

## Research Article

# SiReRAG: Indexing Similar and Related Information for Multihop Reasoning

Nan Zhang<sup>1,2</sup>, Prafulla Kumar Choubey<sup>2</sup>, Alexander Fabbri<sup>2</sup>, Gabriel Bernadett-Shapiro<sup>2</sup>, Rui Zhang<sup>1</sup>, Prasenjit Mitra<sup>1</sup>, Caiming Xiong<sup>2</sup>, Chien-Sheng Wu<sup>2</sup>

1. Pennsylvania State University, United States; 2. Salesforce AI Research

Indexing is an important step towards strong performance in retrieval-augmented generation (RAG) systems. However, existing methods organize data based on either semantic similarity (similarity) or related information (relatedness), but do not cover both perspectives comprehensively. Our analysis reveals that modeling only one perspective results in insufficient knowledge synthesis, leading to suboptimal performance on complex tasks requiring multihop reasoning. In this paper, we propose SiReRAG, a novel RAG indexing approach that explicitly considers both similar and related information. On the similarity side, we follow existing work and explore some variances to construct a similarity tree based on recursive summarization. On the relatedness side, SiReRAG extracts propositions and entities from texts, groups propositions via shared entities, and generates recursive summaries to construct a relatedness tree. We index and flatten both similarity and relatedness trees into a unified retrieval pool. Our experiments demonstrate that SiReRAG consistently outperforms state-of-the-art indexing methods on three multihop datasets (MuSiQue, 2WikiMultiHopQA, and HotpotQA), with an average 1.9% improvement in F1 scores. As a reasonably efficient solution, SiReRAG enhances existing reranking methods significantly, with up to 7.8% improvement in average F1 scores.

## 1. Introduction

Retrieval-augmented generation (RAG) has shown strong potential in augmenting large language models (LLMs) with highly specialized and constantly updated knowledge<sup>[1][2]</sup>. Getting rid of fine-tuning LLMs, it is an efficient method for handling users' queries that require domain knowledge. A typical RAG pipeline may involve chunking, embedding, indexing, retrieval with queries, reranking, and LLM response generation<sup>[3]</sup>.

The indexing step is a prerequisite. It focuses on organizing a large amount of data and serves as an upstream step of retrieval. For example, RAPTOR<sup>[4]</sup> shows a significant performance improvement by

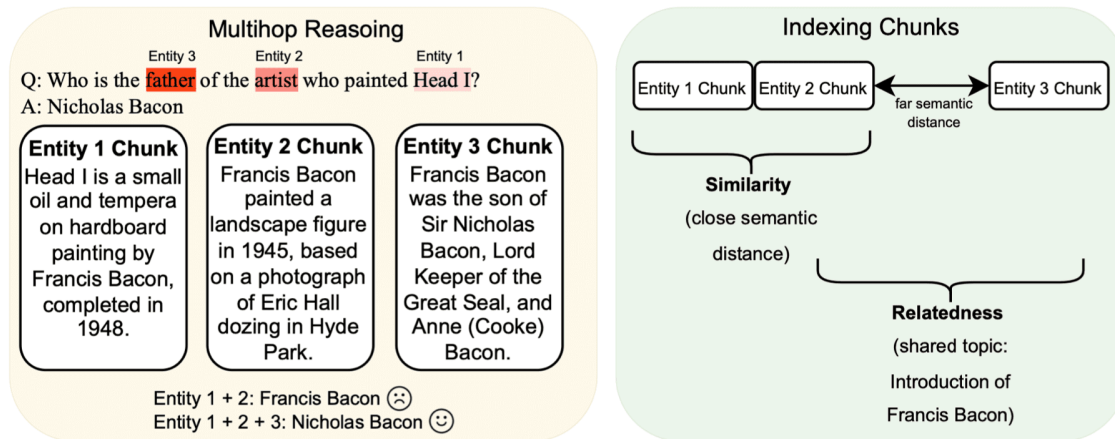
adding recursive summaries to text chunks of a dataset, which demonstrates the potential of adding synthesized information for retrieval. The added recursive summaries combine semantically similar information within a dataset. GraphRAG<sup>[5]</sup>, on the other hand, indices an entire corpus via an entity-guided knowledge graph. It then constructs summaries from closely-related entities and their mentions, synthesizing the connections and relatedness among different pieces of information.

However, none of the existing methods address the importance of indexing from both similarity and relatedness sides, which limits a holistic understanding of the provided dataset. We define **similarity** as the semantic distance of text pieces and **relatedness** as the degree of connection of texts based on signals such as entities and propositions. Indexing similar and related information facilitates more comprehensive knowledge integration than indexing individual kind of information. As shown in Figure 1, a complex question that involves two hops of reasoning requires the retrieval and synthesis of relevant entity chunks. For example, synthesizing entity 1 and 2 chunks would encourage LLMs to generate “Francis Bacon”, which is a common mistake. Entity 1 and 2 chunks have close semantic distance. On the other hand, synthesizing entity 2 and 3 chunks would not maximize the probability of the correct answer even when entity 1 chunk is retrieved, since LLMs may struggle to reason about the painter of “Head I” in long-context environment<sup>[6]</sup>. In this case, entity 2 and 3 chunks are related due to a shared topic. Therefore, it is important to synthesize both similar and related information in order to maximize the chance of retrieving relevant knowledge, in particular, for multihop reasoning questions. We verify the bottleneck of solely modeling similarity or relatedness through quantitative methods in Section 3 and demonstrate that neither perspective yields the optimal performance.

In this paper, we propose SiReRAG, which stands for RAG indexing of similarity and relatedness as shown in Figure 2. On the similarity side, SiReRAG follows RAPTOR<sup>[4]</sup> to build a recursive tree based on chunk similarity. We adopt a shallow tree with 4 levels in total. On the relatedness side, SiReRAG first extracts entities (e.g., “Sonnet 110” and “William Shakespeare”) and fine-grained propositions (e.g., “Sonnet 110 is one of 154 sonnets written by William Shakespeare.”) from each text chunk/document using LLMs. We group these propositions into aggregated ones via entities, simply concatenating them with the original order in chunk/document. These proposition aggregates contain related information, because they mention shared entities. Then, recursive summaries with soft clustering are built on top of those aggregated propositions. Finally, we index both trees by flattening nodes in each tree for retrieval.

We show that SiReRAG is effective on a variety of multihop question answering (QA) datasets including MuSiQue<sup>[7]</sup>, 2WikiMultiHopQA<sup>[8]</sup>, and HotpotQA<sup>[9]</sup>. We find that SiReRAG consistently outperforms the strongest RAG indexing methods, achieving an average F1 score improvement of at least 1.9%. We conduct

an ablation study on several components in SiReRAG and observe that adding proposition aggregates to the similar information within a dataset yields the most improvement, which echoes our motivation. SiReRAG is also a reasonably efficient model, as it does not introduce many lengthy or redundant retrieval candidates.



**Figure 1.** Challenges of existing RAG indexing methods for multihop reasoning. Entity 1 and 2 chunks contain similar information while entity 2 and 3 chunks contain related contents. Since synthesizing information only based on entity 1 and 2 (or entity 2 and 3) will lead to a higher probability of a wrong answer, an indexing method that considers both similarity and relatedness is needed to maximize retrieving relevant knowledge for multihop questions.

## 2. Related Work

### RAG

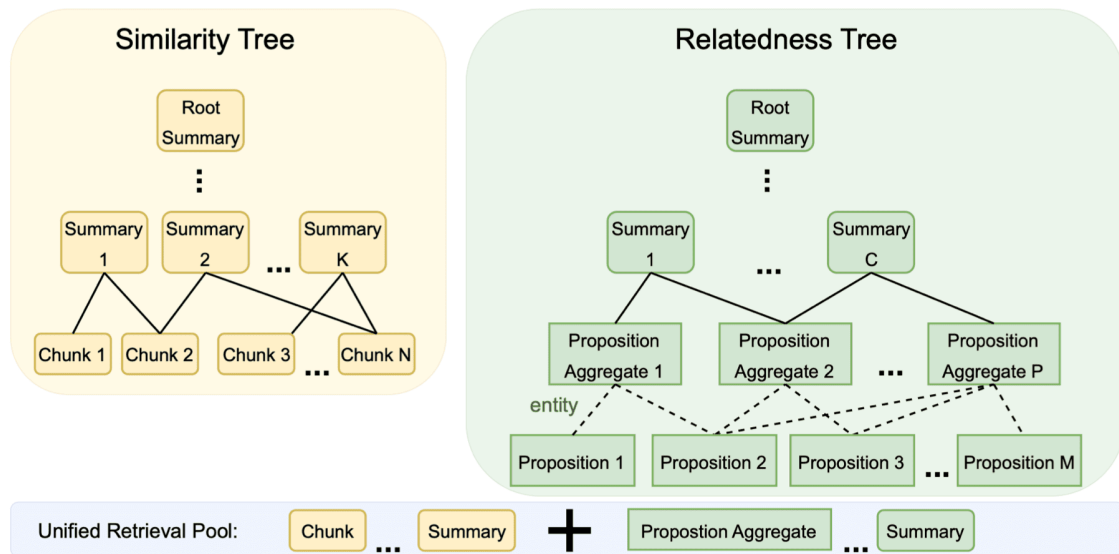
RAG is a framework that integrates retrieval mechanisms into generative models to enhance text generation by leveraging external knowledge. This concept has evolved from earlier retrieval-based methods such as DrQA<sup>[10]</sup> and DPR<sup>[11]</sup>. Instead of separating retrieval and generation phases, researchers also showed the potential of tightly coupling retrieval and generation into an end-to-end framework<sup>[1]</sup>.

Recent advances in retrieval mechanisms include leveraging LLMs as retrievers<sup>[12][13]</sup> and exploring retrieval granularity such as proposition<sup>[14]</sup>. Here each proposition is an atomic expression that contains a factoid presented in natural language, which is similar to the propositions used by SiReRAG. Inspired by the idea of retrieval granularity, we combine several different text granularities (entity, propositions, text

chunk, and summary) to index data. As a representative work on text segmentation, researchers proposed the first supervised approach to generate hierarchical segmentation structures<sup>[15]</sup>.

## RAG Indexing

An earlier work<sup>[1]</sup> shows the benefits of document index on the overall retrieval performance. Most recent works on RAG indexing include RAPTOR that builds a tree with recursive summaries<sup>[4]</sup>, HippoRAG that leverages the hippocampal indexing theory of human long-term memory for deep knowledge integration<sup>[16]</sup>, and GraphRAG that constructs an entity-guided knowledge graph<sup>[5]</sup>. However, all these works overlooked the importance of considering both similarity and relatedness during indexing. Specifically, RAPTOR integrated knowledge only based on similarity, and the other two only considered relatedness to synthesize information. Although these three approaches achieved competitive performance on different kinds of datasets and HippoRAG has the same goal as ours (multihop reasoning benchmarks), SiReRAG is fundamentally different from them in terms of the explicit incorporation of both similar and related knowledge.



**Figure 2. SiReRAG Tree.** We adopt RAPTOR<sup>[4]</sup> to construct the similarity tree (left). On the right, we construct the relatedness tree by clustering the propositions based on their entities to get proposition aggregates and having recursive summaries on top. Note that propositions are not included in the relatedness tree, so their connections to proposition aggregates are marked with dashed lines.

### 3. Bottleneck of Solely Modeling Similarity or Relatedness

To verify our hypothesis of insufficient knowledge integration when solely modeling similarity or relatedness, we perform different kinds of clustering philosophies on a retrieval corpus of MuSiQue<sup>[7]</sup>. Using the same retrieval corpus as HippoRAG<sup>[16]</sup>, we obtain 1000 questions from the validation set of MuSiQue along with their candidate passage clusters (each cluster includes supporting and distractor passages). We include distractor passages for a more realistic setting, since they are semantically close and/or related to the supporting candidates. Treating these clusters as the gold labels for different queries, we run our own clustering on all passages based on either similarity or relatedness.

Following RAPTOR<sup>[4]</sup>, we use Gaussian Mixture Models (GMMs) to perform soft clustering, assuming that a candidate passage can belong to multiple clusters. For similarity, we run GMMs on the deep representations of all passages to find semantically similar groups. For relatedness, we first extract the topic of each passage using OpenAI GPT-4o and then cluster passages based on the representations of their topics, because we assume that related passages share similar topics. To obtain deep representations of either passages or topics, we use OpenAI text-embedding-3-small for a balance of performance and efficiency. To evaluate the overlapping ratio among clusters, we convert every cluster into pairwise connections. For instance, given two clusters [1, 2, 3] and [3, 5], the resulting pairwise connections are as follows: “1-2”, “1-3”, “2-3”, and “3-5”. By computing the number of shared pairwise connections between gold clusters and predicted clusters, we aim to see how many pairwise connections from gold labels are covered by our two clustering philosophies, and we believe this coverage is a key indicator of knowledge integration for RAG indexing. We also report the overlapping ratio between “supporting only” similarity and relatedness to make a point of combining both clustering philosophies.

Table 1 shows the coverage of supporting only or all passages as gold labels. Higher coverage is obtained when we only use supporting passages as gold, which indicates the commonality of supporting ones with respect to distractors. Both philosophies are able to capture this commonality. However, all these percentage scores are low, with the highest one being around one-fifth of the total coverage (19.14%). This indicates a significant insufficiency in knowledge synthesis when we model similarity or relatedness solely. As a result, the chance of retrieving relevant knowledge for multihop reasoning questions could be suboptimal.

Taking a further step in the supporting-only setting, we find that only 50.15% of the correct similarity connections overlap with the correct relatedness connections. Correct similarity connections are those gold connections covered by similarity, while correct relatedness connections are those covered by relatedness. Conversely, 68.85% of the correct relatedness connections overlap with the correct similarity connections.

In other words, predictions based on similarity or relatedness are not identical, and we can potentially leverage both to improve retrieval performance and facilitate a more comprehensive knowledge integration process. Although more customized clustering algorithms of each clustering philosophy can be proposed, combining both similarity and relatedness offers an effective and straightforward solution.

	Similarity Coverage		Relatedness Coverage		Overlapping Ratio	
	Supporting Only	All	Supporting Only	All	Overlap@Similarity	Overlap@Relatedness
<b>Coverage</b>	19.14%	10.70%	13.94%	8.51%	50.15%	68.85%

**Table 1.** Coverage percentage between different clusters. MuSiQue clusters include supporting and distractor passages. The “all” setting treats both as gold, while “supporting only” uses only supporting passages as gold. We first show the coverage of supporting or all passages under two clustering philosophies. We then report the overlapping ratio between “supporting only” similarity and relatedness to motivate our work of combining both philosophies.

## 4. Methodology

We propose SiReRAG, a RAG indexing framework guided by similarity and relatedness. As shown in Figure 2, its left tree integrates information based on similarity while its right tree integrates information based on relatedness. As a first step, we study an alternative tree design to determine whether we can develop a generalized tree structure for similarity and relatedness trees, and beyond (Section 4.1). After the construction of the similarity tree, we extract propositions and their entities from our multihop reasoning dataset and perform clustering based on entities to synthesize related information (Section 4.2). Indexing separate similarity and relatedness trees (Section 4.3), SiReRAG explicitly models both kinds of information within our dataset.

### 4.1. Exploring a hierarchical structure of trees

For efficiency, we stick to a tree structure to organize texts and explore whether a more structured tree design would offer performance improvement. RAPTOR placed all text chunks at the bottom level and recursive summaries at upper levels, but this design does not closely follow the commonsense definition of a tree<sup>[17]</sup>, where multiple levels of data abstraction are provided. The reason is that different text chunks of a document may showcase different levels of abstraction, which serves as the assumption of many

document discourse trees<sup>[18][19]</sup>. For example, text chunks of the introduction section of a research paper are more likely to be more abstractive than those from the methodology section. Therefore, we first explore placing different text chunks on different levels of a tree based on their levels of abstractiveness.

For this analysis, we choose QuALITY dataset<sup>[20]</sup> instead of multihop reasoning ones, because the documents of QuALITY are longer, resulting in more text chunks than using the other datasets. We prompt GPT-4o to identify a two-level hierarchy for all candidate chunks: low (text chunks describing fine-grained details about a topic) and high (text chunks giving an overview of a topic and summarizing fine-grained details) abstractiveness. Summary nodes will still start from the second level (the level above the bottom level), so the second level will contain both summary nodes (of the low-abstractive chunks) and high-abstractive chunks. As in RAPTOR, nodes of the third level and above are recursive summaries of their children. Using the same text chunks across two methods, we compare RAPTOR against our hierarchical design. Since we simply retrieve the top 10 relevant nodes via an embedding model (text-embedding-3-small) when a query arrives, the idea of hierarchical structure of chunks would affect summary nodes due to the differences of their children.

	Average Accuracy	Standard Deviation
<b>RAPTOR</b>	78.88	0.005
<b>Hierarchical Text Chunks</b>	78.76	0.004

**Table 2.** QA performance of having hierarchical text chunks on the validation set of QuALITY. Due to randomization of clustering and generation temperature of LLMs, we run each indexing method for 5 times and compute their average and standard deviation of accuracies.

As shown in Table 2, we do not observe a clear trend of improvement from using hierarchical text chunks, even averaged over 5 different runs. The RAPTOR tree is relatively more efficient, as identifying hierarchical text chunks and document discourse trees would require additional computation. Thus, we conclude that a more structured tree design would not significantly improve performance, as long as the correct information is indexed to answer a question. Since RAPTOR uses Gaussian Mixture Models and representations of text chunks to perform clustering, it sets a nice example of integrating knowledge based on similarity. We adopt it to construct our similarity tree as shown in the left part of Figure 2.

## 4.2. Synthesizing information based on relatedness

For relatedness, we need to synthesize information based on a different philosophy than RAPTOR. Because related information pieces always share some degree of connection (e.g., overlapping subjects), we assume that two text pieces are related if they mention the same entity (e.g., person, location, product, etc). For example, entity 2 and 3 chunks in Figure 1 are unlikely to be clustered in the same group based on similarity, but since they both mention Francis Bacon, we are able to connect them together.

### *Modeling Relatedness with Entity-Specific Propositions:*

To effectively use entities for organizing related content, we first need to determine the appropriate granularity for text pieces. There are three main limitations with directly connecting entities to standard text chunks. First, a chunk often contains information beyond the scope of a specific entity, making it challenging to localize information about one entity, potentially adding noise. Second, aggregating all chunks in an indexing corpus for each entity can result in hundreds of thousands of tokens for each entity, which may lead to long context performance issues, such as losing critical information in the middle<sup>[6]</sup> or experiencing low coverage and citation performance<sup>[21]</sup>. Third, linking with chunks will introduce redundancy as each chunk may be a part of multiple entity clusters. Therefore, inspired by recent works on retrieval granularity<sup>[22][14]</sup>, we propose to use short entity-specific “propositions” to represent fine-grained knowledge about entities and build our relatedness tree.

### *Extracting Propositions and Entities from Documents*

We define a proposition as “a factual statement describing important information (preferably about some entities) from a paragraph”. We extract entities and propositions using the Distill-SynthKG pipeline<sup>[23]</sup>, adapting its SynthKG workflow. Specifically, we first rewrite chunks of documents taken from 10K BAAI/IndustryCorpus documents<sup>1</sup> to resolve entity references. Here we use Meta-Llama-3-70B-Instruct<sup>2[24]</sup>, and the LLM prompt of rewriting chunks is in Figure 6. Then, after we obtain entities from these rewritten chunks, we extract all relevant propositions and their associated entities (using the same LLM) as detailed in the LLM prompt in Figure 7. Finally, we consolidate all propositions and entities as the training data to fine-tune Mistral-7B-Instruct-v0.3<sup>3[25]</sup>. This smaller fine-tuned model is used to extract propositions and their associated entities from our multihop datasets.

Our prompt for extracting propositions and entities does not require every proposition to have associated entities. When we prompt LLMs to generate entities for each proposition, they sometimes produce common nouns as entities for those propositions that lack actually associated entities. This can lead to the clustering



of unrelated propositions based on common nouns, potentially introducing noise into the relatedness tree. For example, as illustrated in Figure 3, we prefer to avoid LLMs generating *Drug* as an entity for *proposition 1* due to its ambiguity. Subsequently, we exclude propositions without associated entities when constructing the relatedness tree, ensuring that only high-quality, entity-linked propositions are utilized. The Table 3 provides a detailed breakdown of key statistics for extracted propositions and entities from the MuSiQue, 2Wiki, and HotpotQA datasets.

					#Props per Entity		
Dataset	#Chunks	#Props.	#Ents	#Pops Aggs.	Avg.	Max.	Min.
MuSiQue	11,656	54,605	50,926	20,788	2.74	168	1
2Wiki	6,119	27,697	29,490	11,108	2.49	195	1
HotpotQA	9,221	47,153	46,856	18,278	2.66	165	1

**Table 3.** Key Statistics for extracted propositions and entities from MuSiQue, 2Wiki, and HotpotQA datasets. We show the number of chunks, propositions, entities, proposition aggregates, and the average, maximum, and minimum number of propositions per entity across the three datasets.

*Proposition 1:* Drug sales are reaching record highs as new therapies are developed and approved.  
*Entities:* []  
*Proposition 2:* Drug sales for Eli Lilly’s Mounjaro and Novo Nordisk’s semaglutide are reaching record highs as new therapies are developed and approved.  
*Entities:* [Eli Lilly, Mounjaro, Novo Nordisk, Semaglutide]

**Figure 3.** Examples of propositions with and without associated entities.

### From Entity-specific propositions to Relatedness Tree

We concatenate related propositions that share the same entity using exact match to form proposition aggregates. We ensure that all propositions from the same document are grouped together and maintain their original order. By treating these proposition aggregates as pseudo-documents, we apply the same clustering pipeline in RAPTOR to obtain recursive summaries at levels above them and build the relatedness tree. Given that most propositions involve multiple entities, each proposition is associated with several entity clusters, offering two key advantages. First, it effectively mimics soft clustering as a single

proposition may belong to multiple aggregates. Secondly, when constructing recursive summaries within the RAPTOR framework, the shared propositions across different aggregates result in high embedding similarity, ensuring that these aggregates remain clustered together even at higher levels in the tree. This relatedness tree forms the right part of SiReRAG as shown in Figure 2. Note that we only index aggregated propositions instead of individual propositions for better inference efficiency.

### 4.3. Indexing similarity and relatedness trees

We propose to construct similarity and relatedness trees independently. This approach ensures that summary nodes in one tree do not access the clusters of the other, leading to a simpler design. There is another slightly complex design in which we allow summary nodes from one tree to access clusters from the other tree. This interaction may enable summary nodes in both trees to inform and enhance each other and improving their informativeness and consequently performance. However, this approach sacrifices the distinction between similarity and relatedness. Additionally, allowing cross-tree interaction leads to more nodes to cluster at each level as well as requires summarization based on a greater number of nodes per cluster, all of which increases the overall complexity of the system. We experimented with both settings and did not observe performance improvement as shown in Appendix A. Therefore, we opted for the simpler first implementation in our evaluation.

Flattening all tree nodes, we place them into a unified retrieval pool. In other words, regardless of a node’s origin (*e.g.*, bottom or upper levels, similarity or relatedness trees), it is added to a single list containing all nodes.

## 5. Experiment Setup

### 5.1. Datasets

To demonstrate the effectiveness of SiReRAG, we select three representative multihop QA datasets: MuSiQue<sup>[7]</sup>, 2WikiMultiHopQA<sup>[8]</sup>, and HotpotQA<sup>[9]</sup>. Using the same corpus as HippoRAG<sup>[16]</sup>, we obtain 1000 questions from each validation set of these three datasets.

### 5.2. Baselines

We select RAPTOR, HippoRAG, and GraphRAG as state-of-the-art retrieval baselines. As discussed above, RAPTOR integrates knowledge based on similarity while the other two approaches focus on relatedness. Specifically, HippoRAG has both indexing and retrieval components, and we use ColBERTv2<sup>[26]</sup> as the retriever of HippoRAG due to its strongest QA performance reported. Although GraphRAG has a different

goal (global questions directed at an entire dataset) than ours, we include it to show its performance on multihop QA datasets. Since the queries in our datasets ask fine-grained details, we use the local search function of GraphRAG instead of its global search.

### 5.3. Evaluation metrics

We use exact match (EM) and F1 scores to measure the QA performance of different models. Both metrics evaluate how accurate a generated answer is with respect to the ground truth. Like RAPTOR<sup>[4]</sup>, we do not assess retrieval performance directly. The reason is that both SiReRAG and RAPTOR create new candidates (*e.g.*, summary and proposition aggregate) in the retrieval pool, so it would be unfair to compare methods in terms of retrieval scores across different pools. Instead, QA performance is the best indicator of the overall capability of both RAG pipelines.

We use the average time per query (TPQ) and the time-pool efficiency ratio (TPER) to measure the efficiency of SiReRAG and RAPTOR, as both methods share a significant portion of their retrieval candidates. Average TPQ measures the average time (in seconds) taken to answer a query, and it represents the inference time of a method. For TPER, it computes the growth of total inference time with respect to the growth of the retrieval pool size between two methods:

$$\text{TPER} = \frac{\text{Inference-Time A} / \text{Inference-Time B}}{\text{Pool-Size A} / \text{Pool-Size B}} \quad (1)$$

Setting SiReRAG as method A and a baseline as method B, we aim to ensure that the growth of inference time does not scale proportionally with the increase in the retrieval pool size. The reason behind is that there are many efficiency considerations (*e.g.*, length and redundancy of retrieval candidates) beyond just the sheer number of retrieval candidates. Parallelization could also be designed to retrieve candidates simultaneously from similarity and relatedness trees, thereby minimizing the effect of retrieval pool size. A TPER value less than 1 indicates reasonable efficiency, whereas a TPER value greater than 1 signifies low efficiency.

### 5.4. Implementation details

To generate final answer, we use GPT-4o and the same prompt (“*answer this question in as fewer number of words as possible.*”) to answer queries for all methods, since we aim to control the instruction-following capabilities across all methods. We use either GPT-3.5-Turbo or GPT-4o as the choice of LLM if any methods involve LLM calls. We use OpenAI’s text-embedding-3-small as the embedding model for all methods. During retrieval, we select top 20 candidates that match the provided query for all methods,

because there is a large number of text chunks in our datasets and SiReRAG is expected to perform better when retrieving more due to the incorporation of proposition aggregates and their recursive summaries.

## 6. Results and Analysis

Our results and analysis aim to answer the following research questions:

- **RQ 1:** How does SiReRAG compare against other state-of-the-art baselines (sec 6.1)?
- **RQ 2:** As an important contribution of SiReRAG, is considering both similarity and relatedness an effective method (sec 6.1 and 6.2)?
- **RQ 3:** What is the effect of each component in SiReRAG(sec 6.2)?
- **RQ 4:** What is the applicability of SiReRAG(sec 6.3)?
- **RQ 5:** With the addition of relatedness tree, is SiReRAG an efficient method (sec 6.4)?

### 6.1. Overall results

Our overall results are presented in Table 4. We show results on more datasets (single-hop QA, other multihop QA, and ambiguous questions) in Appendix E to show the generality of SiReRAG across various complex reasoning tasks. Besides quantitative scores, we also conduct our qualitative analysis in Appendix D.

Model	MuSiQue		2Wiki		HotpotQA		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
HippoRAG (GPT-3.5-Turbo)	32.60	43.78	66.40	74.01	59.90	74.29	52.97	64.03
RAPTOR (GPT-3.5-Turbo)	35.30	47.47	54.90	61.20	58.10	72.48	49.43	60.38
GraphRAG (GPT-4o)	12.10	20.22	22.50	27.49	31.70	42.74	22.10	30.15
RAPTOR (GPT-4o)	36.40	49.09	53.80	61.45	58.00	73.08	49.40	61.21
SiReRAG (GPT-3.5-Turbo)	38.90	52.08	60.40	68.20	62.50	77.36	53.93	65.88
SiReRAG (GPT-4o)	40.50	53.08	59.60	67.94	61.70	76.48	53.93	65.83

**Table 4.** QA performance of SiReRAG and baselines. As elaborated in Section 5.4, GPT-4o is used to handle QA for all models, and we use two different LLMs (specified in the parentheses) to build indexing structures. We highlight the best scores using either LLM for indexing in green color.

### Improvement over baselines

SiReRAG delivers consistent improvement over RAPTOR, HippoRAG, and GraphRAG. With an exception on 2Wiki when comparing against HippoRAG, SiReRAG achieves significantly higher performance than indexing baselines (*e.g.*, approximately 5% higher than RAPTOR on average F1, up to 8.3% improvement of

F1 on MuSiQue than HippoRAG, and more than 20% higher than GraphRAG on average EM and F1). This demonstrates the advantage of SiReRAG on multihop QA and modeling both similarity and relatedness. Specifically, SiReRAG outperforms RAPTOR due to the incorporation of a relatedness tree, and it has better overall performance than HippoRAG, because we explicitly model similarity while HippoRAG prioritizes relatedness signals such as nodes with the most edges. We see that HippoRAG is particularly strong on 2Wiki benchmark, which is also reported in its original paper. Thus, we believe 2Wiki is the best fit of HippoRAG, but it has lower performance scores than SiReRAG on other datasets. One potential reason is that the entity-centric design of 2Wiki may be well-suited for HippoRAG, as noted in the HippoRAG paper.

As for GraphRAG, it considers relatedness solely, and it delivers the worst performance scores on our datasets. After a manual verification, we find that GraphRAG often provides “I don’t know” answers, suggesting that it prefers not to give a concrete answer. Since GraphRAG is designed to handle query-focused summarization of an entire corpus, it is not a competitive method in terms of accuracy for existing multihop QA tasks.

### *Effect of LLM choice*

When comparing the performance of SiReRAG using GPT-4o as the LLM against itself using GPT-3.5-Turbo, we find the QA performance is not significantly affected. This phenomenon also holds on RAPTOR. Since the choice of LLM for SiReRAG and RAPTOR only affects summarization results, we believe GPT-3.5-Turbo is a sufficiently good option for both methods. This allows researchers to pursue a more cost-effective solution with SiReRAG for indexing.

## *6.2. Ablation study*

To dissect SiReRAG, we perform a comprehensive ablation analysis as shown in Table 5. There are several variances, including (A) remove the recursive summary on the relatedness tree; (B) add all the propositions into the retrieval pool, and keep all aggregated propositions and recursive summary on the relatedness tree; (C) same as (B) but remove aggregated propositions; (D) same as (C) but further remove the recursive summary design on the relatedness tree.

Variants	MuSiQue		2Wiki		HotpotQA		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
SiReRAG	40.50	53.08	59.60	67.94	61.70	76.48	53.93	65.83
(A): SiReRAG – Re. Summary	37.50	50.38	57.70	65.75	61.20	75.99	52.13	64.04
(B): SiReRAG + Proposition	39.10	51.80	58.10	65.53	60.80	75.26	52.67	64.20
(C): (B) – Aggregate	34.70	47.82	53.20	60.22	58.90	73.38	48.93	60.47
(D): (C) – Re. Summary	33.90	46.19	53.00	59.13	57.00	71.09	47.97	58.80
(E): Dual Clustering on Chunks	34.80	47.32	53.50	59.93	56.60	71.84	48.30	59.70

**Table 5.** Ablation study of SiReRAG.

### *Entity clustering.*

In (E), we do not maintain a separate relatedness tree and add an additional clustering philosophy to the similarity tree. Specifically, each text chunk in the similarity tree is simplified to, 'This chunk mentions entity 1 and entity 2,' if both entities are extracted by our LLM. We then run GMMs (the same clustering method as RAPTOR) on these simplified chunks. Once the clustering decisions are obtained, we group the original chunks as additional clusters and append these clusters to the similarity tree, allowing higher levels of the tree to incorporate both clustering philosophies. Since entities primarily determine the outcome of this additional clustering approach, we apply entity clustering to model relatedness on the similarity tree. This allows us to eliminate proposition aggregates in order to examine their utility.

### *Findings.*

Overall, we observe performance drops across all variations, highlighting the effectiveness of our design for SiReRAG. First, the recursive summary on the relatedness tree proves beneficial, as seen in both (A) and (D). Interestingly, adding more propositions to retrieval negatively impacts performance, as shown in (B). This indicates adding redundant information into the retrieval pool hurts the QA performance, since we keep the aggregated propositions in SiReRAG. From (C), we also find that aggregated propositions are essential, with their removal resulting in a significant performance decline. This is an important indicator that adding grouped knowledge about relatedness to the similarity tree would offer improvements, which echoes the bottlenecks described in Section 3.

Both (A) and (B) have better performance than RAPTOR (GPT-4o) from Table 4, which indicates the advantage of proposition aggregates. In contrast, although (E) also models both similarity and relatedness, it exhibits a notable decline comparing against SiReRAG. This finding demonstrates the necessity of proposition aggregates of modeling relatedness. Because proposition aggregates reduce noise and

information redundancy more effectively than text chunks as described in Section 4.2, they serve as an effective carrier of related dataset contents.

### 6.3. Applicability of SiReRAG

We analyze how applicable SiReRAG is when a specific retrieval method is chosen. Therefore, we select BM25<sup>[27]</sup> and ColBERTv2<sup>[26]</sup> as additional reranking-based options. We run them on the retrieval pool of SiReRAG to demonstrate its utility. We show how SiReRAG can complement these non-indexing options in Table 6. Results show that having SiReRAG benefits both BM25 and ColBERTv2 significantly, which demonstrates the advantage of our solution as the upstream step of these methods. On the other hand, RAPTOR only improves the QA performance on MiSuQue while showing performance degradation on other datasets. Thus, the utility of SiReRAG surpasses that of RAPTOR in the context of multihop reasoning. We also apply SiReRAG on an iterative retrieval method called self-ask<sup>[28]</sup> and obtain significant performance improvement as shown in Appendix B. Our method showcases wide applicability on multihop QA across various retrieval methods.

Variants	MuSiQue		2Wiki		HotpotQA		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
BM25	25.90	35.88	53.00	58.58	57.70	71.32	45.53	55.26
RAPTOR + BM25	27.00	38.91	50.60	57.00	56.90	70.88	44.83	55.60
SiReRAG + BM25	35.00	47.66	58.20	65.72	61.70	75.88	51.63	63.09
ColBERTv2	34.00	46.80	52.90	59.48	59.00	73.42	48.63	59.90
RAPTOR + ColBERTv2	34.20	47.17	50.30	57.51	57.90	72.31	47.47	59.00
SiReRAG + ColBERTv2	38.10	51.32	56.70	64.74	60.90	75.72	51.90	63.93

Table 6. Applicability of SiReRAG when a specific retrieval method is selected. We feed our non-indexing models with the retrieval pool of SiReRAG and see whether QA performance improves.

### 6.4. Efficiency of SiReRAG

As shown in Table 7, we compare the efficiency of SiReRAG and RAPTOR using the metrics described in Section 5.3. All the values listed involve the time taken to retrieve the top 20 candidates and prompt GPT-4o to answer the query.

RAPTOR requires less inference time than SiReRAG on average, which is expected due to SiReRAG’s larger retrieval pool. However, with slightly longer inference time, SiReRAG has much better performance as discussed previously. To evaluate whether SiReRAG remains a reasonably efficient method, we compute its TPER values to measure its growth of total inference time relative to its growth of retrieval pool size. Since

all its TPER values are well below 1, SiReRAG demonstrates reasonable efficiency without introducing many lengthy or redundant retrieval candidates.

	MuSiQue		2Wiki		HotpotQA		Average	
Model	TPQ↓	TPER↓	TPQ↓	TPER↓	TPQ↓	TPER↓	TPQ↓	TPER↓
RAPTOR	1.560	-	1.437	-	1.502	-	1.500	-
SiReRAG	2.653	0.600	1.974	0.499	2.319	0.517	2.315	0.539

**Table 7.** Efficiency of SiReRAG and RAPTOR. Since both methods share a significant portion of retrieval candidates, we designate SiReRAG as Method A and RAPTOR as Method B in the TPER columns, as defined in Equation 1.

## 7. Conclusion

In this paper, we identify the bottleneck of solely modeling similarity or relatedness when we need to index a multihop reasoning dataset for knowledge integration. To address it, we introduce SiReRAG, an innovative RAG indexing approach that considers both similarity and relatedness. SiReRAG delivers a consistent improvement over state-of-the-art indexing baselines across several multihop QA benchmarks.

## Appendix A. An Alternative Design of Allowing Cross-Tree Interaction

We discuss an alternative design of combining similarity and relatedness trees. Specifically, this design combines nodes from both sides in the same pool for finding additional clusters and performing summarization at every tree level. In other words, we find additional clusters by concatenating the nodes of both trees, which considers cross-tree interaction instead of keeping them separate.

As shown in Table 8, the performance of considering cross-tree interaction is slightly lower than SiReRAG. Therefore, it is more efficient to keep trees separate in order to reduce the overall complexity of the system as discussed in Section 4.3.

## Appendix B. Additional Experiment on Other Non-Indexing Methods

Although there are many existing methods that work on multihop reasoning tasks, SiReRAG is about indexing corpus data under RAG setup. In other words, instead of being our baselines, other non-indexing



works<sup>[28][29]</sup> focus on other dimensions of improving performance on complex reasoning tasks.

To further demonstrate the applicability of SiReRAG, we have run two additional sets of experiments: (1) the closed-book setting, and (2) an iterative retrieval method<sup>[28]</sup> called self-ask specifically designed for multihop reasoning. The closed-book setting means to directly get the final answer from GPT-4o without any retrieval. For self-ask, we prompt GPT-4o in two iterations without using a search engine. In the first iteration, the model is prompted to propose follow-up questions and provide answers to them. In the second iteration, GPT-4o is instructed to answer the final question by incorporating the follow-up thought process. We feed the model with a one-shot example and 10 retrieved candidates that match the final question in both iterations. For the self-ask retrieval pool, we use either all text chunks or SiReRAG.

Similar to Table 6, our experiment in Table 9 on self-ask shows that SiReRAG can complement existing methods for optimal performance. We see that the closed-book setting yields the worst performance, which indicates that LLMs’ parametric knowledge alone does not offer a decent performance on our datasets. Then, we see SiReRAG successfully improves the scores of self-ask, demonstrating its wide applicability. By leveraging SiReRAG’s retrieval pool, we view our method as an augmentation to other non-indexing methods for multihop reasoning, rather than as a competitor.

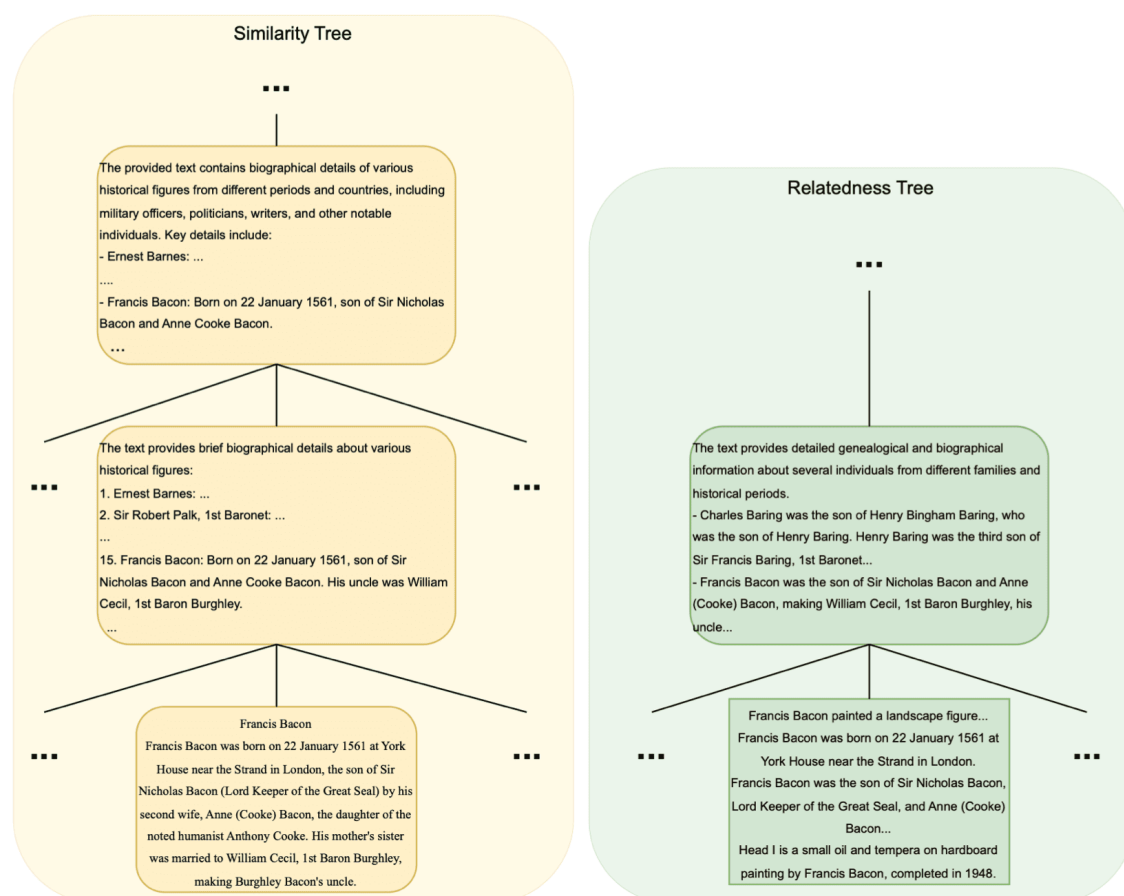
Efficiency-wise, we show TPQ of using both SiReRAG and self-ask in Table 10. By having SiReRAG, the TPQ of self-ask increases by approximately 1.2 seconds over the three datasets. Since self-ask requires two iterations of LLM prompting in our implementation, the increase in TPQ is relatively small compared to the significant performance improvement brought by SiReRAG.

## Appendix C. Retrieval Pool Size

The retrieval pool sizes of SiReRAG on MuSiQue, 2Wiki, and HotpotQA are 35070, 19100, and 29934 respectively. The retrieval pool sizes of RAPTOR on MuSiQue, 2Wiki, and HotpotQA are 12371, 6939, and 10031 respectively. SiReRAG’s retrieval pool size is slightly less than three times the size of RAPTOR’s. Considering the discussion on TPER in Section 6.4, we believe our method is reasonably efficient.

## Appendix D. An Example SiReRAG Tree

Using the question “who is the father of the artist who painted Head I?” as an example, we focus on the relevant part of the SiReRAG tree in Figure 4 to conduct our qualitative analysis.



**Figure 4.** Relevant part of the SiReRAG tree for the question: “who is the father of the artist who painted Head I?”.

Variants	MuSiQue		2Wiki		HotpotQA		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
<b>SiReRAG</b>	40.50	53.08	59.60	67.94	61.70	76.48	53.93	65.83
<b>Cross-tree interaction</b>	40.20	53.06	58.30	65.25	60.30	75.71	52.93	64.67

**Table 8.** QA performance of two designs: separating similarity and relatedness trees (SiReRAG) and allowing cross-tree interaction.

Model	MuSiQue		2Wiki		HotpotQA		Average	
	EM	F1	EM	F1	EM	F1	EM	F1
Closed-book	10.0	22.0	19.0	34.0	29.0	44.0	19.3	33.3
Self-ask	31.20	44.35	55.00	61.99	57.10	71.11	47.77	59.15
SiReRAG + self-ask	36.50	49.12	57.20	65.13	59.70	74.07	51.13	62.77

**Table 9.** QA performance of closed-book and self-ask. We feed self-ask with the unified retrieval pool of SiReRAG and see whether performance benefits from that.

Model	MuSiQue TPQ	2Wiki TPQ	HotpotQA TPQ	Average TPQ
Self-ask	2.72	2.21	2.29	2.41
SiReRAG + self-ask	4.53	3.07	3.30	3.63

**Table 10.** Efficiency of using self-ask with SiReRAG. We show TPQ on MuSiQue, 2Wiki, and HotpotQA datasets.

For the multihop question in Figure 1 (correct answer: Nicholas Bacon), the MuSiQue corpus contains one relevant paragraph stating, “Francis Bacon was born on 22 January 1561 at York House near the Strand in London, the son of Sir Nicholas Bacon...”. This paragraph is a leaf node of the similarity tree shown in Figure 4. The similarity tree for the entire MuSiQue corpus has two more mentions (both mentions are in summary nodes) of Nicholas Bacon, one of which reads: “...Francis Bacon: Born on 22 January 1561, son of Sir Nicholas Bacon and Anne Cooke Bacon...” The addition of our relatedness tree adds two more mentions of Nicholas Bacon: one is in a proposition aggregate (“...Francis Bacon was the son of Sir Nicholas Bacon, Lord Keeper of the Great Seal, and Anne (Cooke) Bacon... Head I is a small oil and tempera on hardboard painting by Francis Bacon, completed in 1948...”), and the other one is a summary node (“The text provides detailed genealogical and biographical information about several individuals from different families and historical periods... Francis Bacon was the son of Sir Nicholas Bacon and Anne (Cooke) Bacon, making William Cecil, 1st Baron Burghley, his uncle...”). Thus, our similarity tree (RAPTOR tree) has three mentions of Nicholas Bacon, but none of them contains Head I information. SiReRAG has five mentions of Nicholas Bacon, and one of them (the proposition aggregate) contains Head I information. This proposition aggregate groups several propositions together via the entity “Francis Bacon”. Because this node is the only retrieval candidate that fully matches the question, retrieving it would maximize the chance of generating

the correct answer. If we use the RAPTOR tree only, we will not have this retrieval candidate. We believe this is an excellent example of how a comprehensive knowledge integration process can enhance the performance of RAG in multihop reasoning.

## Appendix E. Performance on Single-Hop QA, MultiHop-RAG, and Ambiguous Questions

To showcase the generality of SiReRAG on more datasets of complex reasoning tasks, we run the comparison between SiReRAG and RAPTOR on single-hop questions and MultiHop-RAG dataset<sup>[30]</sup>. We also try ASQA dataset<sup>[31]</sup> that contains ambiguous factoid questions. For single-hop questions, we use MuSiQue dataset and collect all the decomposed questions of multihop queries. We filter out some decomposed questions if they are still multihop or are based on another question. As a result, we end up with 502 single-hop questions from MuSiQue. As for MultiHop-RAG, it is a more recent dataset. We filter all unanswerable questions and randomly select 350 “comparison” queries and 350 “inference” queries, which forms a pool of 700 queries in total.

Moreover, the primary difference between the multihop QA datasets and ASQA is that ASQA requires LLMs to reason across multiple perspectives (*e.g.*, disambiguated questions) of an ambiguous question and organize their generation into a coherent and detailed answer. We report scores on ASQA based on all 948 ambiguous questions of its development set.

Table 11 shows the performance scores of RAPTOR and our method. On the single-hop questions, SiReRAG still outperforms RAPTOR, but the lead narrows compared to the scores in Table 4. Since all queries in MuSiQue involve at least two hops, we observe that an increased number of reasoning hops positively impacts SiReRAG’s performance. This is because single-hop questions may not require comprehensive knowledge synthesis, as they only involve retrieving the relevant chunks for the single hop. However, with more hops, we not only need to retrieve relevant chunks but also synthesize them comprehensively. SiReRAG also delivers better performance on MultiHop-RAG, which echoes our main experiment.

Table 11 also displays STR-EM (string exact match), Disambig-F1, and Disambiguation-Rouge metrics for ASQA dataset. Specifically, STR-EM and Disambig-F1 dissect RAG answers to ambiguous questions into multiple perspectives and measure their accuracy with respect to each disambiguated question. Disambiguation-Rouge serves as an overall statistic that incorporates both ROUGE (with respect to the references) and accuracy scores. We see SiReRAG delivers a consistent improvement over RAPTOR, which again demonstrates the benefit of adopting SiReRAG on ambiguous queries.

Our comprehensive selection of datasets demonstrates the generality and contribution of SiReRAG across various complex reasoning tasks, with the most significant improvement observed in multihop QA.

## Appendix F. LLM Prompts

The prompt we use to perform summarization on a cluster of nodes is “summarize the provided text, including as many key details as needed”. This prompt is the same as RAPTOR. In Section 3, we use ‘identify the high-level topic of this paragraph as concise as possible’ to extract the topic of each passage. As mentioned in Section 4.1, the prompt used for identifying a two-level hierarchy for all chunks is shown in Figure 5. As mentioned in Section 4.2, the LLM prompt for rewriting chunks is shown in Figure 6, and the prompt for extracting propositions and entities is shown in Figure 7.

Model	Single-Hop		MultiHop-RAG		ASQA		
	EM	F1	EM	F1	STR-EM	Disambig-F1	Disambiguation-Rouge
RAPTOR	37.10	40.50	59.80	60.46	51.47	40.30	42.72
SiReRAG	37.20	41.49	61.30	61.89	52.81	41.82	43.47

**Table 11.** QA performance on single-hop questions, MultiHop-RAG, and ASQA.

```

## Instructions
1. Group sentences based on their topic and abstractiveness.
2. For each sentence group define metadata consisting of topic and the abstractiveness level.
3. There are two abstractiveness levels:
    a. low: sentences describe fine-grained details about the topic
    b. high: sentences gives an overview of the topic, summarizing fine-grained details.
4. Ensure that each group consists of contiguous sentences.
5. Respond using the below json format.
{
  "group_1":
  {
    "sentences": [list of sentence ids],
    "topic": "topic of the group",
    "abstractiveness": "high"
  },
  "group_2":
  {
    "sentences": [list of sentence ids],
    "topic": "topic of the group",
    "abstractiveness": "low"
  },
  "group_3":
  {
    "sentences": [list of sentence ids],
    "topic": "topic of the group",
    "abstractiveness": "high"
  }
}

```

**Figure 5.** Prompt of identifying a two-level hierarchy for all candidate chunks.

**Previous paragraph from Document:**

Gualala, the isolated Mendocino Coast town with a name that leaves most visitors tongue-tied, is on a new list of the 50 best places to live in the United States. Men's Journal magazine describes Gualala as an outpost of adventure lifestyle in its latest edition, which goes on sale today. The magazine describes Gualala (pronounced wa-LA-la by locals) as one of the below-the-radar places to make a move on before the word gets out. There were five such cities. The others were Homer, Alaska; Newport, Vt.; Logan, Utah; and Walla Walla, Wash. Rolling Stone magazine's Jann Wenner publishes Men's Journal, which has a paid circulation of about 620,000. Gualala joined three other California communities on the magazine's list: Santa Cruz, Mammoth Lakes and Bishop. We were looking for places that combined affordability, proximity to outdoor adventure and a generally undiscovered quality of life, said Erica Kestenbaum, a spokeswoman for Men's Journal.

**Instruction:**

Rewrite the below paragraph by resolving all entity coreferences with the preceding paragraph from document.

- Resolve all inter-sentence pronoun references.
- Make sure that all pronouns in a sentence refers to some named entity with in the same sentence.
- Explicitly mention entity names wherever necessary to remove ambiguity from a sentence. Remember to make each sentence clear and unambiguous.
- For each entity, use only the one most informative name.
- Do not generate anything except the rewritten paragraph.

**Paragraph:**

She said isolation played a factor. In Northern California, it's particularly difficult to find a beautiful coastal setting that isn't entirely overrun, she said. Gualala residents Monday were largely unaware of the magazine listing or the attention it could bring to the old logging town turned tourist center. A few coastal residents chuckled about any notion of affordability, given an influx of newcomers who've driven the median housing price to \$580,000 compared to the median family income of \$47,778. Others recalled an era when the Gualala region was better known for the logging of ancient redwoods, marijuana growing and boisterous beer drinking at the historic Gualala Hotel. Still there was a certain pride to the magazine's designation. Yvette White, a 25-year resident who works at the Gualala Sport; Tackle shop, said she's proud her town made it on the list.

**Output:**

Erica Kestenbaum said isolation played a factor. In Northern California, it's particularly difficult to find a beautiful coastal setting that isn't entirely overrun, Erica Kestenbaum said. Gualala residents Monday were largely unaware of the Men's Journal magazine listing or the attention it could bring to the old logging town turned tourist center. A few coastal residents of Gualala chuckled about any notion of affordability, given an influx of newcomers who've driven the Gualala's median housing price to \$580,000 compared to the median family income of \$47,778. Other Gualala residents recalled an era when the Gualala region was better known for the logging of ancient redwoods, marijuana growing and boisterous beer drinking at the historic Gualala Hotel. Still there was a certain pride to the Men's Journal magazine's designation. Yvette White, a 25-year Gualala resident who works at the Gualala Sport; Tackle shop, said she's proud her town made it on the list.

**Previous paragraph from Document:** [previous paragraph]**Instruction:**

Rewrite the below paragraph by resolving all entity coreferences with the preceding paragraph from document.

- Resolve all inter-sentence pronoun references.
- Make sure that all pronouns in a sentence refers to some named entity with in the same sentence.
- Explicitly mention entity names wherever necessary to remove ambiguity from a sentence. Remember to make each sentence clear and unambiguous.
- For each entity, use only the one most informative name.
- Do not generate anything except the rewritten paragraph.

**Paragraph:** [paragraph ]**Output:**

Figure 6. Prompt for rewriting a paragraph (e.g., a document chunk) by resolving entity coreferences.

Extract all facts from the document. For each fact, also generate all mentioned named entities.

#### Instructions

- Generate only the most informative name for each named entity. Example: if John P. Parker, John Parker are coreferential, only generate John Parker.
- Consistently use the most informative name for each named entity in all facts.
- You must not include pronouns or ambiguous references in facts. Directly include all relevant named entities.
- Respond using json format provided below:

```
{
  "f1": {
    "fact": "A factual statement describing important information (preferably about some entities) from the paragraph",
    "entities": ["entity 1", "entity 2", "entity 3"]
  },
  "f2": {},
}
```

Below is an example for reference.

**Paragraph:** Locked in a heated battle with Novo Nordisk's semaglutide franchise, Eli Lilly's tirzepatide is beginning to come into its own—both with regards to sales and amid attempts to show the dual GIP/GLP-1 agonist can strike out beyond diabetes and obesity. As Mounjaro, tirzepatide won its first FDA nod in Type 2 diabetes back in May 2022. An obesity approval followed last November, with that formulation of tirzepatide adopting the commercial moniker Zepbound. In 2023's fourth quarter, Mounjaro generated a whopping \$2.2 billion in sales, a nearly eight-fold increase over the \$279 million it pulled down during the same stretch in 2022. Year-to-date, the drug brought home around \$5.2 billion in revenues, Lilly said in an earnings release Tuesday. Zepbound, for its part, generated \$175.8 million during its first quarter on the market. Overall, Lilly reeled in around \$9.4 billion in fourth-quarter sales, growing 28% over the \$7.3 billion it made for the quarter in 2022.

**Entities:** Eli Lilly, Novo Nordisk, Tirzepatide, Semaglutide, dual GIP/GLP-1 agonist, FDA, Mounjaro, Zepbound, diabetes, Type 2 diabetes, Obesity

```
Output:
{
  "f1": {
    "fact": "Eli Lilly's tirzepatide is competing with Novo Nordisk's semaglutide franchise.",
    "entities": ["Eli Lilly", "Novo Nordisk", "tirzepatide", "semaglutide"]
  },
  "f2": {
    "fact": "Eli Lilly is trying to show tirzepatide, the dual GIP/GLP-1 agonist, can strike out beyond diabetes and obesity.",
    "entities": ["tirzepatide", "dual GIP/GLP-1 agonist", "diabetes", "obesity"]
  },
  "f3": {
    "fact": "Tirzepatide, under the brand name Mounjaro, received its first FDA approval for Type 2 diabetes in May 2022.",
    "entities": ["tirzepatide", "Mounjaro", "FDA", "Type 2 diabetes"]
  },
  "f4": {
    "fact": "Tirzepatide, under the brand name Zepbound, received an obesity approval in November 2022.",
    "entities": ["tirzepatide", "Zepbound", "obesity"]
  },
  "f5": {
    "fact": "Mounjaro generated $2.2 billion in sales in the fourth quarter of 2023, an eight-fold increase from the $279 million during the same period in 2022.",
    "entities": ["Mounjaro"]
  },
  "f6": {
    "fact": "Mounjaro brought in around $5.2 billion in revenues year-to-date in 2023, Lilly said in an earnings release Tuesday",
    "entities": ["Mounjaro", "Eli Lilly"]
  },
  "f7": {
    "fact": "Zepbound generated $175.8 million in sales in its first quarter on the market.",
    "entities": ["Zepbound"]
  },
  "f8": {
    "fact": "Eli Lilly's fourth-quarter sales were around $9.4 billion, a 28% increase over the $7.3 billion during the same period in 2022.",
    "entities": ["Eli Lilly"]
  }
}
```

Remember to put the response in code block and do not generate any extra information only the JSON output.

Figure 7. Prompt for extracting propositions and entities from a document.

## Notes

Nan Zhang: Work done while interning at Salesforce AI Research.

## Footnotes

<sup>1</sup> <https://huggingface.co/datasets/BAAI/IndustryCorpus>

<sup>2</sup> <https://huggingface.co/meta-llama/Meta-Llama-3-70B-Instruct>

<sup>3</sup> <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

## References

1. Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W, Rocktäschel T, Riedel S, Kiela D. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In: Larochelle H, Ranzat



- o M, Hadsell R, Balcan MF, Lin H, editors. *Advances in Neural Information Processing Systems*. Curran Associates, Inc.; 2020. p. 9459–9474. Available from: [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/6b4\\_93230205f780e1bc26945df7481e5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/6b4_93230205f780e1bc26945df7481e5-Paper.pdf).
2. <sup>△</sup>Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, Dai Y, Sun J, Wang M, Wang H (2024). "Retrieval-Augmented Generation for Large Language Models: A Survey". arXiv. Available from: <https://arxiv.org/abs/2312.10997>.
3. <sup>△</sup>Wang X, Wang Z, Gao X, Zhang F, Wu Y, Xu Z, Shi T, Wang Z, Li S, Qian Q, Yin R, Lv C, Zheng X, Huang X (2024). "Searching for Best Practices in Retrieval-Augmented Generation". arXiv. [arXiv:2407.01219 \[cs.CL\]](https://arxiv.org/abs/2407.01219).
4. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub>, <sup>△</sup><sub>c</sub>, <sup>△</sup><sub>d</sub>, <sup>△</sup><sub>e</sub>, <sup>△</sup><sub>f</sub> Sarthi P, Abdullah S, Tuli A, Khanna S, Goldie A, Manning CD (2024). "Raptor: Recursive abstractive processing for tree-organized retrieval". arXiv preprint arXiv:2401.18059.
5. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub> Edge D, Trinh H, Cheng N, Bradley J, Chao A, Mody A, Truitt S, Larson J (2024). "From local to global: A graph rag approach to query-focused summarization". arXiv preprint arXiv:2404.16130.
6. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub> Liu NF, Lin K, Hewitt J, Paranjape A, Bevilacqua M, Petroni F, Liang P (2024). "Lost in the middle: How language models use long contexts". *Transactions of the Association for Computational Linguistics*. 12: 157–173.
7. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub>, <sup>△</sup><sub>c</sub> Trivedi H, Balasubramanian N, Khot T, Sabharwal A (2022). "MuSiQue: Multihop Questions via Single-hop Question Composition". arXiv. Available from: <https://arxiv.org/abs/2108.00573>.
8. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub> Ho X, Duong Nguyen AK, Sugawara S, Aizawa A. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. In: Scott D, Bel N, Zong C, editors. *Proceedings of the 28th International Conference on Computational Linguistics*. Barcelona, Spain (Online): International Committee on Computational Linguistics; 2020. p. 6609–6625. doi:[10.18653/v1/2020.coling-main.580](https://doi.org/10.18653/v1/2020.coling-main.580). Available from: <https://aclanthology.org/2020.coling-main.580>.
9. <sup>△</sup><sub>a</sub>, <sup>△</sup><sub>b</sub> Yang Z, Qi P, Zhang S, Bengio Y, Cohen W, Salakhutdinov R, Manning CD. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In: Riloff E, Chiang D, Hockenmaier J, Tsujii J, editors. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics; 2018. p. 2369–2380. doi:[10.18653/v1/D18-1259](https://doi.org/10.18653/v1/D18-1259). Available from: <https://aclanthology.org/D18-1259>.
10. <sup>△</sup>Chen D, Fisch A, Weston J, Bordes A. Reading Wikipedia to answer open-domain questions. In: Barzilay R, Kan M-Y, editors. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada: Association for Computational Linguistics; 2017. p. 1870–1879. doi:[10.18653/v1/P17-1171](https://doi.org/10.18653/v1/P17-1171). <https://aclanthology.org/P17-1171>.
11. <sup>△</sup>Karpukhin V, Oguz B, Min S, Lewis P, Wu L, Edunov S, Chen D, Yih W. Dense passage retrieval for open-domain question answering. In: Webber B, Cohn T, He Y, Liu Y, editors. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics; 2020. p. 5397–5408. doi:[10.18653/v1/2020.emnlp-main.438](https://doi.org/10.18653/v1/2020.emnlp-main.438).

020. p. 6769–6781. doi:[10.18653/v1/2020.emnlp-main.550](https://doi.org/10.18653/v1/2020.emnlp-main.550). Available from: <https://aclanthology.org/2020.emnlp-main.550>.
12. <sup>^</sup>Yu W, Iter D, Wang S, Xu Y, Ju M, Sanyal S, Zhu C, Zeng M, Jiang M (2023). "Generate rather than Retrieve: Large Language Models are Strong Context Generators". arXiv. <https://arxiv.org/abs/2209.10063>.
  13. <sup>^</sup>Sun Z, Wang X, Tay Y, Yang Y, Zhou D (2023). "Recitation-Augmented Language Models". arXiv. <https://arxiv.org/abs/2210.01296>.
  14. <sup>a, b</sup>Chen T, Wang H, Chen S, Yu W, Ma K, Zhao X, Yu D, Zhang H (2023). "Dense x retrieval: What retrieval granularity should we use?" arXiv preprint arXiv:2312.06648. [arXiv:2312.06648](https://arxiv.org/abs/2312.06648).
  15. <sup>^</sup>Nair I, Garimella A, Srinivasan BV, Modani N, Chhaya N, Karanam S, Shekhar S. A neural CRF-based hierarchical approach for linear text segmentation. In: Vlachos A, Augenstein I, editors. Findings of the Association for Computational Linguistics: EACL 2023. Dubrovnik, Croatia: Association for Computational Linguistics; 2023. p. 883–893. doi:[10.18653/v1/2023.findings-eacl.65](https://doi.org/10.18653/v1/2023.findings-eacl.65). Available from: <https://aclanthology.org/2023.findings-eacl.65>.
  16. <sup>a, b, c</sup>Gutiérrez BJ, Shu Y, Gu Y, Yasunaga M, Su Y (2024). "HippoRAG: Neurobiologically Inspired Long-Term Memory for Large Language Models". arXiv preprint arXiv:2405.14831. Available from: <https://arxiv.org/abs/2405.14831>.
  17. <sup>^</sup>Zhang J, Silvescu A, Honavar V. Ontology-driven induction of decision trees at multiple levels of abstraction. In: Abstraction, Reformulation, and Approximation: 5th International Symposium, SARA 2002 Kananaskis, Alberta, Canada August 2–4, 2002 Proceedings 5. Springer; 2002. p. 316–323.
  18. <sup>^</sup>Maekawa A, Hirao T, Kamigaito H, Okumura M. Can we obtain significant success in RST discourse parsing by using large language models? In: Graham Y, Purver M, editors. Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers). St. Julian's, Malta: Association for Computational Linguistics; 2024. p. 2803–2815. Available from: <https://aclanthology.org/2024.eacl-long.171>.
  19. <sup>^</sup>Liu Z, Shi K, Chen NF (2021). "DMRST: A joint framework for document-level multilingual RST discourse segmentation and parsing". arXiv preprint arXiv:2110.04518. [arXiv:2110.04518](https://arxiv.org/abs/2110.04518).
  20. <sup>^</sup>Pang RY, Parrish A, Joshi N, Nangia N, Phang J, Chen A, Padmakumar V, Ma J, Thompson J, He H, Bowman S. "QuALITY: Question Answering with Long Input Texts, Yes!" In: Carpuat M, de Marneffe MC, Meza Ruiz IV, editors. Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Seattle, United States: Association for Computational Linguistics; 2022. p. 5336–5358. doi:[10.18653/v1/2022.naacl-main.391](https://doi.org/10.18653/v1/2022.naacl-main.391). Available from: <https://aclanthology.org/2022.naacl-main.391>.

21. <sup>A</sup>Laban P, Fabbri AR, Xiong C, Wu C-S (2024). "Summary of a haystack: A challenge to long-context llms and rag systems". arXiv preprint arXiv:2407.01370.
22. <sup>A</sup>Liu Y, Fabbri AR, Zhao Y, Liu P, Joty S, Wu C-S, Xiong C, Radev D (2023). "Towards interpretable and efficient automatic reference-based summarization evaluation". arXiv preprint arXiv:2303.03608. Available from: <https://arxiv.org/abs/2303.03608>.
23. <sup>A</sup>Choubey PK, Su X, Luo M, Peng X, Xiong C, Le T, Rosenman S, Lal V, Mui P, Ho R, Howard P, Wu CS (2024). "Distill-SynthKG: Distilling Knowledge Graph Synthesis Workflow for Improved Coverage and Efficiency". arXiv. arXiv:2410.16597 [cs.CL].
24. <sup>A</sup>AI@Meta (2024). "Llama 3 Model Card". [https://github.com/meta-llama/llama3/blob/main/MODEL\\_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md).
25. <sup>A</sup>Jiang AQ, Sablayrolles A, Mensch A, Bamford C, Chaplot DS, de las Casas D, Bressand F, Lengyel G, Lample G, Saulnier L, Lavaud LR, Lachaux MA, Stock P, Le Scao T, Lavril T, Wang T, Lacroix T, El Sayed W (2023). "Mistral 7B". arXiv. Available from: <https://arxiv.org/abs/2310.06825>.
26. <sup>a</sup>, <sup>b</sup>Santhanam K, Khattab O, Saad-Falcon J, Potts C, Zaharia M. ColBERTv2: Effective and efficient retrieval via a lightweight late interaction. In: Carpuat M, de Marneffe MC, Meza Ruiz IV, editors. Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Seattle, United States: Association for Computational Linguistics; 2022. p. 3715–3734. doi:10.18653/v1/2022.naacl-main.272. Available from: <https://aclanthology.org/2022.naacl-main.272>.
27. <sup>A</sup>Robertson SE, Walker S. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In: SIGIR'94: Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, organised by Dublin City University. Springer; 1994. p. 232–241.
28. <sup>a</sup>, <sup>b</sup>, <sup>c</sup>Press O, Zhang M, Min S, Schmidt L, Smith N, Lewis M. Measuring and narrowing the compositionality gap in language models. In: Bouamor H, Pino J, Bali K, editors. Findings of the Association for Computational Linguistics: EMNLP 2023. Singapore: Association for Computational Linguistics; 2023. p. 5687–5711. doi:10.18653/v1/2023.findings-emnlp.378. Available from: <https://aclanthology.org/2023.findings-emnlp.378>.
29. <sup>A</sup>Islam SB, Rahman MA, Hossain KSMT, Hoque E, Joty S, Parvez MR. Open-RAG: Enhanced retrieval augmented reasoning with open-source large language models. In: Al-Onaizan Y, Bansal M, Chen Y-N, editors. Findings of the Association for Computational Linguistics: EMNLP 2024. Miami, Florida, USA: Association for Computational Linguistics; 2024. p. 14231–14244. doi:10.18653/v1/2024.findings-emnlp.831. <https://aclanthology.org/2024.findings-emnlp.831>.
30. <sup>A</sup>Tang Y, Yang Y (2024). "MultiHop-RAG: Benchmarking Retrieval-Augmented Generation for Multi-Hop Queries". arXiv. <https://arxiv.org/abs/2401.15391>.

31. <sup>^</sup>Stelmakh I, Luan Y, Dhingra B, Chang MW. "ASQA: Factoid Questions Meet Long-Form Answers." In: Goldberg Y, Kozareva Z, Zhang Y, editors. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics; 2022. p. 8273–8288. doi:[10.18653/v1/2022.emnlp-main.566](https://doi.org/10.18653/v1/2022.emnlp-main.566). Available from: <https://aclanthology.org/2022.emnlp-main.566>.

## Declarations

**Funding:** No specific funding was received for this work.

**Potential competing interests:** No potential competing interests to declare.