluation Criteria:
1471
1472   • Comprehensiveness:  
How much detail does the answer provide to cover all aspects
1473 and sub-questions implied by the original query?
1474
1475   • Diversity:  
How varied and rich is the answer in providing different perspectives,
1476 evidence sources, or reasoning paths?
1477
1478   • Empowerment:  
How well does the answer help the reader understand the topic
1479 and make informed judgments or decisions?
1480
1481 Instructions:
1482 Compare Answer 1 and Answer 2 for each criterion.
1483 Choose the better answer and explain why.
1484 Select an overall winner based on balance across all three.
1485
1486 Input:
1487 Question: {query}
1488 Answer 1: {answer1}
1489 Answer 2: {answer2}
1490
1491 Output Format (JSON):
1492
1493 {
1494   "Comprehensiveness": {
1495     "Winner": "Answer 1 or Answer 2",
1496     "Explanation": "..."
1497   },
1498   "Diversity": {
1499     "Winner": "Answer 1 or Answer 2",
1500     "Explanation": "..."
1501   },
1502   "Empowerment": {
1503     "Winner": "Answer 1 or Answer 2",
1504     "Explanation": "..."
1505   },
1506   "Overall Winner": {
1507     "Winner": "Answer 1 or Answer 2",
1508     "Explanation": "..."
1509   }
1510 }
1511

```

**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.