ReGraphRAG: Reorganizing Fragmented Knowledge Graphs for Multi-Perspective Retrieval-Augmented Generation

Yebeen Lee

February 27, 2026

INFONET Journal Club

Paper details

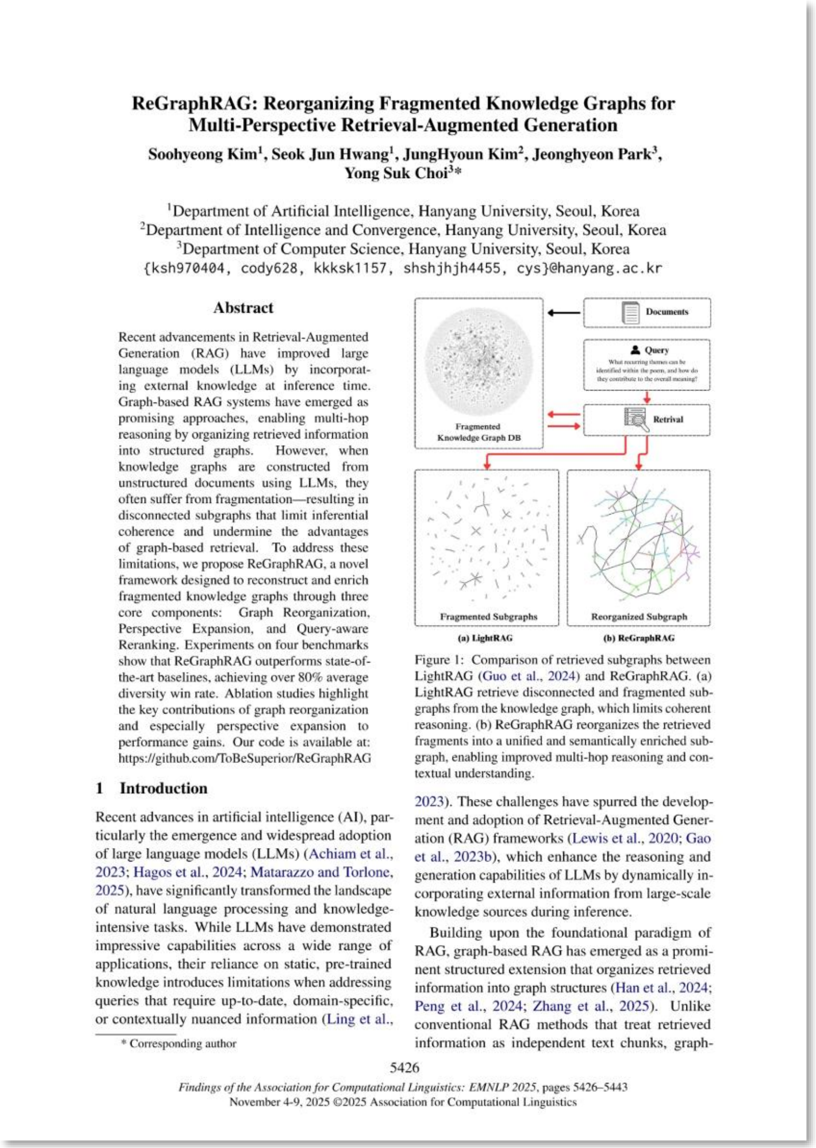

ReGraphRAG: Reorganizing Fragmented Knowledge Graphs for Multi-Perspective Retrieval-Augmented Generation

Soohyeong Kim[1], Seok Jun Hwang[1], JungHyoun Kim[2], Jeonghyeon Park[3], Yong Suk Choi[3]*

[1]Department of Artificial Intelligence, Hanyang University, Seoul, Korea
[2]Department of Intelligence and Convergence, Hanyang University, Seoul, Korea
[3]Department of Computer Science, Hanyang University, Seoul, Korea
{ksh970404, cody628, kkksk1157, shshjhjh4455, cys}@hanyang.ac.kr

Abstract

Recent advancements in Retrieval-Augmented Generation (RAG) have improved large language models (LLMs) by incorporating external knowledge at inference time. Graph-based RAG systems have emerged as promising approaches, enabling multi-hop reasoning by organizing retrieved information into structured graphs. However, when knowledge graphs are constructed from unstructured documents using LLMs, they often suffer from fragmentation—resulting in disconnected subgraphs that limit inferential coherence and undermine the advantages of graph-based retrieval. To address these limitations, we propose ReGraphRAG, a novel framework designed to reconstruct and enrich fragmented knowledge graphs through three core components: Graph Reorganization, Perspective Expansion, and Query-aware Reranking. Experiments on four benchmarks show that ReGraphRAG outperforms state-of-the-art baselines, achieving over 80% average diversity win rate. Ablation studies highlight the key contributions of graph reorganization and especially perspective expansion to performance gains. Our code is available at: https://github.com/ToBeSuperior/ReGraphRAG

Figure 1: Comparison of retrieved subgraphs between LightRAG (Guo et al., 2024) and ReGraphRAG. (a) LightRAG retrieve disconnected and fragmented subgraphs from the knowledge graph, which limits coherent reasoning. (b) ReGraphRAG reorganizes the retrieved fragments into a unified and semantically enriched subgraph, enabling improved multi-hop reasoning and contextual understanding.

1 Introduction

Recent advances in artificial intelligence (AI), particularly the emergence and widespread adoption of large language models (LLMs) (Achiam et al., 2023; Hagos et al., 2024; Matarazzo and Torlone, 2025), have significantly transformed the landscape of natural language processing and knowledge-intensive tasks. While LLMs have demonstrated impressive capabilities across a wide range of applications, their reliance on static, pre-trained knowledge introduces limitations when addressing queries that require up-to-date, domain-specific, or contextually nuanced information (Ling et al., 2023). These challenges have spurred the development and adoption of Retrieval-Augmented Generation (RAG) frameworks (Lewis et al., 2020; Gao et al., 2023b), which enhance the reasoning and generation capabilities of LLMs by dynamically incorporating external information from large-scale knowledge sources during inference.

Building upon the foundational paradigm of RAG, graph-based RAG has emerged as a prominent structured extension that organizes retrieved information into graph structures (Han et al., 2024; Peng et al., 2024; Zhang et al., 2025). Unlike conventional RAG methods that treat retrieved information as independent text chunks, graph-

* Corresponding author

Findings of the Association for Computational Linguistics: EMNLP 2025, pages 5426–5443
November 4-9, 2025 ©2025 Association for Computational Linguistics

- **Title:** "ReGraphRAG: Reorganizing Fragmented Knowledge Graphs for Multi-Perspective Retrieval-Augmented Generation"
- **Authors:** Soohyeong Kim , Seok Jun Hwang
- **Affiliation:** Hanyang University, Seoul, Korea
- **Publication:** EMNLP, 2025
- **Link:** https://aclanthology.org/2025.findings-emnlp.290/

Introduction

Motivation

Key problems

1. Text-based RAG (NaiveRAG / HyDE)

- Retrieval is based on **text chunks**
- Cannot explicitly leverage **structural relations**
- Limited capability for true **multi-hop reasoning**

2. Graph-based RAG (GraphRAG / LightRAG)

- **fragmentation remains unresolved**
- Retrieval often relies on **1-hop,** fragmented evidence

Tends to compensate by retrieving **more nodes/context**

Leads to weakened **coherent multi-hop reasoning flow**

Motivation

The retrieval process itself becomes fragmented, reducing the efficiency of downstream reasoning and generation.

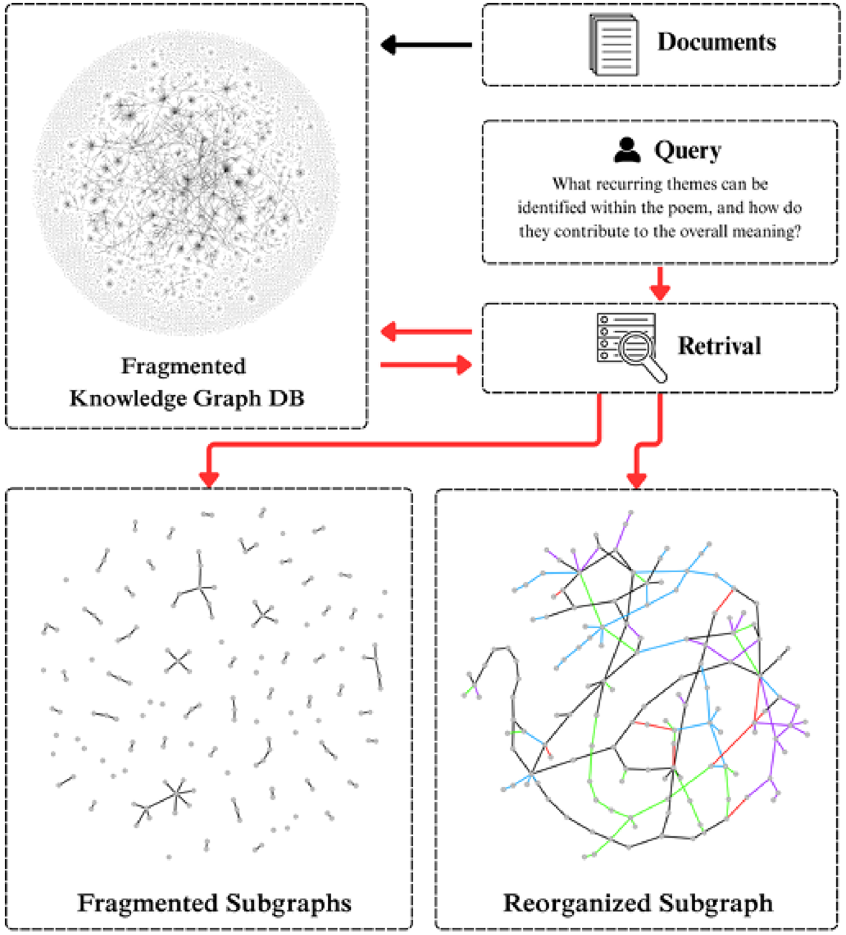

(a) LightRAG (b) ReGraphRAG

propose **ReGraphRAG**

Method

Method steps

Perspective Expansion

Graph Reorganization

Query-aware Reranking

Prompt Instruction

Perspective Expansion

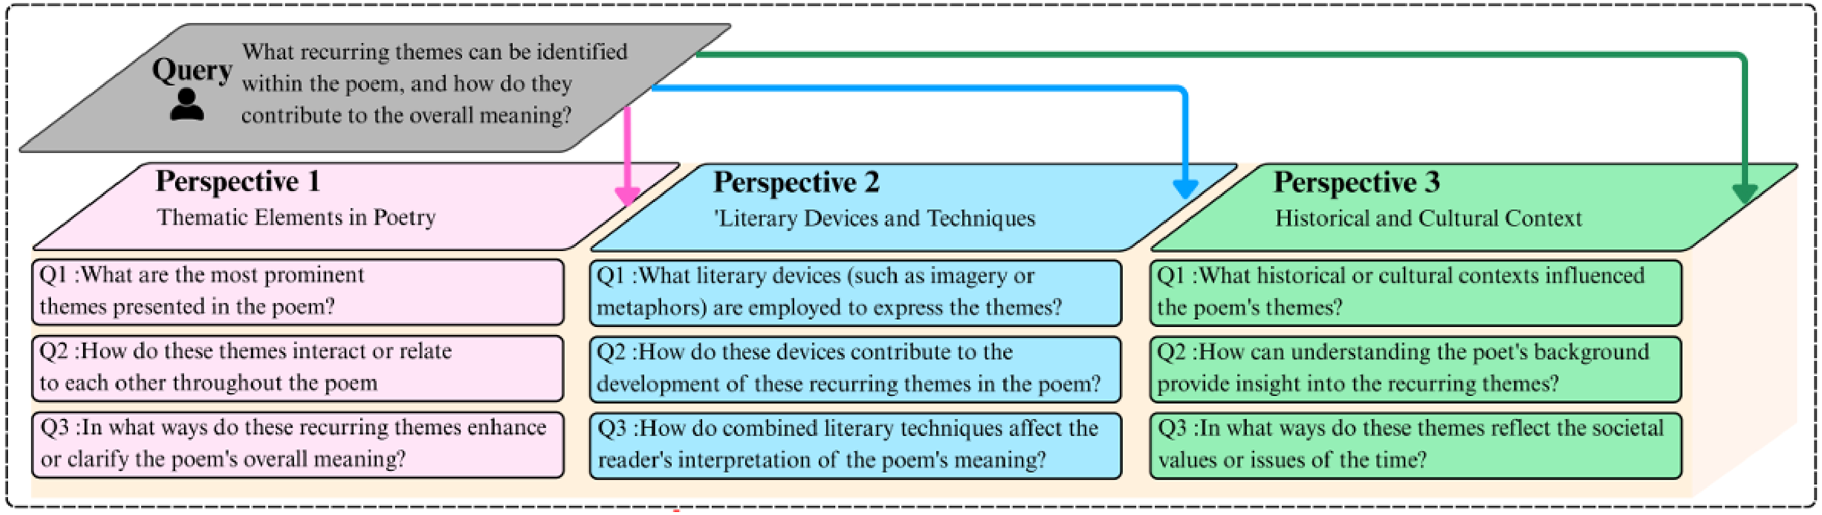

Perspective Expansion

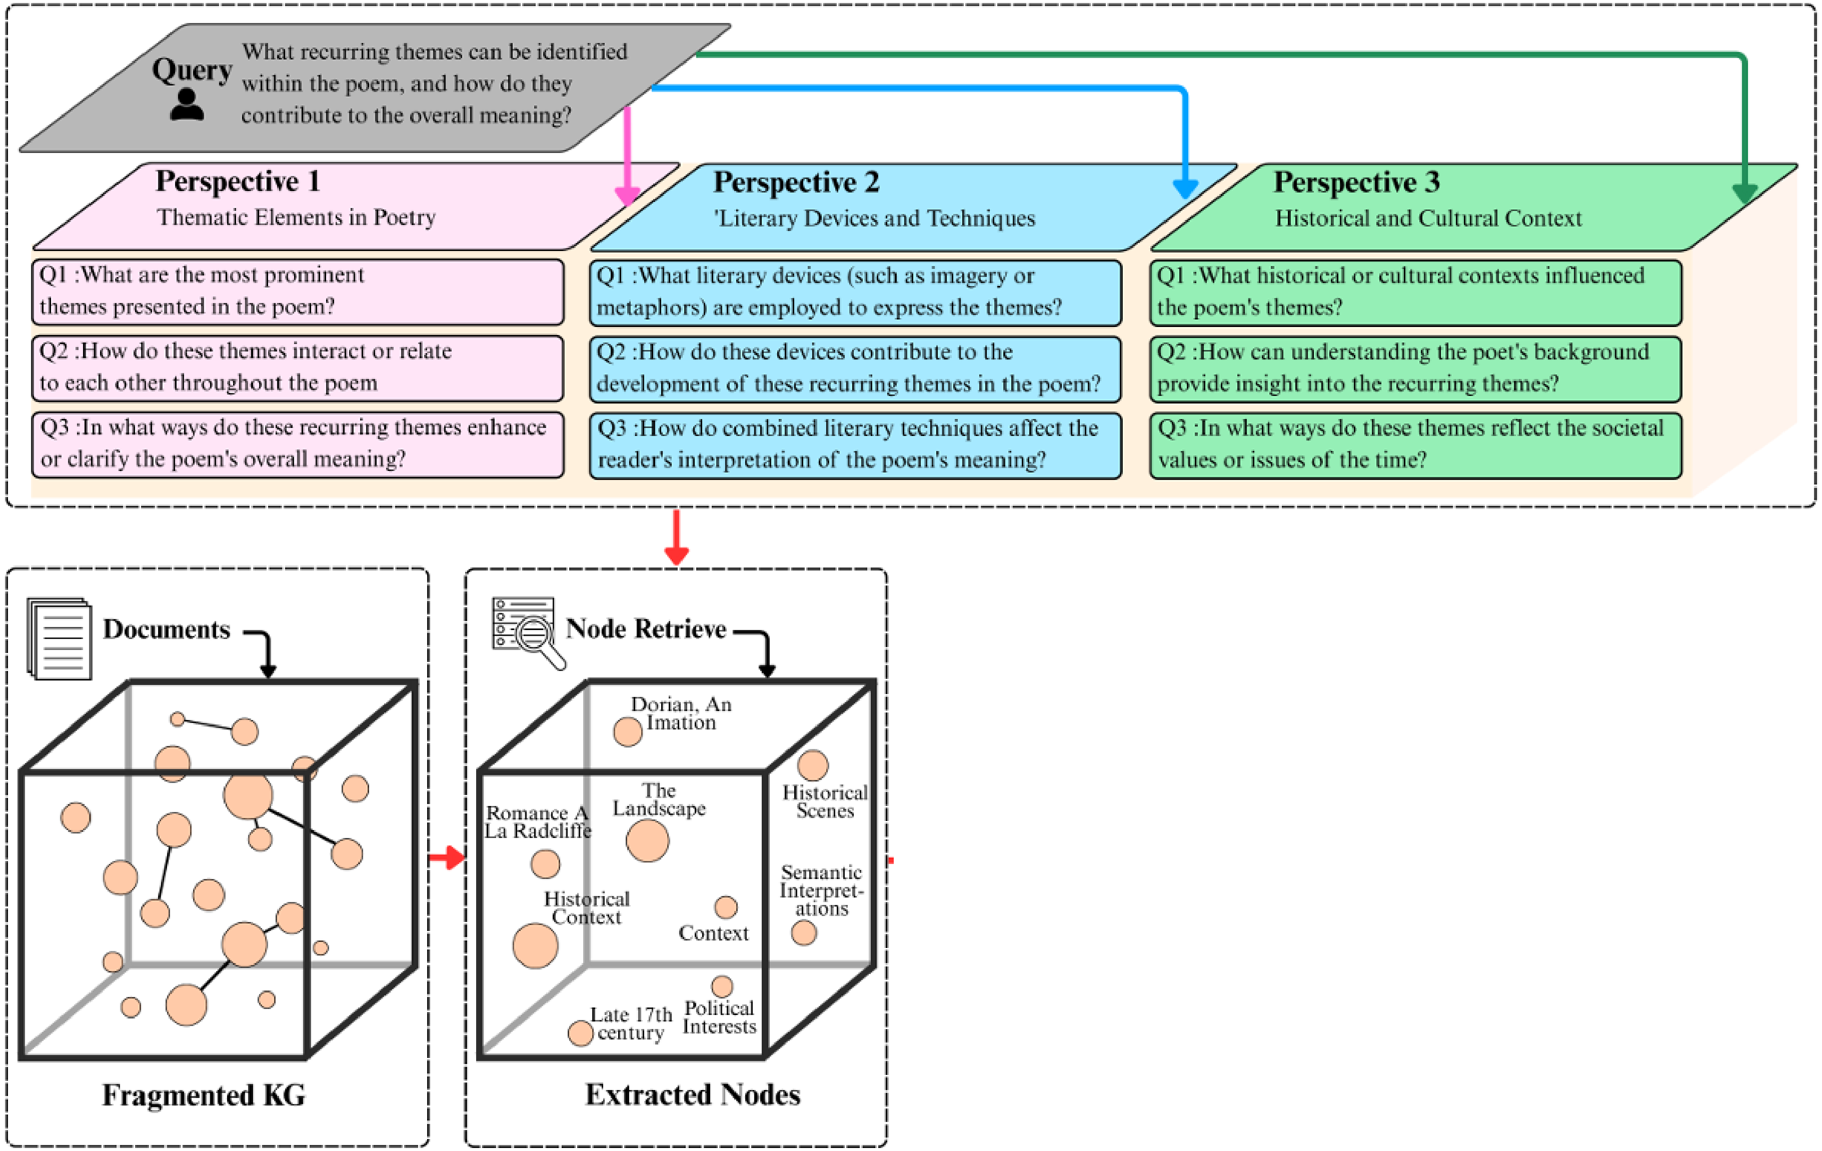

Graph Reorganization

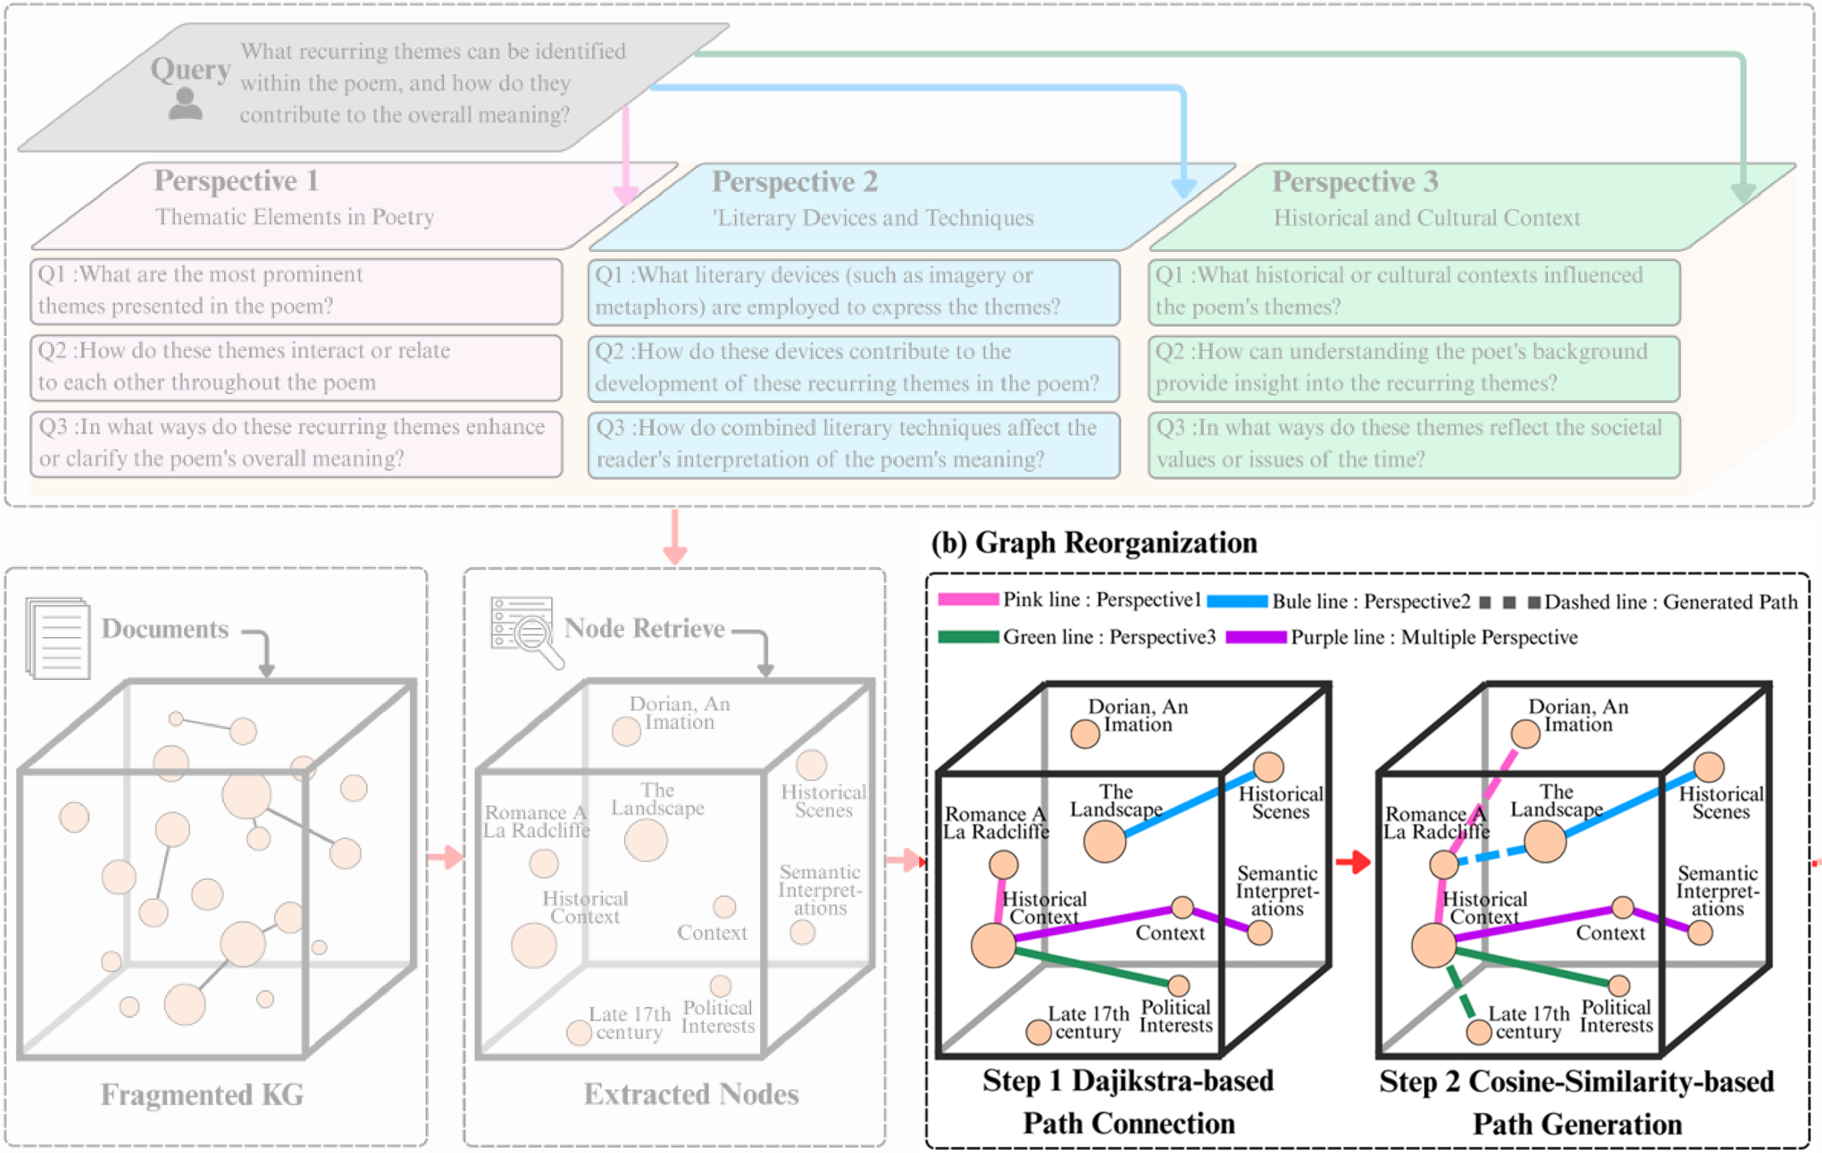

Graph Reorganization

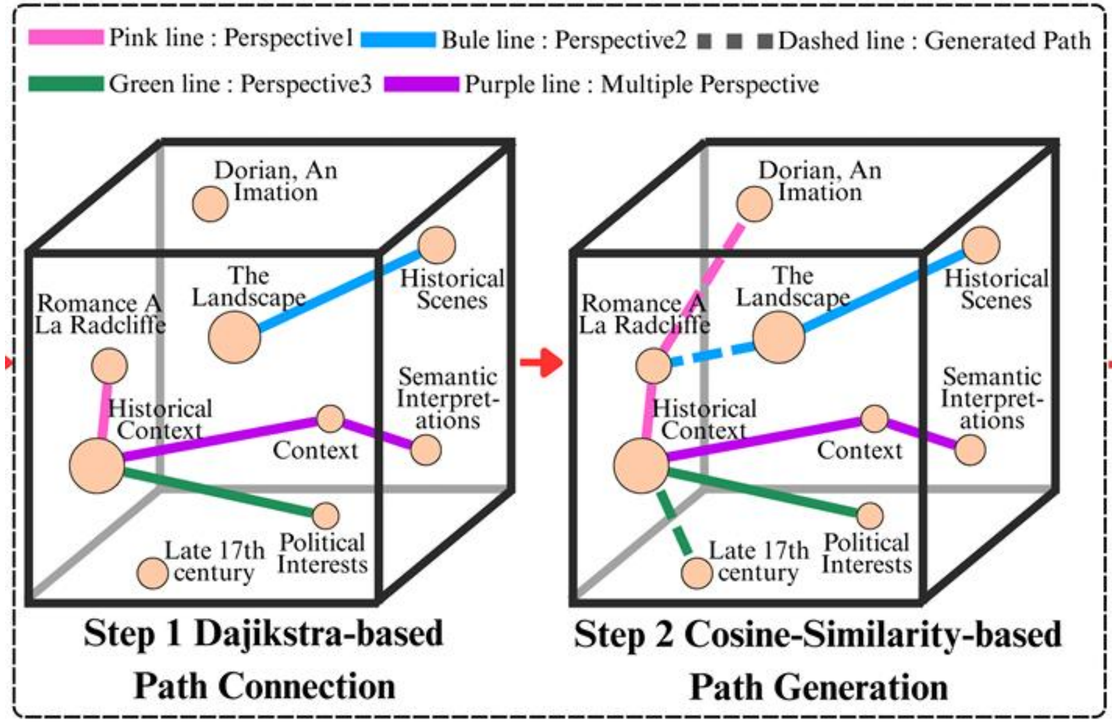

Algorithm 1 Graph Reorganization

Require: Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, list of subgraphs $\{\mathcal{S}_1, \ldots, \mathcal{S}_k\}$ where $\mathcal{S}_i = (\mathcal{V}_i, \mathcal{E}_i)$
Ensure: Connected and merged graph $\mathcal{R}$

$\mathcal{P} \leftarrow \emptyset$ ▷ Step 1: Merge subgraphs based on Dijkstra-based shortest path distances
for $i \leftarrow 1$ to $k-1$ **do**
  for $j \leftarrow i+1$ to k **do**
    $\mathcal{V}_i \leftarrow$ entities of $\mathcal{S}_i$, $\mathcal{V}_j \leftarrow$ entities of $\mathcal{S}_j$
    if paths exist between any $u \in \mathcal{V}_i$ and $v \in \mathcal{V}_j$ **then**
      $p^* \leftarrow$ shortest path between some $u \in \mathcal{V}_i$ and $v \in \mathcal{V}_j$ using Dijkstra
      Add $(i, j, p^*, |p^*|)$ to $\mathcal{P}$
    end if
  end for
end for
Sort $\mathcal{P}$ by path length ascending
for each (i, j, p^*) in $\mathcal{P}$ **do**
  if $\mathcal{S}_i$ and $\mathcal{S}_j$ are not yet merged **then**
    Merge $\mathcal{S}_i$ and $\mathcal{S}_j$, including nodes and edges from p^*
  end if
end for
▷ Step 2: Iteratively merge most similar subgraph pairs based on cosine similarity
$\mathcal{U} \leftarrow$ list of unmerged subgraphs
Extract node embeddings $\Phi(\mathcal{V})$ from subgraphs in $\mathcal{U}$
Compute cosine similarity matrix $\mathcal{C}$ from $\Phi(\mathcal{V})$
while there are disconnected subgraphs in $\mathcal{U}$ **do**
  $(\mathcal{S}_i, \mathcal{S}_j) \leftarrow$ most similar pair in $\mathcal{C}$
  $(u, v) \leftarrow$ most similar entity pair between $\mathcal{S}_i$ and $\mathcal{S}_j$
  Add edge (u, v) to $\mathcal{G}$
  Merge $\mathcal{S}_i$ and $\mathcal{S}_j$ into $\mathcal{S}_{ij}$
  $\mathcal{U} \leftarrow \mathcal{U} \setminus \{\mathcal{S}_i, \mathcal{S}_j\} \cup \{\mathcal{S}_{ij}\}$
  Update $\mathcal{C}$ accordingly
end while
return final connected graph $\mathcal{R}$

Query-aware Reranking & Prompt Instruction

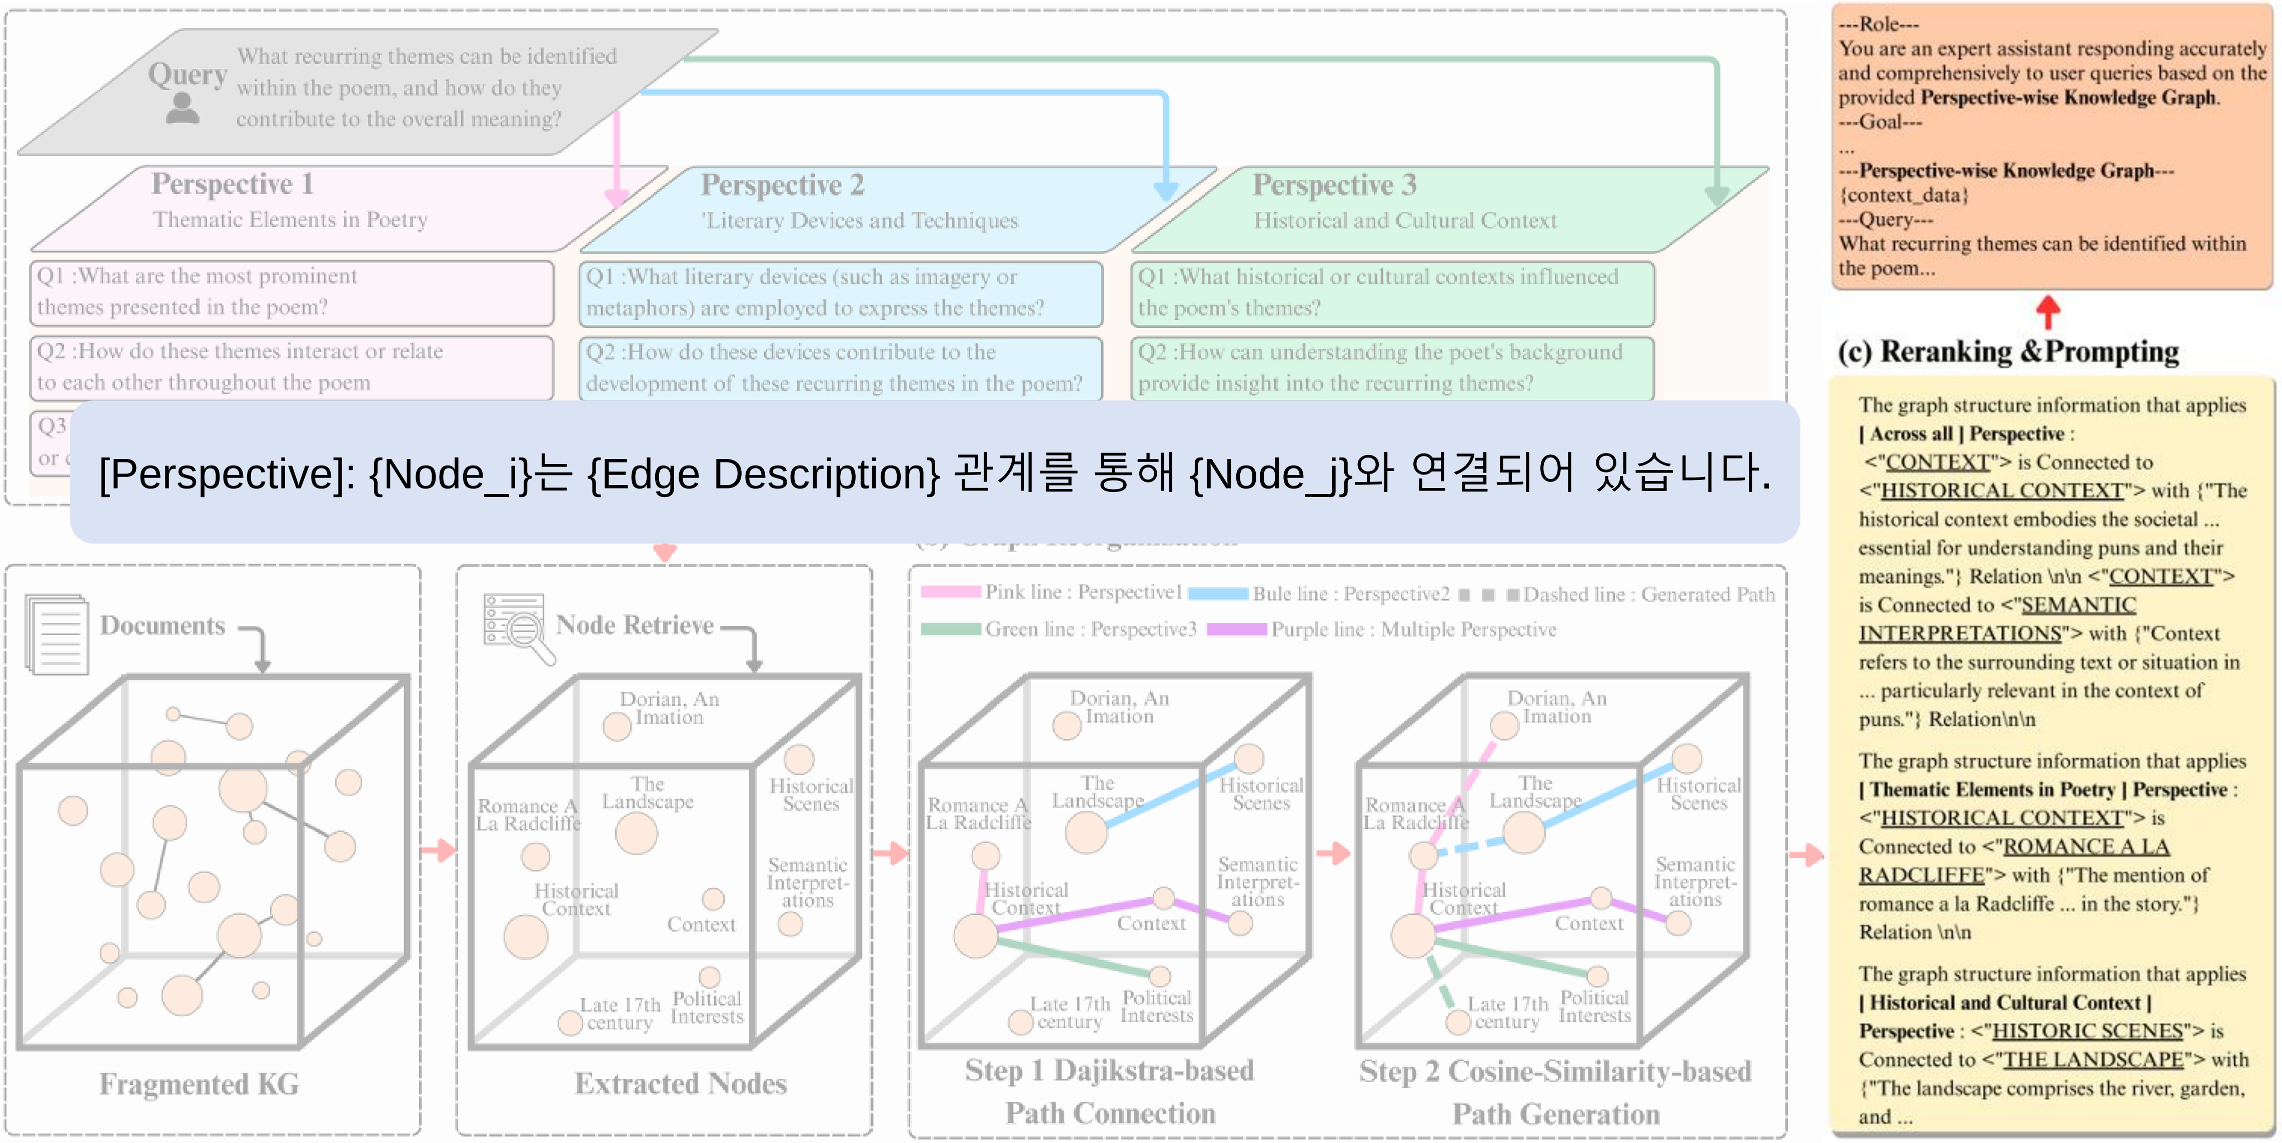

[Perspective]: {Node_i}는 {Edge Description} 관계를 통해 {Node_j}와 연결되어 있습니다.

GIST

Experiment

Experiment

- Baseline: **NaiveRAG** (Gao et al., 2023b), **GraphRAG** (Edge et al., 2024), **HyDE** (Gao et al., 2023a), 그리고 **LightRAG** (Guo et al., 2024)

- Dataset and Metrics: **Ultradomain** (Qian et al., 2025)

- LLM 기반 평가: pairwise comparison
  - Comprehensiveness, Diversity, Empowerment, Overall quality -> win rate

Experiment

	Agriculture		Legal		CS		Mix		Average Win Rate	
	Naïve	ReGraphRAG	Naïve	ReGraphRAG	Naïve	ReGraphRAG	Naïve	ReGraphRAG	Naïve	ReGraphRAG
Comprehensiveness	23.2%	**76.8%**	28.8%	**71.2%**	18.4%	**81.6%**	28.5%	**71.5%**	24.7%	**75.3%**
Diversity	8.0%	**92.0%**	4.8%	**95.2%**	4.0%	**96.0%**	9.8%	**90.2%**	6.7%	**93.4%**
Empowerment	18.4%	**81.6%**	20.8%	**79.2%**	13.6%	**86.4%**	26.8%	**73.2%**	19.9%	**80.1%**
Overall	17.6%	**82.4%**	23.2%	**76.8%**	13.6%	**86.4%**	26.0%	**74.0%**	20.1%	**79.9%**
	HyDE	ReGraphRAG	HyDE	ReGraphRAG	HyDE	ReGraphRAG	HyDE	ReGraphRAG	HyDE	ReGraphRAG
Comprehensiveness	31.2%	**68.8%**	36.0%	**64.0%**	32.0%	**68.0%**	26.8%	**73.2%**	31.5%	**68.5%**
Diversity	17.6%	**82.4%**	15.2%	**84.8%**	7.2%	**92.8%**	10.6%	**89.4%**	12.7%	**87.4%**
Empowerment	26.4%	**73.6%**	28.8%	**71.2%**	26.4%	**73.6%**	26.0%	**74.0%**	26.9%	**73.1%**
Overall	27.2%	**72.8%**	30.4%	**69.6%**	25.6%	**74.4%**	24.4%	**75.6%**	26.9%	**73.1%**
	GraphRAG	ReGraphRAG	GraphRAG	ReGraphRAG	GraphRAG	ReGraphRAG	GraphRAG	ReGraphRAG	GraphRAG	ReGraphRAG
Comprehensiveness	20.0%	**80.0%**	35.2%	**64.8%**	32.0%	**68.0%**	35.8%	**64.2%**	30.8%	**69.3%**
Diversity	10.4%	**89.6%**	22.4%	**77.6%**	8.0%	**92.0%**	19.5%	**80.5%**	15.1%	**84.9%**
Empowerment	19.2%	**80.8%**	32.8%	**67.2%**	30.4%	**69.6%**	42.3%	**57.7%**	31.2%	**68.8%**
Overall	17.6%	**82.4%**	33.6%	**66.4%**	30.4%	**69.6%**	36.6%	**63.4%**	29.6%	**70.5%**
	LightRAG	ReGraphRAG	LightRAG	ReGraphRAG	LightRAG	ReGraphRAG	LightRAG	ReGraphRAG	LightRAG	ReGraphRAG
Comprehensiveness	20.0%	**80.0%**	30.4%	**69.6%**	23.2%	**76.8%**	26.8%	**73.2%**	25.1%	**74.9%**
Diversity	8.0%	**92.0%**	6.4%	**93.6%**	8.0%	**92.0%**	8.1%	**91.9%**	7.6%	**92.4%**
Empowerment	16.8%	**83.2%**	27.2%	**72.8%**	17.6%	**82.4%**	26.8%	**73.2%**	22.1%	**77.9%**
Overall	15.2%	**84.8%**	26.4%	**73.6%**	17.6%	**82.4%**	25.2%	**74.8%**	21.1%	**78.9%**

Experiment

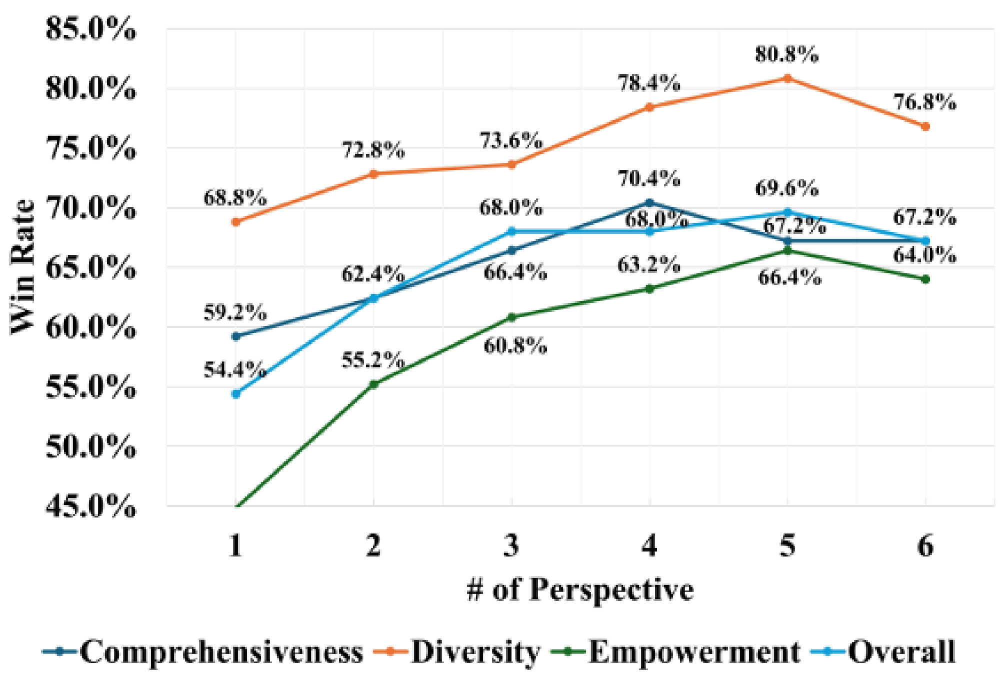

Figure 3: Win rate of ReGraphRAG compared to GraphRAG as **the number of perspectives** increases, measured across four evaluation dimensions: Comprehensiveness, Diversity, Empowerment, and Overall.

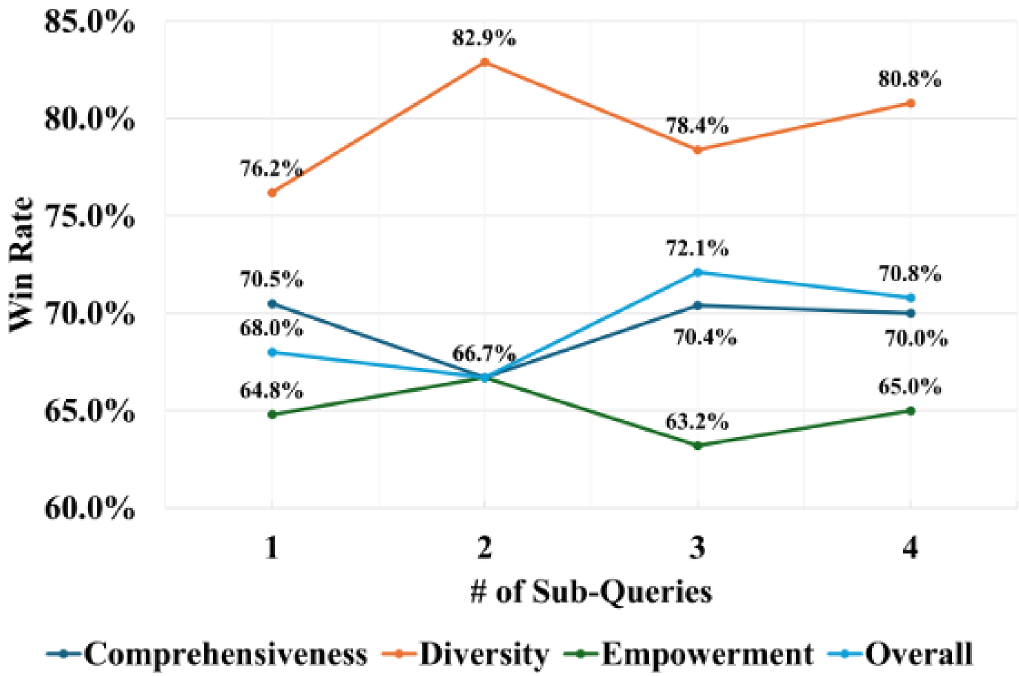

Figure 4: Win rate of ReGraphRAG compared to GraphRAG as **the number of sub-queries** per perspective increases, evaluated across Comprehensiveness, Diversity, Empowerment, and Overall dimensions.

Ablation study

	Agriculture		Legal		CS		Mix		Average Win Rate	
	w/o P-exp	ReGraphRAG	w/o P-exp	ReGraphRAG	w/o P-exp	ReGraphRAG	w/o P-exp	ReGraphRAG	w/o P-exp	ReGraphRAG
Comprehensiveness	38.4%	**61.6%**	42.4%	**57.6%**	41.6%	**58.4%**	42.4%	**57.6%**	41.2%	**58.8%**
Diversity	29.6%	**70.4%**	33.6%	**66.4%**	27.2%	**72.8%**	38.4%	**61.6%**	32.2%	**67.8%**
Empowerment	34.4%	**65.6%**	35.2%	**64.8%**	40.0%	**60.0%**	40.0%	**60.0%**	37.4%	**62.6%**
Overall	35.2%	**64.8%**	36.0%	**64.0%**	40.0%	**60.0%**	40.8%	**59.2%**	38.0%	**62.0%**
	w/o Reorg	ReGraphRAG	w/o Reorg	ReGraphRAG	w/o Reorg	ReGraphRAG	w/o Reorg	ReGraphRAG	w/o Reorg	ReGraphRAG
Comprehensiveness	48.8%	**51.2%**	48.8%	**51.2%**	46.8%	**53.2%**	48.0%	**52.0%**	48.1%	**51.9%**
Diversity	48.8%	**51.2%**	48.0%	**52.0%**	47.6%	**52.4%**	48.4%	**51.6%**	48.2%	**51.8%**
Empowerment	44.8%	**55.2%**	47.8%	**52.2%**	43.5%	**56.5%**	41.8%	**58.2%**	44.5%	**55.5%**
Overall	46.4%	**53.6%**	46.0%	**54.0%**	45.2%	**54.8%**	46.4%	**53.6%**	46.0%	**54.0%**
	w/o Rerank	ReGraphRAG	w/o Rerank	ReGraphRAG	w/o Rerank	ReGraphRAG	w/o Rerank	ReGraphRAG	w/o Rerank	ReGraphRAG
Comprehensiveness	**52.0%**	48.0%	**52.8%**	47.2%	**52.0%**	48.0%	**52.0%**	48.0%	**52.2%**	47.8%
Diversity	**52.0%**	48.0%	**56.0%**	44.0%	**54.4%**	45.6%	**56.1%**	43.9%	**54.6%**	45.4%
Empowerment	**52.8%**	47.2%	**54.4%**	45.6%	**52.0%**	48.0%	**57.7%**	42.3%	**54.2%**	45.8%
Overall	**54.4%**	45.6%	**54.4%**	45.6%	**52.8%**	47.2%	**54.5%**	45.5%	**54.0%**	46.0%

Discussion

- Possible cause: The observed behavior may stem not only from the reordering module itself, but also from intrinsic graph-structure characteristics or limitations in prompt tuning/design.

- Prompt limitation: Our prompting strategy preserves the graph format mainly as triplets and does not explicitly model deep multi-hop prioritization (§4.4).

- Implication: As a result, structural reasoning cues (e.g., paths/chains) may not be sufficiently exposed, potentially diluting the impact of reordering.

Limitation

Limited effectiveness of the query-aware reordering module	In ablation, removing the reordering stage sometimes **improved performance**, contrary to expectations.	The current strategy—**cosine similarity between edge embeddings and query representations**—may fail to fully capture the **multi-hop reasoning potential** provided by graph structure. Also, the **interaction between prompt design (graph-structured prompting) and reordering** remains underexplored.
Computational overhead	The **multi-perspective expansion** (decomposing the original query into multiple perspectives and generating multiple sub-queries per perspective) increases both **retrieval and inference time**.	While it improves diversity and semantic coverage, it can reduce practicality in **real-time** or **resource-constrained** settings.

-