

# Core-based Hierarchies for Efficient GraphRAG

Jakir Hossain  
University at Buffalo  
Buffalo, NY, USA  
mh267@buffalo.edu

Ahmet Erdem Sariyüce\*  
University at Buffalo  
Buffalo, NY, USA  
erdem@buffalo.edu

## Abstract

Retrieval-Augmented Generation (RAG) enhances large language models by incorporating external knowledge. However, existing vector-based methods often fail on global sensemaking tasks that require reasoning across many documents. GraphRAG addresses this by organizing documents into a knowledge graph with hierarchical communities that can be recursively summarized. Current GraphRAG approaches rely on Leiden clustering for community detection, but we prove that on sparse knowledge graphs, where average degree is constant and most nodes have low degree, modularity optimization admits exponentially many near-optimal partitions, making Leiden-based communities inherently non-reproducible. To address this, we propose replacing Leiden with  $k$ -core decomposition, which yields a deterministic, density-aware hierarchy in linear time. We introduce a set of lightweight heuristics that leverage the  $k$ -core hierarchy to construct size-bounded, connectivity-preserving communities for retrieval and summarization, along with a token-budget-aware sampling strategy that reduces LLM costs. We evaluate our methods on real-world datasets including financial earnings transcripts, news articles, and podcasts, using three LLMs for answer generation and five independent LLM judges for head-to-head evaluation. Across datasets and models, our approach consistently improves answer comprehensiveness and diversity while reducing token usage, demonstrating that  $k$ -core-based GraphRAG is an effective and efficient framework for global sensemaking.

## CCS Concepts

• Information systems → Retrieval models and ranking.

## Keywords

GraphRAG,  $k$ -core, Leiden, RAG, information retrieval

## ACM Reference Format:

Jakir Hossain and Ahmet Erdem Sariyüce. 2026. Core-based Hierarchies for Efficient GraphRAG. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770855.3818007>

## Resource Availability:

The source code of this paper has been made publicly available at <https://doi.org/10.5281/zenodo.20500254>.

\*Amazon Web Services. This publication describes work performed at University at Buffalo, and is not associated with Amazon Web Services.



This work is licensed under a Creative Commons Attribution 4.0 International License. *KDD '26, Jeju Island, Republic of Korea*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2259-2/2026/08  
<https://doi.org/10.1145/3770855.3818007>

## 1 Introduction

Retrieval-Augmented Generation (RAG) enhances large language models (LLMs) with external knowledge, enabling more grounded and accurate responses to complex queries [2, 8, 19, 20, 39]. It is particularly useful when the corpus is too large to fit within an LLM’s context window [18, 21]. In a typical RAG setup, a small set of relevant records is retrieved and fed, along with the query, to the LLM for response generation. RAG has been widely adopted in domains such as healthcare, law, finance, and education, improving faithfulness, reducing hallucinations, mitigating privacy risks, and enhancing robustness [22, 36, 37, 40, 41].

Most RAG systems rely on lexical/semantic retrieval over text and are effective only for queries answerable from a few documents. These approaches struggle with multi-hop reasoning or queries requiring synthesis across semantically diverse documents or global corpus understanding [18, 21]. For instance, answering “How have patient outcomes in cancer treatment evolved in response to multi-center clinical trials over the last 15 years?” demands aggregating evidence across many sources, which conventional RAG methods cannot easily support. Similarly, a financial analyst asking “What common strategies have semiconductor companies adopted in response to supply chain disruptions over the past decade?” must synthesize themes across hundreds of earnings call transcripts, a task that no single retrieved passage can answer. We refer to such queries as *global sensemaking* tasks: they require the system to identify recurring themes, reconcile conflicting perspectives, and synthesize evidence distributed across the entire corpus, capabilities that go beyond standard RAG.

Graph-based Retrieval-Augmented Generation (GraphRAG) has recently emerged as a promising approach for global sensemaking tasks by representing the corpus as a knowledge graph and organizing it into a hierarchy of communities that can be recursively summarized [11]. The GraphRAG framework introduced by Edge et al. relies on modularity-based community detection, such as Leiden [33]. However, Leiden often produces hierarchies that are shallow, overly fragmented, and/or dominated by a few large communities. Building on an observation by Good, de Montjoye, and Clauset [10], we prove that when the average degree is constant and most nodes have low degree—the regime of typical knowledge graphs—the number of near-optimal modularity partitions is exponential in the graph size (Theorem 1). This implies that Leiden-based communities are inherently non-reproducible on such graphs, compounding known issues with resolution limits and sensitivity to initialization. In practice, this instability causes communities to unpredictably merge or fragment semantically meaningful structures.

Motivated by these limitations, we explore  *$k$ -core-based alternatives* for hierarchical graph construction in GraphRAG. The  $k$ -core decomposition organizes a network into nested layers of increasing minimum degree [27]. Each node is assigned a core number: the

largest  $k$  for which it belongs to a subgraph where every node has at least  $k$  neighbors. This creates a deterministic, density-aware hierarchy that can be computed in a single  $O(|E|)$  pass. Unlike modularity-driven methods, the  $k$ -core hierarchy naturally captures progressively denser and more cohesive substructures, making it well suited for global sensemaking tasks that require integrating evidence across multiple interconnected regions of the corpus. In knowledge graphs specifically, higher  $k$ -cores correspond to entities connected through multiple distinct relational paths, providing a natural proxy for topical centrality that modularity’s comparison to a degree-preserving null model cannot capture.

We operationalize this hierarchy through a set of lightweight, interpretable community construction heuristics that operate over different levels of the  $k$ -core decomposition. Our approach constructs balanced hierarchical communities while explicitly controlling cluster sizes to respect LLM context constraints. These heuristics enable comprehensive and diverse summaries without excessive token usage. Additionally, we introduce a token-budget-aware sampling strategy that further reduces LLM costs while preserving retrieval quality. We make the following key contributions:

- We introduce  $k$ -core decomposition as a drop-in replacement for Leiden in GraphRAG, yielding deterministic, density-aware hierarchies in linear time.
- We prove that modularity optimization on sparse graphs admits exponentially many near-optimal partitions (Theorem 1), formally explaining why Leiden-based community detection is unreliable on knowledge graphs.
- We propose multiple hierarchical heuristic strategies that exploit different regions of the induced  $k$ -core hierarchy to balance coverage, granularity, and efficiency.
- We conduct extensive evaluations on three real-world datasets, using three LLMs for answer generation and five independent LLM judges for head-to-head evaluation.

## 2 Related Work and Background

In this section, we review the literature on Graph-based RAG and remind key concepts of GraphRAG workflows and  $k$ -core decomposition.

**Graph-based RAG.** Traditional vector-based RAG retrieves relevant passages to augment LLM responses, but is largely limited to single-hop, fact-based queries where the answer resides within one or a few contiguous passages. When reasoning must span multiple documents or synthesize information across a corpus, these approaches fall short [11, 24].

To address this, several methods incorporate graph structures into the retrieval process. One line of work constructs query-focused subgraphs from knowledge graphs and performs reasoning over them. Systems such as KG-GPT [17], G-Retriever [12], GRAG [15], and HLG [9] follow this paradigm, extracting relevant subgraphs and using LLMs to reason over the retrieved structure, thereby improving multi-hop answer quality. A complementary line of work enhances the retrieval step itself with graph signals, using graph topology to guide which documents or passages are selected [24, 29, 35]. However, both subgraph-reasoning and graph-enhanced retrieval methods are primarily designed for queries that require a bounded number of reasoning hops. They do not address

*global sensemaking*, tasks where the answer requires understanding themes, patterns, or relationships that emerge only when considering an entire corpus.

GraphRAG [7] addresses this gap by applying Leiden-based hierarchical community detection on a knowledge graph and using the resulting communities for global sensemaking (see Section 2.1 for details). A key limitation, however, is that modularity-based clustering methods like Leiden can produce highly uneven hierarchies, where community sizes vary widely across levels. This imbalance can lead to inconsistent summary quality and makes it difficult to control the granularity of retrieved context. Our work addresses this limitation by replacing community detection with  $k$ -core decomposition, which yields a naturally nested and more balanced hierarchy of dense subgraphs.

### 2.1 Community-based GraphRAG Overview

Edge et al. [7] introduced a graph-based approach for global sense-making question answering by a community-driven GraphRAG workflow. Here, we briefly describe the two main stages of the GraphRAG workflow: indexing and query-time answer generation, followed by their evaluation conditions.

**Indexing.** Large text corpora are first split into manageable chunks, from which entities, relationships, and claims are extracted using LLMs. These elements are aggregated into a knowledge graph, where nodes represent entities or claims and edges represent relationships, weighted by frequency. The graph is hierarchically partitioned into communities using Leiden community detection [33], recursively identifying sub-communities until reaching leaf-level communities that cannot be further split. Leaf-level communities are summarized first by prioritizing elements according to edge prominence, adding source and target node descriptions, edges, and related claims into the LLM context until the token limit is reached. For higher-level communities, if all element summaries fit within the context window, they are summarized directly; otherwise, sub-community summaries iteratively replace longer element-level descriptions until they fit. This process, from knowledge graph construction to hierarchical community summaries, is referred to as indexing.

**Query-Time Answer Generation.** Given a user query, indexed community summaries are used to generate a final answer through a multi-stage Map-Reduce process. First, summaries are randomly shuffled and divided into chunks of a pre-specified token size to ensure relevant information is distributed across multiple contexts rather than concentrated in a single window. Next, intermediate answers are generated for each chunk in parallel, with the LLM also assigning a helpfulness score between 0 and 100. These intermediate answers are then sorted by helpfulness and iteratively added into a new context window until the token limit is reached. The LLM uses this aggregated context to produce the final global answer. By leveraging the hierarchical structure of the communities, this approach allows answers to be generated from different levels and enables robust aggregation of information from all relevant sub-communities.

**Evaluation for GraphRAG.** Evaluating RAG systems for multi-document reasoning presents distinct challenges beyond standard

QA benchmarks. Recent work assesses output quality along dimensions such as comprehensiveness, diversity, empowerment, and directness [7]. Edge et al. tested six conditions, including four hierarchical community levels (C0–C3), source-text summarization, and a vector-based RAG approach. The hierarchical configurations differ in granularity. **C0** uses only top-level communities, yielding the smallest set of summaries, **C1** uses communities one level below the root for a slightly finer decomposition, and so on. Results on two datasets show that all community-level approaches (C0–C3) outperform both source-text summarization and vector RAG for global sensemaking. Among these, **C2** and **C3** consistently achieve the best performance, hence we consider those two for our evaluation.

## 2.2 Hierarchical $k$ -core Decomposition

In this work, we consider the knowledge graph as an unweighted graph  $G = (V, E)$ , where  $V$  denotes the set of nodes and  $E$  denotes the set of edges. The degree of a node  $u$  within a subgraph  $S \subseteq G$  is denoted  $\deg(u, S)$ . A  $k$ -core is the maximal connected subgraph  $S \subseteq G$  in which every vertex has at least  $k$  neighbors, i.e.,  $\deg(u, S) \geq k \forall u \in S$ . Each node is assigned a *core number*, the largest  $k$  for which it belongs to a  $k$ -core. The  $k$ -shell is the set of nodes with core number  $k$  [4]. The  $k$ -cores for all  $k$  can be computed by recursively removing nodes with degree less than  $k$  and their edges, which runs in  $O(|E|)$  time [1]. This process of identifying all  $k$ -cores and assigning core numbers is referred to as core decomposition. Since nodes can belong to multiple  $k$ -cores with different  $k$  values, the decomposition naturally forms a hierarchy of nested dense subgraphs. This hierarchy can be represented as a tree, with nodes denoting subgraphs and edges encoding containment relationships. The root represents the entire graph (1-core), while child nodes correspond to subgraphs with increasing core numbers. Internal nodes may contain denser subgraphs with higher core values, forming a multi-level structure that captures the nested organization of densely connected regions within the graph.

## 3 Why Modularity Optimization is Unreliable on Sparse Knowledge Graphs?

Modularity-based community detection methods such as Leiden [33] compare observed within-community edge density to the expectation under a degree-preserving null model (the configuration model), an approach that works well in dense networks but degrades in sparse ones. Good, de Montjoye, and Clauset [10] provided a systematic characterization of this degradation. They showed that the modularity function  $Q$  typically admits an exponentially large number of distinct high-scoring partitions while lacking a clear global maximum: a phenomenon they termed the *degeneracy of modularity*. Using real-world metabolic networks, they demonstrated that these near-optimal partitions can fundamentally disagree on key properties such as the composition of the largest modules and the distribution of module sizes, meaning that any optimizer (including Louvain and Leiden) is effectively selecting one solution from a vast equivalence class, guided more by random seeds and tie-breaking rules than by genuine structure. They further showed that  $Q_{max}$  depends strongly on both network size and the number of modules,

and that degeneracy is most severe in sparse and hierarchically structured networks.

**Modularity Degeneracy.** We give a key structural observation: the low-degree nodes (degree at most a fixed constant  $d$ ) have vanishing modularity sensitivity in sparse graphs, because their few edges contribute negligibly to the objective. Denote by  $n_{\leq d} = |\{i \in V : k_i \leq d\}|$  the number of such nodes. For small  $d$ , these nodes can be reassigned across communities with negligible impact on modularity, creating an exponentially large degenerate set.

Recall that the modularity of a partition  $\sigma$  of a graph  $G = (V, E)$  with  $n$  nodes and  $m$  edges is  $Q(\sigma) = \frac{1}{2m} \sum_{i,j} [A_{ij} - k_i k_j / 2m] \delta(\sigma_i, \sigma_j)$ , where  $k_i$  is the degree of node  $i$ . Let  $Q^* = \max_{\sigma} Q(\sigma)$  denote the optimal modularity. For  $\varepsilon > 0$ , we define the  $\varepsilon$ -degeneracy of  $G$  as the number of structurally distinct partitions achieving near-optimal modularity:  $\mathcal{D}(\varepsilon) = |\{\sigma \in \mathcal{P}(G) : Q^* - Q(\sigma) < \varepsilon\}|$ .

**THEOREM 1 (MODULARITY DEGENERACY IN SPARSE GRAPHS).** *Let  $G$  be a graph with  $n$  nodes,  $m$  edges, average degree  $\bar{k} = 2m/n = O(1)$ , and  $n_{\leq d} = \Theta(n)$ . Then for any  $\varepsilon > d(2 + \bar{k})/(2m)$ ,*

$$\mathcal{D}(\varepsilon) \geq 2^{n_{\leq d}/(d+1)}.$$

*In particular, the number of near-optimal partitions is exponential in  $n$ , and the tolerance threshold  $\varepsilon$  required to trigger this blowup is  $O(1/n)$ .*

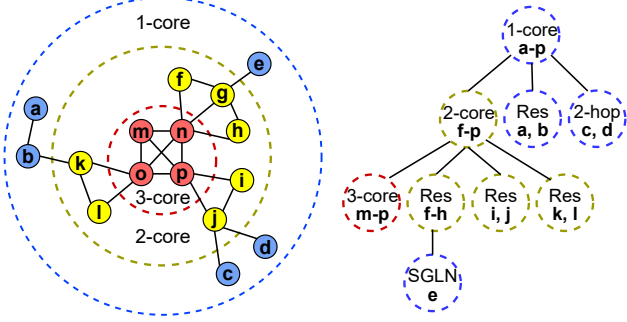
Proof sketch, the full proof, and experimental validation are given in Appendix A.

Knowledge graphs (KG) are particularly well-suited to  $k$ -core decomposition over modularity-based methods for two structural reasons. First, their heavy-tailed but low-mean degree distributions mean that  $n_{\leq d} = \Theta(n)$ , placing them squarely in the regime where Theorem 1 bites hardest. Second, KG edges are semantically meaningful rather than stochastic, so the  $k$ -core criterion, requiring each node to have at least  $k$  neighbors, directly captures genuine relational richness in a way that comparison to a degree-preserving null model does not.

## 4 A Robust Alternative: $k$ -core Decomposition

Theorem 1 exposes a fundamental instability: for a graph with low average degree, any modularity optimizer, Leiden included, selects one partition from an exponentially large near-optimal plateau. Different seeds, tie-breaking rules, or minor edge perturbations yield structurally different communities.

The  $k$ -core decomposition sidesteps this problem entirely. It is *unique*: for every graph  $G$  and integer  $k$ , the  $k$ -core  $H_k$  is the unique maximal subgraph with minimum degree  $\geq k$ , computed by the deterministic peeling process; no optimization landscape, no stochasticity. It is *robust*: adding or removing a few edges changes the shells by at most a proportional number of nodes, whereas Theorem 1 implies that even a single edge can shift the optimal partition across the entire degenerate set. And  $k$ -core decomposition provides a *natural hierarchy*: the nested shells  $H_1 \supseteq H_2 \supseteq \dots$  reflect structural connectivity rather than comparison to a null model. In the sparse regime where  $k_i \cdot k_j / 2m \approx 0$  for most pairs, the null model is nearly vacuous yet  $k$ -core structure remains informative. The 2-core captures the backbone of multiply-connected nodes while degree-1 nodes form the periphery. In knowledge graphs, this distinction is semantically meaningful: multiple relational paths



**Figure 1:**  $k$ -core decomposition (left) and corresponding hierarchy tree produced by RkH (right).

between entities signal genuine topical relatedness. Last, but not least, computing  $k$ -core is much cheaper than Leiden; a single pass over the graph suffices to compute all and a hierarchy [25].

In GraphRAG, community assignments determine which entities are summarized together and which retrieval units are constructed. Exponential degeneracy means Leiden-based pipelines produce non-reproducible summaries: the same knowledge graph under different random seeds yields different communities and different retrieval behavior, with communities that are either too fragmented (splitting apart a natural topic) or too arbitrary (grouping unrelated peripheral nodes together because the optimization found a marginal modularity gain). The  $k$ -core decomposition eliminates this variance. Its nested shells provide a small number of progressively tighter subgraphs where the innermost cores represent the most interconnected, and therefore most semantically central concepts, while the outer shells supply context. This maps naturally onto a summarization hierarchy: summarize the dense core first, then expand outward, producing stable, deterministic retrieval units from which summaries and indices can be reliably built.

**Remark.** A natural refinement of  $k$ -core is the  $k$ -truss [6], which requires every edge to participate in at least  $k - 2$  triangles, yielding tighter subgraphs. However, knowledge graphs are triangle-poor: edges typically connect entities through distinct relation types (e.g., *born\_in*, *capital\_of*), producing bipartite-like local structure with global clustering coefficients well below 0.05. Even the 3-truss discards the vast majority of the graph, making the degree-based  $k$ -core criterion the right granularity for sparse relational graphs. A brief evaluation is provided in Appendix C.

Motivated by these advantages, we propose a set of heuristics for efficient global sensemaking in GraphRAG. These heuristics leverage the  $k$ -core hierarchy to generate size-constrained clusters, merge small two-hop clusters, and handle residual components.

#### 4.1 Handling Residuals in $k$ -core Hierarchy

Our first heuristic builds on the  $k$ -core hierarchy while addressing practical issues such as singleton nodes and oversized clusters by separating dense cores from sparse residuals. At each level, dense components are recursively partitioned into size-bounded, connectivity-preserving clusters, while low-core residuals are handled separately to avoid distorting core structure. The result is a deterministic hierarchy that captures both central regions and peripheral context, enabling reliable hierarchical summarization.

#### Algorithm 1: RkH: Residual-aware $k$ -core Hierarchy ( $G, M$ )

---

**Require:** Graph  $G = (V, E)$ , max cluster size  $M$   
**Ensure:** Hierarchical clusters  $C$

- 1:  $G \leftarrow \text{LCC}(G)$ ; remove self-loops
- 2: Compute core number  $c(v)$  for all  $v \in V$
- 3: Initialize queue  $Q \leftarrow \{(V)\}$  and level  $\ell \leftarrow 1$
- 4: Initialize cluster set  $C \leftarrow \emptyset$ , Global singleton  $\mathcal{G}_s \leftarrow \emptyset$
- 5: **while**  $\ell \leq \max c(v)$  **do**
- 6:    $Q' \leftarrow \emptyset$
- 7:   **for** each cluster  $S \in Q$  **do**
- 8:      $S_c \leftarrow \{v \in S \mid c(v) \geq \ell\}$ ,  $S_r \leftarrow S \setminus S_c$
- 9:     **for** each connected component  $R$  of  $G[S_c]$  **do**
- 10:        $R' \leftarrow \begin{cases} \{R\}, & |R| \leq M \\ \text{SPLIT}(G, R, M) & \text{otherwise} \end{cases}$
- 11:        $C \leftarrow C \cup R'$ ,  $Q' \leftarrow Q' \cup R'$
- 12:     **end for**
- 13:     **for** each connected component  $R$  of  $G[S_r]$  **do**
- 14:        $C \leftarrow C \cup \begin{cases} \{R\}, & |R| \leq M \\ \text{SPLIT}(G, R, M) & \text{otherwise} \end{cases}$
- 15:     **end for**
- 16:      $S_{\text{single}} \leftarrow \{R \in C \mid |R| = 1\}$
- 17:      $\mathcal{H} \leftarrow \{R \subseteq S_{\text{single}} \mid R \text{ is 2-hop connected in } G\}$
- 18:      $\mathcal{G}_s \leftarrow S_{\text{single}} \setminus \bigcup_{R \in \mathcal{H}} R$
- 19:     **for** each 2-hop component  $R$  of  $\mathcal{H}$  **do**
- 20:        $C \leftarrow C \cup \begin{cases} \{R\}, & |R| \leq M \\ \text{SPLIT-2HOP}(G, R, M) & \text{otherwise} \end{cases}$
- 21:     **end for**
- 22:   **end for**
- 23:    $Q \leftarrow Q'$ ,  $\ell \leftarrow \ell + 1$
- 24: **end while**
- 25: Attach each of  $\mathcal{G}_s$  to neighboring clusters in  $C$
- 26: **return**  $C$

---

Algorithm 1 presents the resulting **Residual-aware  $k$ -core Hierarchy** heuristic, RkH in short. It takes two parameters: the input graph ( $G$ ), and the maximum cluster size ( $M$ ). RkH begins by extracting the largest connected component and removing self-loops (line 1), as done in GraphRAG [7]. It then computes the core number for each node (line 2). The hierarchy construction is initialized by setting  $\ell = 1$ , enqueueing full node set  $V$  into  $Q$ , and initializing an empty cluster set  $C$  and global singleton set  $\mathcal{G}_s$  (lines 3–4).

The graph is then processed iteratively over increasing core levels  $\ell$  until the maximum core is reached (line 5), splitting each cluster  $S$  into core ( $c(v) \geq \ell$ ) and residual ( $c(v) < \ell$ ) nodes (line 8). Core components smaller than or equal to  $M$  are added directly to  $C$  and the queue; larger components are split into size-bounded clusters via SPLIT before adding (lines 9–11). The maximum cluster size ( $M$ ) is set by dividing the input token limit by the average tokens per node, providing a practical estimate that keeps most clusters within the allowed input size. Residual nodes are handled similarly, either added directly or SPLIT as needed (lines 13–14); since these residual subgraphs correspond to leaf clusters, they are not pushed to the next-level queue. The SPLIT procedure breaks each oversized cluster into smaller, size-bounded clusters by greedily growing each cluster from a high-degree seed node, prioritizing neighbors that maximize internal connectivity. This ensures that dense regions remain intact while respecting the maximum cluster size; full algorithmic details are provided in Algorithm 3 in Appendix B.1.

Now, we extract singleton clusters from the existing set  $C$  (line 16) into  $S_{\text{single}}$  and form new clusters  $\mathcal{H}$  consisting of nodes that are 2-hop connected in  $G$  (line 17). A set  $R$  is 2-hop connected if each node  $u \in R$  has a path of length  $\leq 2$  in  $G$  to some other node  $v \in R$ . Any remaining singleton nodes not included in  $\mathcal{H}$  are added to global set  $\mathcal{G}_s$  (line 18). Now, each 2-hop component from  $\mathcal{H}$  is either added directly to  $C$  or split into smaller clusters via SPLIT-2HOP if it exceeds size  $M$  (lines 19–20).

The SPLIT-2HOP procedure greedily splits a 2-hop connected component, ensuring that nodes sharing common anchors are grouped together. Starting from a seed node with the largest anchor set, the algorithm greedily grows a cluster by adding nodes that share anchors with the current cluster, continuing until the cluster reaches the maximum size or no eligible nodes remain. Anchor nodes connected to multiple cluster members are then included to maintain structural coherence. The process produces size-bounded clusters that respect both 2-hop connectivity and anchor relationships. Full details are provided in Algorithm 4 in Appendix B.2.

After all clusters at the current level are processed, the queue is updated with the newly formed clusters, and the level counter is incremented (line 23). Once all core levels have been processed, a final attachment step assigns each node in  $\mathcal{G}_s$  to a neighboring cluster in  $C$ , ensuring that no isolated nodes remain in the hierarchy (line 25). Finally, the algorithm returns the hierarchical cluster set  $C$  (line 26). This approach ensures size-bounded, connectivity-preserving clusters while capturing both dense core structures and sparse residual nodes.

Figure 1 shows an example hierarchy produced by RkH. The 1-core contains all nodes (a–p). At the next level, the 2-core consists of nodes (f–p), while the remaining nodes (a–e) form different clusters, including a connected component (a–b), a 2-hop connected group (c–d), and a singleton (e). At the 3-core level, only nodes (m–p) remain, and the residual nodes (f–l) are partitioned into three connected components (f–h), (i–j), and (k–l). Note that the singleton (e) is attached to its neighboring cluster (f–h) at the end. This recursive decomposition yields a hierarchy where higher  $k$ -cores capture denser subgraphs and residuals are organized into progressively finer-grained structures, as formalized in Algorithm 1.

## 4.2 Handling Small Clusters

Due to the sparse nature of knowledge graphs, many tiny clusters (often containing only two nodes) can arise from previous heuristic. As a result, when GraphRAG evaluates community relevance during query answering, these small clusters usually receive low scores. Consequently, they are often excluded from the final answer generation, which can affect overall sensemaking performance. We propose to explicitly merge such small clusters.

**Merging 2-hop Small Clusters.** In RkH, some 2-hop clusters generated in lines 19–20 may contain only two nodes. These very small clusters can fragment the hierarchy and reduce connectivity if left unprocessed. To address this, we apply a post-processing step after Algorithm 1 that merges small 2-hop clusters into larger clusters whenever possible, or creates new clusters when no suitable neighbors exist. The merging procedure is detailed in Algorithm 2,

---

### Algorithm 2: M2hC : Merge 2-hop Clusters ( $G, C$ )

---

**Require:** Graph  $G = (V, E)$ , existing cluster set  $C$

**Ensure:** Updated cluster set  $C$  with small clusters merged

```

1:  $\mathcal{S} \leftarrow$  All 2-hop clusters of size 2 in  $C$ 
2:  $\mathcal{L} \leftarrow C \setminus \mathcal{S}$ 
3: while  $\mathcal{S} \neq \emptyset$  do
4:   Pick small cluster  $C_s \in \mathcal{S}$  with the most neighbors in  $\mathcal{L}$ 
5:   Count neighbors of  $C_s$  in each cluster of  $L \in \mathcal{L}$ 
6:   if any neighbors exist then
7:     Find  $L_{best} \in \mathcal{L}$  with most neighbors of  $C_s$ 
8:      $L_{best} \leftarrow L_{best} \cup C_s$ 
9:   else
10:     $\mathcal{L} \leftarrow \mathcal{L} \cup C_s$ 
11:   end if
12:   Remove  $C_s$  from  $\mathcal{S}$ 
13: end while
14: return  $\mathcal{L}$ 

```

---

referred to as M2hC. First, all 2-hop clusters of size two are separated from the remaining clusters (lines 1–2). While small clusters remain (line 3), the algorithm selects the small cluster with the most connections to existing clusters (line 4) and counts its neighbors in each cluster (line 5). If neighbors exist, the small cluster is merged in-place into the cluster with which it shares the most edges (lines 6–8). Otherwise, a new cluster is created for the small cluster (line 9). After each iteration, the processed small cluster is removed from the pool (line 12). Finally, the algorithm returns the updated cluster set with all small clusters integrated (line 14).

**Merging Residual Clusters.** In addition to the previous heuristics, we observe that many residual connected components produced during RkH (lines 13–14) have size two. To address this, we introduce an extended heuristic, referred to as MRC. MRC is conceptually similar to M2hC, but instead of focusing solely on 2-hop clusters, it handles all small residual clusters of size two. The only modification occurs in line 1 of Algorithm 2, where both residual clusters and 2-hop clusters of size two are collected into the set  $\mathcal{S}$  for processing. The rest of the procedure remains unchanged: each small cluster is merged into a neighboring parent cluster when possible; otherwise, a new cluster is created only if no suitable neighbors exist.

Note that, for M2hC and MRC, we focus specifically on clusters of size two, as they account for the majority of small clusters observed in practice. Clusters of size three or larger occur infrequently, and trying to merge would not yield measurable improvements; moreover, merging them would unnecessarily inflate neighboring clusters and weaken size constraints.

## 4.3 Token Efficiency via Sampling

In large hierarchical graph decompositions, leaf-level communities can be densely connected, and many nodes/edges within the same community frequently carry overlapping information. Passing all of this content to an LLM for retrieval or summarization can therefore be costly in terms of tokens, leading to inefficiency and redundancy. To address this challenge, we propose the *Round-Robin Token-Constrained Selection* (RRTC) heuristic, which efficiently reduces token usage while preserving the most informative edges

**Table 1: Summary statistics of datasets used in our experiments, showing model, dataset, number of documents, and graph information ( $|V|$ : nodes,  $|E|$ : edges,  $K_m$ : maximum core number). Post-cutoff data is used for GPT-3.5-turbo (cutoff: Sep 1, 2021) and GPT-4o-mini (cutoff: Oct 1, 2023), while full data is used for GPT-5-mini.**

Model	Dataset	# Docs	$ V $	$ E $	$K_m$
GPT-3.5-turbo (post-cutoff)	podcast	36	1791	2651	6
	news	609	10655	15368	6
	semiconductor	183	4324	9546	13
GPT-4o-mini (post-cutoff)	podcast	13	628	882	5
	news	331	5312	7316	6
	semiconductor	63	2029	4249	9
GPT-5-mini (full)	podcast	72	3323	7761	7
	news	609	4499	7941	8
	microsoft	45	790	1702	7

from each community. RRTC operates on leaf-level communities produced by any hierarchy heuristic (e.g., RkH, M2hC, or MRC) and selects a representative subset of edges under a fixed token budget. Within each community, edges are ranked by the combined degree of their endpoints to reflect overall prominence. Selection then proceeds in a round-robin fashion across leaf-level communities, traversing from higher to lower  $k$ -shells, until the token budget is exhausted. By sampling a representative subset of edges from each community, RRTC captures the essential information without including all nodes/edges. Our design is motivated by the  $k$ -shell-based influential node identification methods [14, 34], where round-robin selection across  $k$ -shells identifies diverse key nodes.

## 5 Experimental Setup

In this section, we describe the experimental setup, including the datasets and evaluation approach. Since our study targets the global sensemaking task, we selected three datasets requiring reasoning over multiple documents. The first two align with the original study of Edge et al., while the third also requires multi-document reasoning. Dataset sizes range from approximately 1M to 6M tokens, all processed using the same knowledge graph construction pipeline used by [7]. Details of these datasets are reported in Table 1.

**Podcast transcripts (podcast).** Public transcripts of *Behind the Tech with Kevin Scott*, a podcast featuring conversations between Microsoft CTO Kevin Scott and various thought leaders in science and technology [26]. The corpus contains 72 documents with approximately 1 million tokens in total.

**News articles (news).** A benchmark dataset of news articles spanning multiple categories, including entertainment, business, sports, technology, health, and science [31]. It is a multi-document QA dataset containing 609 documents and nearly 1.4 million tokens.

**S&P 500 earnings transcripts.** This dataset consists of earnings call transcripts from publicly listed S&P 500 companies spanning 2013–2025 [3], with over 20,000 documents totaling 232M tokens. For GPT-3.5-turbo and GPT-4o-mini, to focus on multi-document synthesis and keep the dataset manageable, we use only post-cutoff transcripts from the Semiconductor industry (**semiconductor**), yielding 183 and 63 documents, respectively. For GPT-5-mini, since the full semiconductor dataset is large (nearly 6M tokens), we use Microsoft earnings calls from 2013–2024 (**microsoft**) to limit corpus size.

## 5.1 Evaluation Criteria

Following the methodology of [7], we generated 125 sensemaking questions using GPT-5-mini (details are in Appendix D.1). Consistent with prior work [7, 16, 32], we evaluate answers using two primary criteria: **Comprehensiveness**, which measures how thoroughly an answer addresses all relevant aspects of a question, and **Diversity**, which measures the extent to which an answer captures varied perspectives and insights. Edge et al. [7] also report *Empowerment* (ability to support informed judgment) and *Directness* (degree to which an answer directly addresses the question); results for these metrics are provided in the extended version [13].

**Retrieval approaches and evaluation.** We focus on C2 and C3 from Section 2, which perform best overall, and evaluate all  $k$ -core heuristics—RkH, M2hC, and MRC—against them. Higher  $k$ -core levels typically yield finer-grained communities, improving answer quality via more detailed summaries (details are provided in the extended version [13]). Hence, for all heuristics, we consider the Leaf and immediate parent levels, labeled LF and L1 (e.g., RkH LF). For evaluation, we use a head-to-head framework where LLMs compare answers to select a winner, loser, or tie. To improve reliability, we use multiple evaluators and perform repeated assessments, with majority voting determining the final outcome—suitable for global sensemaking tasks without gold-standard references (see Appendix D.2 for details).

**Configuration.** For fair comparison with [7], we performed indexing using a 600-token chunking window with 100-token overlap and an 8k-token context window for community summaries, following their approach. We utilized the public Azure OpenAI endpoint and the OpenAI API for GPT-3.5-turbo, GPT-4o-mini, and GPT-5-mini. For the multiple-judge evaluation, GPT-5-mini was accessed via Azure OpenAI, while the other four LLMs (Gemini 3 Pro Preview, Gemini 2.5 Pro, Qwen3 Next 80B, and DeepSeek v3.2) were accessed via GCP. All LLM-based evaluations were performed on a GCP virtual machine equipped with 32 GB memory and 16 vCPUs (Intel Haswell architecture). Detailed prompts for graph construction, community summaries, and global answer generation (following [7]), along with our code, are available here: <https://github.com/erdemUB/KDD26>.

## 6 Results and Analysis

Knowledge cutoff is a key consideration in RAG evaluation [28, 38], as models may have already seen the corpus if it predates their training data. To minimize such contamination, we prioritize **post-cutoff data** when selecting both datasets and evaluation models.

Our primary evaluation (Section 6.1) uses GPT-3.5-turbo and GPT-4o-mini across three post-cutoff datasets: `podcast`, `news`, and `semiconductor`. This setup ensures a fair comparison of GraphRAG configurations under realistic retrieval and sensemaking conditions, while minimizing the influence of memorized knowledge. To further verify that our heuristics generalize to more recent, stronger models, we conduct a secondary evaluation using GPT-5-mini on the full datasets, to account for the limited post-cutoff content



**Table 3: GPT-5-mini full: Head-to-head win rates (%), for comprehensiveness and diversity metrics against Leiden C2 level from Edge et al. [7]. LF indicates leaf-level communities, and L1 indicates the level immediately above the leaf.**

GPT-5-mini Comprehen.	podcast			news			microsoft		
	Win	Loss	Tie	Win	Loss	Tie	Win	Loss	Tie
RkH L1	41	47	12	56	42	2	43	43	14
RkH LF	50	48	2	50	45	5	45	40	15
M2hC L1	40	52	8	40	34	26	40	45	15
M2hC LF	52	42	6	40	40	20	48	50	2
MRC L1	48	50	2	42	38	20	45	48	7
MRC LF	54	44	2	42	36	22	45	35	20
Diversity	Win	Loss	Tie	Win	Loss	Tie	Win	Loss	Tie
RkH L1	40	44	16	50	44	6	48	50	2
RkH LF	50	48	2	45	43	12	38	45	17
M2hC L1	37	51	12	40	38	22	45	38	17
M2hC LF	56	36	8	40	37	23	52	46	2
MRC L1	48	50	2	36	40	24	55	42	3
MRC LF	54	44	2	37	37	26	40	43	17

LLM judges for evaluation, while we strengthen our evaluation considering multiple judges/runs, randomize answer order to reduce bias, and apply post-cutoff filtering to mitigate memorization effects. Additionally, we conducted a small human evaluation on 25 randomly selected questions per dataset on the GPT-3.5 setup. The results show that human judgments agree with the LLM majority vote in 96% of cases overall (Cohen’s  $k = 0.94$ ), confirming that LLM-based comparisons reliably reflect answer quality.

**Sensitivity Analysis of the Maximum Cluster Size.** We performed a sensitivity analysis on the podcast dataset to evaluate the effect of the maximum cluster size parameter,  $M$ , which is set heuristically as explained in Section 4 to ensure that important information is not truncated. +5/+10 change in  $M$  reduces the win rate by 2–3%, likely due to context overflow, while  $-5/-10$  change in  $M$  increases the number of clusters and LLM calls without affecting overall performance. These results validate our choice of  $M$ .

**Runtime Analysis.** Graph construction is identical across methods; runtime differences arise from community/hierarchy construction and the number of LLM calls, with more communities increasing cost. We measured wall-clock indexing time across the three datasets and find that  $k$ -core reduces hierarchy construction overhead by approximately 31% on average compared to Leiden, demonstrating its practical efficiency. Additionally, as each LLM call has a similar cost, total runtime scales with the number of communities. Since  $k$ -core yields approximately 30–35% fewer LLM calls on average (see Table 5), it further improves end-to-end efficiency.

## 6.2 Evaluation on Full Data by GPT-5-mini

In our secondary experiments, we use GPT-5-mini to examine whether the trends observed with earlier models persist with a stronger, more recent evaluator. Due to GPT-5-mini’s later training window, none of the datasets have sufficient post-cutoff content. Hence, we evaluate on the full news, podcast, and microsoft (since semiconductor would create an excessively large corpus).

Table 3 reports head-to-head results between Leiden C2 and our  $k$ -core heuristics; C3 results are given in the extended version [13]. While our heuristics continue to show improvements, the high proportion of ties suggests that prior knowledge reduces the

**Table 4: Wilcoxon signed-rank test  $p$ -values for comprehensiveness (GPT-3.5-turbo) across datasets and baselines. Values with  $p < 0.005$  are highlighted in bold.**

Cond.	podcast		news		semiconductor	
	C2	C3	C2	C3	C2	C3
RkH L1	<0.001	<0.001	1.000	<0.001	1.000	<0.001
RkH LF	<0.001	<0.001	<0.001	0.370	0.982	0.363
M2hC L1	<0.001	<0.001	<0.001	<0.001	<0.001	0.006
M2hC LF	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
MRC L1	0.007	1.000	<0.001	0.069	<0.001	0.073
MRC LF	0.982	<0.001	0.007	1.000	<0.001	<0.001

discriminative power of head-to-head evaluation. Across datasets and metrics, leaf-level variants generally outperform L1 by 2–6 percentage points, with M2hC LF showing the strongest gains for diversity on podcast and MRC LF performing best for comprehensiveness on microsoft. While M2hC LF still achieves the strongest overall performance, MRC LF edges out M2hC LF on podcast and microsoft. Overall win rates remain around 45–55%, reflecting narrower margins than with GPT-3.5-turbo and slightly wider than with GPT-4o-mini. Several L1 configurations fall below 50%, indicating that gains are concentrated in leaf-level variants. Nevertheless, the directional consistency with earlier evaluators suggests that  $k$ -core-based heuristics retain an advantage, albeit a modest one, even under a stronger model whose prior knowledge reduces the discriminative power of head-to-head evaluation.

## 6.3 Statistical Analysis

To assess statistical significance, we follow the procedure of [7] and report  $p$ -values for pairwise comparisons of each heuristic against C2/C3 baselines on the comprehensiveness metric using GPT-3.5-turbo post-cutoff data in Table 4 (details and diversity results are in the extended version [13]). Values with  $p < 0.005$  are highlighted in bold. Overall, RkH and M2hC are significant on podcast and news, while M2hC and MRC perform best on semiconductor. Notably, M2hC LF is the only heuristic significant ( $p < 0.005$ ) against both C2 and C3 across all datasets.

**Summary.** The results above show that the performance gap between our heuristics and Leiden-based C2/C3 is systematic rather than incidental. Knowledge graphs exhibit extremely sparse graph structures; e.g., average degrees in our datasets range from 2.88 to 4.42 and 55–60% of nodes have a degree of just 1. This high prevalence of degree-1 nodes means that the majority of nodes have only a single connection, providing insufficient local structure for Leiden’s modularity optimization to form meaningful communities, as the algorithm relies on dense local neighborhoods to effectively aggregate nodes into clusters. As Theorem 1 shows, Leiden’s degeneracy on such sparse graphs makes it unlikely to find the “right” partition, as there are exponentially many near-equivalent solutions, most of which group peripheral nodes incoherently. In contrast,  $k$ -core-based heuristics produce a deterministic hierarchy that is structurally grounded, explaining the consistent advantage of our methods across GPT-3.5-turbo, GPT-4o-mini, and GPT-5-mini evaluations. Note that nearly 45% of clusters from RkH have size 2, so M2hC and MRC help improve performance by merging these, reducing fragmentation. Clusters of size 3 or 4 account for

**Table 5: Community statistics across datasets, showing the number of communities and the percentage of source text used at different hierarchy levels and methods.**

Dataset	# Tokens	Metric	C2	C3	RkH LF	M2hC LF	MRC LF
podcast	180,267	# Comm.	291	301	351	311	194
		Coverage (%)	68.93	70.87	83.70	75.47	58.08
news	1,340,078	# Comm.	2,019	2,021	2,205	1,863	1,357
		Coverage (%)	65.90	71.26	76.46	65.74	55.08
semic.	438,819	# Comm.	627	696	847	705	468
		Coverage (%)	66.81	72.67	90.18	76.45	59.94

only 10–12%, and merging them offers little benefit while inflating neighboring clusters, so they were excluded from our small-cluster definition. Under GPT-3.5-turbo, these differences are statistically significant across all three datasets ( $p < 0.005$  for M2hC LF), and directionally consistent under GPT-4o-mini and GPT-5-mini, where stronger prior knowledge narrows the margins but does not eliminate the advantage.

#### 6.4 Token Utilization and Impact of RRTC

Finally, we analyze how the evaluated methods consume tokens, and how our round-robin token-constrained selection (RRTC) mechanism (see Section 4.3) performs. We first analyze the number of communities produced by each heuristic, along with the percentage of source tokens contained in those communities. These are the communities and tokens that are used to answer queries via LLM calls: fewer communities reduce the number of LLM calls, and lower token counts reduce overall LLM token usage. Results on GPT-3.5-turbo post-cutoff data are shown in Table 5 for leaf-level, which typically contain more communities and use more tokens than L1. Notably, MRC reduces both the number of communities and the tokens used: it consistently produces the fewest communities and lowest coverage (55–60%), meaning it aggressively consolidates information. These reductions indicate that our heuristics use fewer tokens and require fewer LLM calls, while achieving comparable or improved performance on global sensemaking, as shown in the previous section.

#### Impact of Round-Robin Token-Constrained Selection (RRTC).

Table 6 reports results for M2hC LF with RRTC across three datasets, compared against Leiden C2/C3 under three edge budgets: 80%, 70%, and 60% of the original graph. The table reports token usage relative to C2/C3 and head-to-head win rates against them for comprehensiveness and diversity, where win rates are computed against losses only (ties excluded); thus, a 55% win rate implies a 45% loss rate.

Across datasets, RRTC maintains competitive performance even when selecting only 60% of graph edges. Comprehensiveness win rates remain in the 50–56% range, with the highest values on semiconductor at the 80% budget, reaching 56% and 58% against C2 and C3, respectively. Diversity is similarly preserved, often exceeding 55% on podcast and semiconductor despite substantial edge reduction. As the edge budget decreases from 80% to 60%, win rates gradually decline, with comprehensiveness on podcast dropping below 50%, indicating that a 60% budget may be too aggressive for some dataset–metric combinations. Overall, RRTC reduces token usage by up to ~40% relative to Leiden C2/C3 while maintaining competitive comprehensiveness and diversity, providing a practical, token-efficient alternative under moderate budget constraints.

**Table 6: RRTC (M2hC Leaf) vs. Leiden C2/C3: Head-to-head win rates (%) on comprehensiveness and diversity**

Dataset	Edge budget (%)	Relative token usage (%)		Comp. win rate (%)		Div. win rate (%)	
		C2	C3	C2	C3	C2	C3
podcast	80	90	88	57	55	64	55
	70	79	77	54	52	56	52
	60	68	66	47	46	51	49
news	80	81	75	54	54	55	58
	70	71	65	49	54	52	56
	60	60	56	52	50	46	50
semic.	80	92	85	56	58	55	60
	70	81	74	53	53	50	57
	60	69	64	55	50	52	54

## 7 Conclusion

We proved that modularity optimization is inherently unreliable on the sparse graphs typical of GraphRAG knowledge graphs, and proposed  $k$ -core-based heuristics that yield deterministic, size-bounded hierarchies in linear time. Across three datasets, three generator LLMs, and five LLM judges, our approach consistently improves comprehensiveness and diversity while reducing token usage compared to Leiden-based GraphRAG. Future work includes extending the framework to dynamic knowledge graphs and exploring its applicability beyond global sensemaking tasks. In addition, alternative edge-selection mechanisms for token-budgeted retrieval and summarization, including PageRank-based methods and approaches inspired by [14], remain an interesting direction for future research.

## Acknowledgments

Hossain and Saryüce were supported by NSF-2107089 and NSF-2236789 awards. This work used resources from the Center for Computational Research at the University at Buffalo [5] and the SUNY AI Platform on Google Cloud [30].

## References

- [1] Vladimir Batagelj and Matjaž Zaversnik. 2011. Fast algorithms for determining (generalized) core groups in social networks. *Advances in Data Analysis and Classification* 5, 2 (2011), 129–145.
- [2] Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. Query focused abstractive summarization: Incorporating query relevance, multi-document coverage, and summary length constraints into seq2seq models. *arXiv preprint arXiv:1801.07704* (2018).
- [3] Michele Ca’ Zorzi, Gianluigi Lopardo, and Ana-Simona Manu. 2025. Verba Volant, Transcripta Manent: What Corporate Earnings Calls Reveal About the AI Stock Rally. 3093 (2025). <https://glopardo.com/corporatetalks/>
- [4] Shai Carmi, Shlomo Havlin, Scott Kirkpatrick, Yuval Shavitt, and Eran Shir. 2007. A model of Internet topology using  $k$ -shell decomposition. *Proceedings of the National Academy of Sciences* 104, 27 (2007), 11150–11154.
- [5] Center for Computational Research, University at Buffalo. 2026. <http://hdl.handle.net/10477/79221>. Accessed: 2026-05-28.
- [6] J. Cohen. 2008. Trusses: Cohesive subgraphs for social network analysis. National Security Agency Technical Report (2008).
- [7] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130* (2024).
- [8] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-Augmented Generation for Large Language Models: A Survey. *arXiv:2312.10997* [cs.CL] <https://arxiv.org/abs/2312.10997>
- [9] Abdellah Ghassel, Ian Robinson, Gabriel Tanase, Hal Cooper, Bryan Thompson, Zhen Han, Vassilis Ioannidis, Soji Adeshina, and Huzefa Rangwala. 2025. Hierarchical Lexical Graph for Enhanced Multi-Hop Retrieval. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2* (Toronto

- ON, Canada) (KDD '25). Association for Computing Machinery, New York, NY, USA, 4457–4466. <https://doi.org/10.1145/3711896.3737233>
- [10] Benjamin H Good, Yves-Alexandre De Montjoye, and Aaron Clauset. 2010. Performance of modularity maximization in practical contexts. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 81, 4 (2010), 046106.
- [11] Haoyu Han, Yu Wang, Harry Shomer, Kai Guo, Jiayuan Ding, Yongjia Lei, Mahantesh Halappanavar, Ryan A. Rossi, Subhabrata Mukherjee, Xianfeng Tang, Qi He, Zhigang Hua, Bo Long, Tong Zhao, Neil Shah, Amin Javari, Yinglong Xia, and Jiliang Tang. 2025. Retrieval-Augmented Generation with Graphs (GraphRAG). arXiv:2501.00309 [cs.LG] <https://arxiv.org/abs/2501.00309>
- [12] Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, and Bryan Hooi. 2024. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. *Advances in Neural Information Processing Systems* 37 (2024), 132876–132907.
- [13] Jakir Hossain and Ahmet Erdem Sariyüce. 2026. Core-based Hierarchies for Efficient GraphRAG. arXiv:2603.05207 [cs.LG] <https://arxiv.org/abs/2603.05207>
- [14] Jakir Hossain, Sucheta Soundarajan, and Ahmet Erdem Sariyüce. 2023. Quantifying node-based core resilience. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 259–276.
- [15] Yuntong Hu, Zhihan Lei, Zheng Zhang, Bo Pan, Chen Ling, and Liang Zhao. 2025. GRAG: Graph Retrieval-Augmented Generation. arXiv:2405.16506 [cs.LG] <https://arxiv.org/abs/2405.16506>
- [16] Boyoung Kim, Dosung Lee, Sumin An, Jinseong Jeong, and Paul Hongsuck Seo. 2025. ReTAG: Retrieval-Enhanced, Topic-Augmented Graph-Based Global Sensemaking. In *Findings of the Association for Computational Linguistics: EMNLP 2025*. 22249–22277.
- [17] Jiho Kim, Yeonsu Kwon, Yohan Jo, and Edward Choi. 2023. Kg-gpt: A general framework for reasoning on knowledge graphs using large language models. *arXiv preprint arXiv:2310.11220* (2023).
- [18] Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. 2024. In search of needles in a 11m haystack: Recurrent memory finds what llms miss. *arXiv preprint arXiv:2402.10790* (2024).
- [19] Md Tahmid Rahman Laskar, Enamul Hoque, and Jimmy Huang. 2020. Query focused abstractive summarization via incorporating query relevance and transfer learning with transformer models. In *Canadian conference on artificial intelligence*. Springer, 342–348.
- [20] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* 33 (2020), 9459–9474.
- [21] Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the association for computational linguistics* 12 (2024), 157–173.
- [22] Fatma Miladi, Valéry Psyché, and Daniel Lemire. 2024. Leveraging gpt-4 for accuracy in education: A comparative study on retrieval-augmented generation in moocs. In *International Conference on Artificial Intelligence in Education*. Springer, 427–434.
- [23] Tiago P Peixoto. 2014. Hierarchical block structures and high-resolution model selection in large networks. *Physical Review X* 4, 1 (2014), 011047.
- [24] Boci Peng, Yun Zhu, Yongchao Liu, Xiaohu Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2025. Graph retrieval-augmented generation: A survey. *ACM Transactions on Information Systems* 44, 2 (2025), 1–52.
- [25] Ahmet Erdem Sariyüce and Ali Pinar. 2016. Fast hierarchy construction for dense subgraphs. 10, 3 (Nov. 2016), 97–108.
- [26] Kevin Scott. 2024. Behind the Tech. <https://www.microsoft.com/en-us/behind-the-tech>. Accessed: 2026-01-01.
- [27] Stephen B Seidman. 1983. Network structure and minimum degree. *Social networks* 5, 3 (1983), 269–287.
- [28] Minjae Seo, Wonwoo Choi, Myoungsung You, and Seungwon Shin. 2025. AutoPatch: Multi-Agent Framework for Patching Real-World CVE Vulnerabilities. *arXiv preprint arXiv:2505.04195* (2025).
- [29] Jiashuo Sun, Chengjin Xu, Luminyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M Ni, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph. *arXiv preprint arXiv:2307.07697* (2023).
- [30] SUNY AI Platform on Google Cloud. 2026. <https://ai-platform.suny.edu/>. Accessed: 2026-05-28.
- [31] Yixuan Tang and Yi Yang. 2024. Multihop-rag: Benchmarking retrieval-augmented generation for multi-hop queries. *arXiv preprint arXiv:2401.15391* (2024).
- [32] Yash Tiwari, Owais Ahmad Lone, and Mayukha Pal. 2025. OntoRAG: Enhancing Question-Answering through Automated Ontology Derivation from Unstructured Knowledge Bases. *arXiv preprint arXiv:2506.00664* (2025).
- [33] Vincent A Traag, Ludo Waltman, and Nees Jan Van Eck. 2019. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific reports* 9, 1 (2019), 1–12.
- [34] Min Wang, Wanchun Li, Yuning Guo, Xiaoyan Peng, and Yingxiang Li. 2020. Identifying influential spreaders in complex networks based on improved k-shell method. *Physica A: Statistical Mechanics and its Applications* 554 (2020), 124229.
- [35] Yu Wang, Nedim Lipka, Ryan A Rossi, Alexa Siu, Ruiyi Zhang, and Tyler Derr. 2024. Knowledge graph prompting for multi-document question answering. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 38. 19206–19214.
- [36] Nirmalie Wiratunga, Ramitha Abeyratne, Lasal Jayawardena, Kyle Martin, Stewart Massie, Ikechukwu Nkisi-Orji, Ruwan Weerasinghe, Anne Liret, and Bruno Fleisch. 2024. CBR-RAG: case-based reasoning for retrieval augmented generation in LLMs for legal question answering. In *International Conference on Case-Based Reasoning*. Springer, 445–460.
- [37] Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Bowen Jin, May Dongmei Wang, Joyce Ho, and Carl Yang. 2024. Ram-ehr: Retrieval augmentation meets clinical predictions on electronic health records. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 754–765.
- [38] Zhuoyi Yang, Yurun Song, Iftekhar Ahmed, and Ian Harris. 2026. Fine-Tuning vs. RAG for Multi-Hop Question Answering with Novel Knowledge. *arXiv preprint arXiv:2601.07054* (2026).
- [39] Jin-ge Yao, Xiaojun Wan, and Jianguo Xiao. 2017. Recent advances in document summarization. *Knowledge and Information Systems* 53, 2 (2017), 297–336.
- [40] Boyu Zhang, Hongyang Yang, Tianyu Zhou, Muhammad Ali Babar, and Xiaoyang Liu. 2023. Enhancing financial sentiment analysis via retrieval augmented large language models. In *Proceedings of the fourth ACM international conference on AI in finance*. 349–356.
- [41] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* 1, 2 (2023).

## A Proof of Theorem 1 and Experimental Validation

### Proof sketch.

**Step 1: Low-degree nodes are weakly coupled.** A node with degree  $k_i \leq d$  participates in at most  $d$  edges. Moving it between communities changes modularity by at most  $O(d/m)$ : the adjacency term shifts by at most  $d/m$ , and the null-model penalty shifts by  $O(k_i \cdot K_r / (2m)^2)$  where  $K_r$  is the total degree of the destination community. In the sparse regime where  $m = \Theta(n)$ , this is  $O(1/n)$ , so the node’s community assignment barely affects the objective.

**Step 2: Weakly coupled nodes can be independently reassigned.** Reassigning one low-degree node perturbs another’s sensitivity by only  $O(1/n^2)$ . By selecting a large independent set among the low-degree nodes, which is straightforward since these nodes have bounded degree, we obtain  $\Theta(n)$  nodes whose reassignments are approximately independent. Each can be placed in at least two communities without meaningful modularity loss, yielding  $2^{\Theta(n)}$  distinct near-optimal partitions.

**Step 3: Sparsity makes this severe.** When  $\bar{k} < 3$ , standard results on Poisson and power-law degree distributions guarantee that  $n_{\leq d} = \Theta(n)$ . This is the regime of knowledge graphs in GraphRAG. In contrast, for dense graphs with large  $\bar{k}$ ,  $n_{\leq d}/n \rightarrow 0$  and most nodes are strongly coupled to their communities, collapsing the degeneracy.

### Full proof.

**PROOF.** We first remind modularity in its standard community-sum form:

$$Q(\sigma) = \sum_c \left[ \frac{e_c}{m} - \left( \frac{K_c}{2m} \right)^2 \right],$$

where  $e_c$  counts edges internal to community  $c$  and  $K_c = \sum_{j \in c} k_j$  is its total degree. Moving node  $i$  from community  $s$  to community  $r$  produces a new partition  $\sigma^{i \rightarrow r}$ . We write  $d_i^{(c)}$  for the number

of neighbors of  $i$  in community  $c$ . We also define the *modularity sensitivity* of node  $i$  under partition  $\sigma$  is  $\Delta_i(\sigma) = \max_{r \neq \sigma_i} |Q(\sigma) - Q(\sigma^{i \rightarrow r})|$ , where  $\sigma^{i \rightarrow r}$  moves  $i$  to community  $r$  with all other assignments fixed. Node  $i$  is  $\varepsilon$ -weakly coupled if  $\Delta_i(\sigma) < \varepsilon$ .

Next, we give two lemmas that will help with the proof

**LEMMA 2 (MOVING A LOW-DEGREE NODE).** *For any partition  $\sigma$  and any node  $i$  with degree  $k_i \leq d$ , moving  $i$  to any other community changes modularity by at most*

$$\Delta_i(\sigma) \leq \frac{2k_i}{m} + \frac{k_i^2}{2m^2} = O\left(\frac{1}{n}\right).$$

**PROOF.** Moving  $i$  from community  $s$  to  $r$  affects only those two terms in the sum. There are two contributions:

*Edge term.* Community  $s$  loses  $d_i^{(s)}$  internal edges;  $r$  gains  $d_i^{(r)}$ . Both are at most  $k_i$ , so

$$|\Delta_{\text{edge}}| = \frac{|d_i^{(r)} - d_i^{(s)}|}{m} \leq \frac{k_i}{m}.$$

*Degree-penalty term.* The total degree of  $s$  drops by  $k_i$  and that of  $r$  rises by  $k_i$ . Expanding the squares and simplifying:

$$|\Delta_{\text{penalty}}| = \frac{|k_i[2(K_r - K_s) + 2k_i]|}{(2m)^2} \leq \frac{k_i}{m} + \frac{k_i^2}{2m^2},$$

using  $|K_r - K_s| \leq 2m$ .

Adding the two and substituting  $k_i \leq d$  and  $m = n\bar{k}/2 = \Theta(n)$  gives  $\Delta_i \leq 2d/m + d^2/(2m^2)$ . Both terms are  $O(1/n)$ .  $\square$

**LEMMA 3 (NON-ADJACENT LOW-DEGREE NODES).** *Let  $i, j$  be non-adjacent nodes with degrees  $\leq d$ . Reassigning  $j$  changes node  $i$ 's sensitivity by at most*

$$|\Delta_i(\sigma) - \Delta_i(\sigma')| \leq \frac{2d^2}{(2m)^2} = O\left(\frac{1}{n^2}\right).$$

**PROOF.** Since  $i$  and  $j$  share no edge, reassigning  $j$  does not change  $i$ 's neighbor counts  $d_i^{(c)}$  in any community, so the edge term for  $i$  is unaffected. The only effect is that  $j$ 's move shifts the total degrees  $K_r, K_s$  by  $\pm k_j \leq d$ , which perturbs  $i$ 's penalty term by at most  $k_i \cdot 2d/(2m)^2 \leq 2d^2/(2m)^2 = O(1/n^2)$ .  $\square$

Now, we give the proof of Theorem 1. The argument has three parts.

**1. Find a large independent set of weakly coupled nodes.** By Lemma 2, every node with  $k_i \leq d$  has sensitivity  $O(1/n)$ , so for  $\varepsilon > d(2+\bar{k})/(2m)$  it is  $\varepsilon$ -weakly coupled. Let  $n_{\leq d}$  denote the number of such nodes. The subgraph they induce has maximum degree  $d$ , so a greedy algorithm yields an independent set  $S$  of size

$$|S| \geq \frac{n_{\leq d}}{d+1}.$$

**2. Every combination of reassignments stays near-optimal.** Start from an optimal partition  $\sigma^*$ . Each node  $i \in S$  can stay or move to an alternative community: two choices per node. For any subset  $T \subseteq S$  of nodes that are reassigned, the total modularity loss

is bounded by summing the individual sensitivities plus pairwise cross-terms (using Lemma 3):

$$|Q^* - Q(\sigma_T)| \leq \underbrace{\sum_{i \in T} \Delta_i(\sigma^*)}_{\text{individual}} + \underbrace{\sum_{\{i,j\} \subseteq T} \frac{2d^2}{(2m)^2}}_{\text{cross-terms}}.$$

Bounding each sum with  $|T| \leq |S| \leq n_{\leq d}/(d+1)$  and  $m = n\bar{k}/2$ :

$$\text{Individual sum} \leq \frac{n_{\leq d}}{d+1} \cdot \frac{d(2+\bar{k})}{n\bar{k}} \leq \frac{d(2+\bar{k})}{(d+1)\bar{k}} =: C_1,$$

$$\text{Cross-term sum} \leq \frac{1}{2} \left(\frac{n_{\leq d}}{d+1}\right)^2 \cdot \frac{2d^2}{n^2\bar{k}^2} \leq \frac{d^2}{(d+1)^2\bar{k}^2} =: C_2.$$

Both  $C_1$  and  $C_2$  are constants (independent of  $n$ ), so setting  $\varepsilon' = C_1 + C_2$  keeps every such partition within  $\varepsilon'$  of  $Q^*$ .

**3. Count and conclude.** Each of the  $|S|$  nodes has two valid assignments, giving  $2^{|S|}$  structurally distinct near-optimal partitions:

$$\mathcal{D}(\varepsilon') \geq 2^{|S|} \geq 2^{n_{\leq d}/(d+1)}.$$

In the sparse regime ( $\bar{k} < 3$ ), standard degree distributions guarantee  $n_{\leq d} = \Theta(n)$ . For example, under a Poisson distribution with  $\bar{k} = 2$  and  $d = 2$ , the fraction of low-degree nodes is  $5e^{-2} \approx 0.68$ ; under a power law with exponent  $\gamma = 2.5$ , degree-1 nodes alone comprise  $\sim 74\%$  of the graph. Therefore  $\mathcal{D}(\varepsilon') \geq 2^{\Theta(n)}$ .  $\square$

## A.1 Empirical Analysis of Leiden Instability

To complement Theorem 1, we conducted a multi-seed stability analysis (10 runs with different random seeds) of Leiden across 3 datasets (GPT-3.5 setup). Across runs, the cluster statistics show noticeable variability. For example, on the Podcast dataset, the number of clusters ranges from 2550 to 2584, while the largest cluster size ranges between 780 and 1069 nodes. The average node-level Adjusted Rand Index (ARI) is 0.94, indicating around 6% disagreement in node-pair assignments. While moderate, this level of variation is significant in GraphRAG as community assignments directly define retrieval units and contexts for the LLM: the Jaccard similarity of retrieved nodes across seeds is 0.73, on average. These findings provide direct empirical evidence of retrieval-level instability in Leiden clustering and further support Theorem 1.

## B Splitting Oversized Components

### B.1 SPLIT: Splitting Large Connected Components

When a connected component  $R$  exceeds the maximum cluster size  $M$ , it is split into smaller clusters using the SPLIT procedure. The algorithm (Algorithm 3) grows each cluster greedily from a seed node with the highest degree (line 3), adding neighboring nodes that maximize internal connectivity (line 6) until the cluster reaches size  $M$ . Each completed cluster is added to the cluster set (line 10) and removed from the remaining node set (line 11). The output is a set of size-constrained clusters that preserve internal connectivity within  $C$ .

**Algorithm 3:** SPLIT ( $G, R, M$ )

---

**Require:** Graph  $G = (V, E)$ , connected comp.  $R$ , max size  $M$   
**Ensure:** List of split clusters  $C$

- 1: Initialize cluster set  $C \leftarrow \emptyset$
- 2: **while**  $R \neq \emptyset$  **do**
- 3:   Pick seed node  $s \in R$  with highest degree
- 4:    $S \leftarrow \{s\}$ ,  $frontier \leftarrow N(s) \cap R$
- 5:   **while**  $|S| < M$  **and**  $frontier \neq \emptyset$  **do**
- 6:     Pick  $v \in frontier$  maximizing  $N(v) \cap S$
- 7:      $S \leftarrow S \cup v$
- 8:      $frontier \leftarrow N(v) \cap (R \setminus S)$
- 9:   **end while**
- 10:    $C \leftarrow C \cup S$
- 11:   Remove  $S$  from  $R$
- 12: **end while**
- 13: **return**  $C$

---

## B.2 SPLIT-2HOP: Splitting Large 2-Hop Connected Components

When a 2-hop connected component  $H$  exceeds the maximum cluster size  $M$ , it is partitioned into smaller clusters using the SPLIT-2HOP procedure (Algorithm 4). The algorithm first computes the anchor set as the union of neighbors of all 2-hop nodes and precomputes anchor neighborhoods for each node (lines 2–3). Clusters are grown greedily by selecting a seed node with the largest anchor set (line 5) and initializing a cluster with this seed (line 6). Candidate nodes that share at least one anchor with the seed are added to the frontier (line 7). The cluster is expanded by repeatedly selecting the frontier node with the highest overlap in anchor connectivity with the current cluster (lines 8–10), until the cluster reaches size  $M$  or no eligible nodes remain. After cluster growth, anchor nodes connected to at least two nodes in the cluster are selected and included (line 12), and the resulting cluster is added to the output set (line 13). The assigned nodes are then removed from the remaining 2-hop groups set (line 14). This process repeats until all nodes in the 2-hop groups are assigned.

## C Alternative Community Partition Methods

To our knowledge, prior GraphRAG studies do not evaluate alternative community detection methods, making Leiden the de facto comparison baseline. We also explored hierarchical block structures [23] and non-modularity-based approaches such as  $k$ -truss, but found them consistently weaker than Leiden and therefore omitted them from the experimental results (additional discussion on  $k$ -truss is provided in Section 4). For example, on the Semiconductor dataset (GPT-3.5 setup), Leiden outperforms  $k$ -truss in  $\sim 82$ – $85\%$  of comparisons, and  $k$ -core outperforms  $k$ -truss in  $\sim 90$ – $93\%$  of comparisons.

## D Evaluation Criteria Details

### D.1 Question Generation

Following methodology described in [7], we generated 125 sense-making questions using GPT-5-mini. Consistent with their approach, we first prompted GPT-5-mini to create personas of hypothetical users for each corpus and then generated 5 tasks per user.

**Algorithm 4:** SPLIT-2HOP ( $G, H, M$ )

---

**Require:** Graph  $G = (V, E)$ , 2-hop groups  $H$ , max size  $M$   
**Ensure:** List of split clusters  $C$

- 1: Initialize cluster set  $C \leftarrow \emptyset$
- 2: Compute anchors  $A \leftarrow \bigcup_{u \in H} N(u)$
- 3: Precompute  $A_u \leftarrow N(u) \cap A$  for  $u \in H$
- 4: **while**  $H \neq \emptyset$  **do**
- 5:   Pick  $seed \in H$  with max  $|A_{seed}|$
- 6:    $S \leftarrow \{seed\}$
- 7:    $frontier \leftarrow \{v \in H \setminus \{seed\} \mid A_v \cap A_{seed} \neq \emptyset\}$
- 8:   **while**  $|S| < M$  **and**  $frontier \neq \emptyset$  **do**
- 9:     Pick  $v \in frontier$  maximizing  $\sum_{u \in S} |A_v \cap A_u|$
- 10:      $S \leftarrow S \cup \{v\}$ ;
- 11:      $frontier \leftarrow (frontier \cup \{u \in H \mid A_u \cap A_v \neq \emptyset\}) \setminus S$
- 12:   **end while**
- 13:   // Select anchors linked to  $\geq 2$  nodes in  $S$
- 14:    $A' \leftarrow \{v \in A \text{ s.t. } N(v) \cap S \geq 2\}$
- 15:    $C \leftarrow C \cup \{S \cup A'\}$
- 16:    $H \leftarrow H \setminus S$
- 17: **end while**
- 18: **return**  $C$

---

For each user-task pair, we use GPT-5-mini to generate high-level questions that require understanding of the entire corpus without relying on low-level fact retrieval. This process ensures that the resulting questions assess comprehensive, corpus-wide reasoning, aligned with prior work.

### D.2 Evaluation Approach

We evaluate models using a head-to-head framework, where LLMs compare pairs of answers to identify a winner, loser, or tie. To enhance reliability, multiple evaluators are used, and repeated assessments are performed. The final outcome is determined via majority voting, making this approach suitable for global sensemaking tasks without gold-standard references. Each comparison proceeds as follows:

- (1) Generate answers to the same query using two approaches.
- (2) Randomize answer order and present the pair to the evaluator.
- (3) Each of five independent LLM evaluators (GPT-5-mini, Gemini 3 Pro Preview, Gemini 2.5 Pro, Qwen3 Next 80B, and DeepSeek v3.2) assesses the same answer pair three times.
- (4) Apply majority voting within each evaluator to determine their decision.
- (5) Apply a second majority vote across evaluators to determine the final judgment.

This procedure ensures robust, repeatable head-to-head comparisons across methods.