

# Hierarchical Retrieval with Evidence Curation for Open-Domain Financial Question Answering on Standardized Documents

Jaeyoung Choe, Jihoon Kim, Woohwan Jung

Department of Applied Artificial Intelligence, Hanyang University  
{cgy9100, skygl, whjung}@hanyang.ac.kr

## Abstract

Retrieval-augmented generation (RAG) based large language models (LLMs) are widely used in finance for their excellent performance on knowledge-intensive tasks. However, standardized documents (e.g., SEC filing) share similar formats such as repetitive boilerplate texts, and similar table structures. This similarity forces traditional RAG methods to misidentify near-duplicate text, leading to duplicate retrieval that undermines accuracy and completeness. To address these issues, we propose the **Hierarchical Retrieval with Evidence Curation (HiREC)** framework. Our approach first performs hierarchical retrieval to reduce confusion among similar texts. It first retrieve related documents and then selects the most relevant passages from the documents. The evidence curation process removes irrelevant passages. When necessary, it automatically generates complementary queries to collect missing information. To evaluate our approach, we construct and release a **Large-scale Open-domain Financial (LOFin)** question answering benchmark that includes 145,897 SEC documents and 1,595 question-answer pairs. Our code and data are available at <https://github.com/deep-over/LOFin-bench-HiREC>.

## 1 Introduction

Retrieval-augmented generation (RAG) (Lewis et al., 2020) with large language models (LLMs) have significantly improved performance in knowledge-intensive tasks. Due to its ability to improve both factual accuracy and timeliness, extensive research (Yepes et al., 2024; Sarmah et al., 2024) has investigated applying RAG in searching financial information where accurate and up-to-date information is crucial for decision-making.

Financial documents, such as annual reports (10-K) in SEC filings, are highly structured and follow standardized templates across companies and periods, often containing similar tables and repetitive

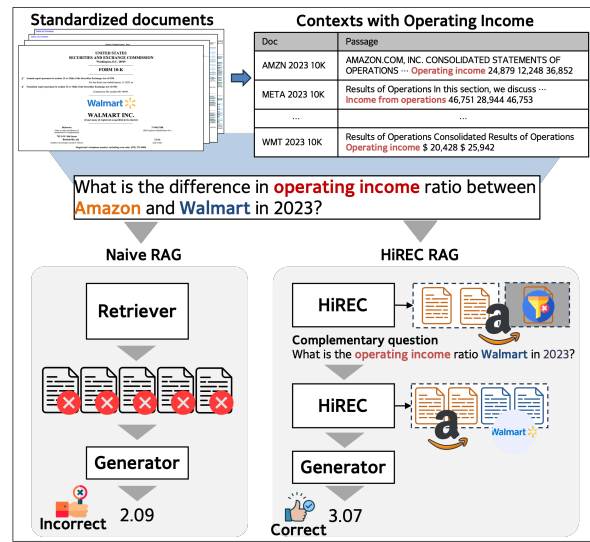


Figure 1: Comparison of a naive RAG approach and HiREC .

narratives. As shown in Figure 1, 2023 10-K reports from Amazon, Meta, and Walmart exhibit nearly identical table structures with similar titles and indicators, differing mainly in numerical values. Consequently, when asked, *What is the difference in operating income ratio between Amazon and Walmart in 2023?*, a conventional RAG system may struggle to distinguish among these similar passages, retrieving irrelevant or redundant information that leads to inaccurate answers.

To tackle these challenges from the standardized format of financial documents, we propose the HiREC (Hierarchical Retrieval and Evidence Curation) framework. HiREC consists of two main components: hierarchical retrieval and evidence curation. The hierarchical retrieval first retrieves related documents and then selects the most relevant passages from the documents, thereby reducing confusion from near-duplicate text. As illustrated in Figure 1, narrowing the candidate set to documents from Amazon and Walmart enables the system to focus on highly relevant passages. However, the prevalence of comparative questions in

financial QA often leads to incomplete retrieval of essential evidence. To address this, the evidence curation stage filters out irrelevant passages and generates complementary queries when necessary. For example if only Amazon operating income is retrieved then a complementary query is generated to fetch Walmart operating income so that all necessary evidence is gathered for an accurate answer.

To evaluate our approach, we assess QA performance in an open-domain setting. Existing financial question-answering benchmarks (Islam et al., 2023; Lai et al., 2024) rely on small-scale corpora that include at most about 1,300 documents and very limited test sets, which do not reflect realistic financial scenarios. We propose LOFin (Large-scale Open-domain Financial QA), a comprehensive financial question-answering benchmark built on a large-scale corpus containing approximately 145,000 SEC filings from companies in the S&P 500. LOFin includes 1,595 open-domain question-answering test instances and addresses challenges in standardized document retrieval, such as near-duplicate tables and repetitive narratives, that are not evident in smaller datasets. In addition, we release the entire benchmark as open-source to support future research in the field.

Experimental results indicate that HiREC improves performance by at least 13% compared to existing RAG methods. In addition, our framework consistently outperforms commercial LLMs with web search engines such as SearchGPT (OpenAI, 2025) and Perplexity (Perplexity, 2023).

Our contributions are threefold:

- We propose a hierarchical retrieval approach that retrieves related documents and identifies the most pertinent passages, thereby reducing confusion caused by near-duplicate content in standardized financial documents.
- We introduce an evidence curation process that filters out irrelevant passages and generates complementary queries when necessary, effectively supplementing missing information for accurate financial QA.
- We present LOFin, a large-scale, realistic financial QA benchmark that exposes challenges in standardized document retrieval, and we release it as open-source.

## 2 Related Works

**Financial QA and RAG** Financial QA has advanced significantly, with recent benchmarks em-

phasizing numeric reasoning and table understanding. TAT-QA (Zhu et al., 2021) and FinQA (Chen et al., 2021) provide single-page tabular contexts, while DocFinQA (Reddy et al., 2024) and DocMath-Eval (Zhao et al., 2024) extend this to multi-page settings. However, these benchmarks remain closed-domain, limiting their applicability for RAG systems. Open-domain benchmarks have also been proposed. For instance, consider Financebench (Islam et al., 2023) and SEC-QA (Lai et al., 2024). However, these datasets are built on small-scale document collections and suffer from limited test set sizes or the lack of publicly available fixed test sets.

With the emergence of financial QA, research on financial RAG has also been progressing. There are graph-based methods (Barry et al., 2025; Sarmah et al., 2024) tailored to the financial domain as well as hybrid approaches (Wang, 2024). In addition, studies on financial document chunking offer valuable insights (Yepes et al., 2024).

**Retrieval augmented generation** RAG has evolved in diverse ways (Gao et al., 2023; Zhang et al., 2025). Hierarchical retrieval was proposed for cases where sections are clearly demarcated, typically employing a two-step document-passage process (Arivazhagan et al., 2023; Chen et al., 2024). The difference from previous studies is that while they segment documents into sections, we segment at the level of individual documents. There are also studies that utilize filtering to enhance the quality of the contexts input into RAG (Zhuang et al., 2024; Wang et al., 2024b). Unlike previous approaches that required training, we manage quality using only an LLM. Iterative retrieval is typically proposed for multi-hop QA. A standard iterative method uses the context retrieved in the first step as part of the query for subsequent iterations (Trivedi et al., 2022; Shao et al., 2023). Self-RAG (Asai et al., 2023) also conducts iterative retrieval that includes the generation process. However, our approach does not utilize previously retrieved context because we specifically need to discover missing information.

## 3 Large-scale Open-domain Financial QA

In this section, we present LOFin, a benchmark that overcomes the limitations of current financial QA datasets by expanding the retrieval corpus and increasing open-domain QA pairs (see Table 1).

### 3.1 Large-scale Document Collection

To reflect a real-world scenario where retrieval and QA must be performed over a large volume of documents, we collected a comprehensive set of SEC filings. Specifically, we gathered 10-K, 10-Q, and 8-K filings from the SEC EDGAR<sup>1</sup> system, covering S&P 500 companies from October 2001 to April 2025.

We converted the HTML documents to PDF using the wkhtmltopdf<sup>2</sup> library. Following the approach of Islam et al., 2023, we used the PyMuPDF library<sup>3</sup> to extract text at the page level from these PDFs, excluding reports that lacked proper page separation. The final corpus consists of 145,897 reports from 516 companies.

### 3.2 Open-domain QA Pair Construction

In this section, we detail our process for constructing open-domain QA pairs by leveraging three established financial QA benchmarks: FinQA (Chen et al., 2021), Financebench (Islam et al., 2023), and SEC-QA (Lai et al., 2024).

We begin by converting the closed-setting questions from FinQA into an open-domain format. First, we exclude any test questions for which evidence documents were not collected due to page separation or collection period issues (35 out of 1147). Next, we transform the remaining questions using GPT-4o by appending relevant period and company information. For example, the question *what are the total operating expenses for 2016?* is transformed into *Could you provide the total operating expenses reported by Lockheed Martin for the year 2016?*, with Lockheed Martin explicitly added to enhance context. The conversion prompt used is provided in Appendix B.2.

Subsequently, we identify candidate evidence pages using a two-step process. First, we compute BM25 similarity scores between the FinQA gold table context and the content of each page in the candidate document, explicitly considering the distinct numerical values in the table. Then, we employ an NLI model<sup>4</sup> to measure the semantic similarity between the FinQA gold context (excluding table content) and each page. If both methods select the same top candidate page we accept it as evidence and verify its correctness. Otherwise

<sup>1</sup><https://www.sec.gov/search-filings>

<sup>2</sup><https://wkhtmltopdf.org/>

<sup>3</sup><https://pymupdf.readthedocs.io/>

<sup>4</sup><https://huggingface.co/sentence-transformers/nli-mpnet-base-v2>

Benchmark	Open	Multi-Doc	# QAs	# Docs
TAT-QA			1,669	-
FinQA			1,147	-
DocFinQA			922	-
Financebench	✓		150	368
SEC-QA	✓	✓	N/A	1,315
LOFin	✓	✓	1,595	145,897

Table 1: Comparison of financial QA benchmarks. “# QAs” is the test set size, and “# Docs” is the number of documents in the retrieval corpus. SEC QA does not have a fixed QA count, as it provides a question generation framework.

we manually annotate the correct page to ensure accurate evidence labeling (see Appendix B.3 for further details).

Financebench is designed for an open-domain setting and its questions are adopted without modification. In contrast both FinQA and Financebench mainly include single-document questions which limits their ability to evaluate multi-document retrieval and reasoning. To address this limitation, we adopt multi-document question templates from SEC-QA. We select the four that are designed for multi-document scenarios. We then manually craft the associated questions and annotate the answers along with the corresponding evidence pages. This process enhances our open-domain QA pairs with challenges that require multi-document and multi-hop reasoning.

Following the ACL ARR review, we finalized the LOFin benchmark with a total of 1,595 QA pairs (initially 1,389), reflecting the addition of 205 newly created QA pairs based on recent SEC filings. We refer to the initial version as LOFin-1.4k and the expanded version as LOFin-1.6k to clearly distinguish between the two. These newly added questions follow the same annotation protocol and are designed to promote more complex reasoning across multiple documents. For details of the expanded LOFin-1.6k benchmark, see Appendix A; for the SEC-QA templates used in its construction, refer to Appendix B.4.

## 4 Hierarchical Retrieval with Evidence Curation (HiREC) Framework

In this section, we introduce the HiREC framework. Figure 2 shows the overall framework and process. The framework comprises two main components: hierarchical retrieval and evidence curation. During the hierarchical retrieval stage a hierarchical approach retrieves passages  $\mathcal{P}_r$  that are relevant to

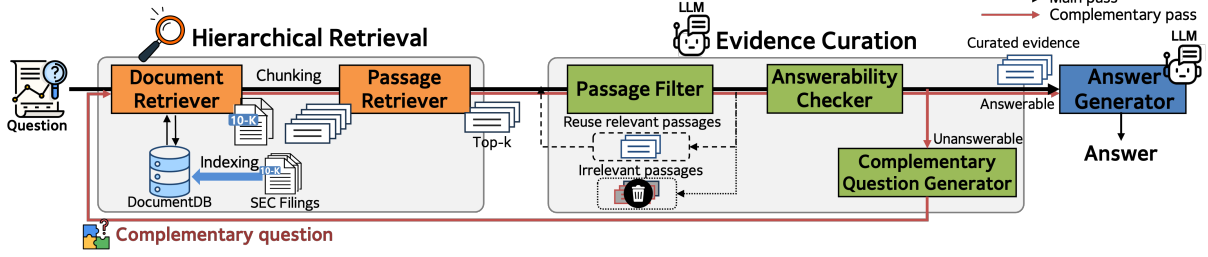


Figure 2: Overview of hierarchical retrieval with evidence curation framework

the question  $q$ . In the evidence curation process the retrieved passages are filtered to retain only those directly pertinent and then evaluated to determine whether they provide sufficient information to answer the question. If the information is insufficient a complementary question  $q_c$  is generated to reinitiate the retrieval process (complementary pass). Otherwise if the evidence is sufficient the filtered passage set  $\mathcal{P}_f$  is forwarded to the Answer Generator to produce the final answer (main pass). When the maximum iteration  $i_{\max}$  is reached evidence curation halts and the passages retrieved using  $q_c$  are merged with the passages filtered passages to generate the answer. The pseudocode of HiREC is described in Algorithm 1. For the LLM-based components of our framework, we use instruction-style prompts tailored to each module’s objective. Appendix D provides the full prompt templates and design process.

## 4.1 Hierarchical Retrieval

Standardized documents use uniform templates with repetitive structures and similar content, which makes retrieving distinct and relevant passages challenging. We address this issue using a hierarchical approach. First we retrieve documents relevant to the question (4.1.1) to narrow the search space, and then we select pertinent passages within those documents (4.1.2).

### 4.1.1 Document Retriever

**Document indexing.** Documents contain extensive context, and their standardized format makes it difficult to capture all important details in a single vector. To address this, we extract and index key distinguishing information when retrieving documents. In financial reports, the cover page provides essential details such as the company name, report type, and fiscal period. For each document  $d \in \mathcal{D}$ , we generate a cover page summary  $d'$  using an LLM (the prompt is detailed in Appendix D.1), pre-compute its embedding with a bi-encoder (Wang

### Algorithm 1 HiREC framework

---

**Require:** A question  $q$ , a corpus  $\mathcal{D}$ , maximum number of iterations  $i_{\max}$

- 1:  $\mathcal{P}_f \leftarrow \emptyset$  // Initialization
- 2:  $\mathcal{P}_r \leftarrow \text{HIERARCHICALRETRIEVAL}(q, \mathcal{D})$  // Sec 4.1
- 3: **for**  $i = 1..i_{\max}$  **do**
- 4:    $(\mathcal{P}_f, q_c, y) \leftarrow \text{EVIDENCECURATION}(q, \mathcal{P}_r)$  // Sec 4.2
- 5:   **if**  $y$  is Answerable **then**
- 6:     **return** ANSWERGENERATION( $q, \mathcal{P}_f$ ) // Sec 4.3
- 7:   **end if**
- 8:    $\mathcal{P}_r \leftarrow \text{HIERARCHICALRETRIEVAL}(q_c, \mathcal{D})$  // Sec 4.1
- 9:    $\mathcal{P}_r \leftarrow \mathcal{P}_r \cup \mathcal{P}_f$
- 10: **end for**
- 11: **return** ANSWERGENERATION( $q, \mathcal{P}_r$ ) // Sec 4.3

---

et al., 2024a), and index the resulting vector as  $\mathbf{v}_d = E^D(d')$  in the document store.

**Document retrieval process.** Our document retrieval process consists of three stages. First, given a question  $q$ , we use an LLM to convert it into a refined query  $q'$ . This conversion reduces issues caused by extraneous financial terms (e.g., tickers or service names like Google or Facebook) that may hinder effective retrieval. (Details of the transformation prompt are provided in Appendix D.2.) Second, we perform dense retrieval. We compute the vector representation of  $q'$  using the same bi-encoder employed during document indexing, yielding  $\mathbf{v}_{q'} = E^D(q')$ . The relevance between  $q'$  and each document  $d$  is then determined via  $s_{q',d}^D = \mathbf{v}_{q'}^\top \mathbf{v}_d$  (Karpukhin et al., 2020); based on these scores, we retrieve  $k'_D$  candidate documents  $\mathcal{D}_{\text{cand}}$ . Finally, we rerank these candidates using a cross-encoder (He et al., 2021) by computing  $\text{CrossEncoder}^D(q', d)$  and select the top  $k_D$  documents  $\mathcal{D}_r$ . This multi-stage process ultimately produces the final set of documents  $\mathcal{D}_r$  that are most relevant to the original question  $q$ .

### 4.1.2 Passage Retriever

Within the final set of retrieved documents  $\mathcal{D}_r$ , the passage retriever evaluates each passage  $p$  by computing a score using a cross-encoder,  $\text{CrossEncoder}^P(q, p)$ . It then selects the top  $k_P$  passages to form the final set  $\mathcal{P}_r$ . By reducing the



number of passages processed by the cross-encoder, we enable real-time computation while still taking advantage of its superior ability to capture inter-sentence relationships compared to a bi-encoder. However, passage retrievers that are pretrained on general text typically have difficulty handling financial tables. For instance, when retrieving tables, attributes such as titles, periods, and indicators are more important than the numerical values, yet these cues are not well captured. To address this, we finetuned the model on table data using FinQA training set where tables serve as evidence. Specifically, for each question  $q$  and its associated evidence document  $d$ , we denote the tables on the evidence page as the evidence passage set  $\mathcal{X}$ . For each evidence passage  $p \in \mathcal{X}$ , we sample  $n_{neg}$  negative passages, where negative passages are defined as tables that appear on pages other than the evidence page. The cross-encoder is then trained with a binary cross-entropy loss (Nogueira and Cho, 2019) defined as

$$\mathcal{L} = \sum_{(q,p) \in \mathcal{X}} \left[ -\log(\text{CrossEncoder}^P(q,p)) - \sum_{p' \in \mathcal{P}^-} \log(1 - \text{CrossEncoder}^P(q,p')) \right],$$

where the cross-encoder applies an internal sigmoid to produce scores in  $[0, 1]$ .

## 4.2 Evidence Curation

Financial questions often involve comparisons between different time periods or companies. Even when the retrieval process selects relevant passages, some critical information may be missing. Furthermore, retrieved passages can contain irrelevant data that hinders overall performance (Liu et al., 2024; Xu et al., 2024). To overcome these issues, we introduce an evidence curation process that filters out incorrect data and fills information gaps by initiating additional retrieval when necessary. Our process comprises three modules: a passage filter that removes irrelevant passages (Section 4.2.1), an answerability checker that assesses evidence sufficiency (Section 4.2.2), and a complementary question generator that formulates a supplementary query if necessary (Section 4.2.3). We use an LLM to perform all three tasks in a single response (see Appendix D.3 for the prompt).

### 4.2.1 Passage Filter

The passage filter removes passages from the retrieved set  $\mathcal{P}_r$  that are not relevant to the question, yielding a filtered passage set  $\mathcal{P}_f$  containing at most

$k'_p$  passages. The filter considers both newly retrieved passages and those previously identified as relevant from earlier iterations. This step is crucial, as the inclusion of noisy, irrelevant passages can lead the LLM to generate inaccurate responses.

### 4.2.2 Answerability Checker

The answerability checker evaluates whether the filtered passages  $\mathcal{P}_f$  provide sufficient evidence to answer the question. If they are deemed adequate,  $\mathcal{P}_f$  and the question are forwarded to the answer generation stage; otherwise, the lack of sufficient information triggers a complementary iteration to retrieve additional data.

### 4.2.3 Complementary Question Generator

The complementary question generator examines the filtered passages  $\mathcal{P}_f$  to identify gaps in the evidence needed to answer the question. It then generates a supplementary query  $q_c$ , which is used as input for the next retrieval.

## 4.3 Answer Generation

In the answer generation stage, the relevant passages and the original question serve as inputs to a reasoning process that derives the final answer. For questions requiring numerical calculations, a Program-of-Thought (PoT) (Chen et al., 2022) reasoning method is employed, while for text-based inferential questions, a Chain-of-Thought (CoT) (Wei et al., 2022) approach is applied. This dual strategy is particularly effective for financial documents, which are rich in numerical data and tables, ensuring comprehensive and accurate reasoning. See Appendix D.4 for prompt details, adapted from DocMath-Eval (Zhao et al., 2024).

# 5 Experiments

## 5.1 Experimental Settings

**Dataset.** We evaluate the open-domain QA methods on the LOFin benchmark proposed in Section 3. All main experiments in this section are conducted on LOFin-1.4k, which contains 1,389 QA pairs. Results on the updated benchmark LOFin-1.6k are reported separately in Appendix A.

For the purpose of the main experimental analysis, we categorize the question-answer pairs into three groups based on their format and context:

- **Numeric (Table):** The answer is a number from a table or can be calculated from numbers in tables.

Dataset	Numeric (Table)		Numeric (Text)		Textual		Average		
	Page Recall	Answer Acc	Page Recall	Answer Acc	Page Recall	Answer Acc	Page Recall	Answer Acc	k
GPT-4o (Zero-shot)	-	3.82	-	2.93	-	35.00	-	13.92	-
Perplexity (Perplexity, 2023)	-	2.51	-	5.13	-	24.00	-	10.55	-
Self-RAG (Asai et al., 2023)	19.96	1.86	24.18	4.03	12.75	17.00	18.96	7.63	10.0
RQ-RAG (Chan et al., 2024)	18.61	1.97	19.05	2.56	17.96	20.50	18.54	8.34	36.0
IRCoT (Trivedi et al., 2022) $\diamond$	28.17	19.10	34.62	27.84	12.67	20.00	25.15	22.31	20.0
HybridSearch (Wang, 2024) $\diamond$	26.75	19.10	32.05	30.77	14.37	27.00	24.39	25.62	10.0
HHR (Arivazhagan et al., 2023) $\diamond$	37.67	26.53	40.29	32.97	21.98	26.50	33.31	28.67	10.0
Dense (Karpukhin et al., 2020) $\diamond$	37.69	23.69	40.48	32.97	26.18	31.00	34.78	29.22	10.0
HiREC (Ours) $\diamond$	<b>50.17</b>	<b>37.23</b>	<b>53.48</b>	<b>48.35</b>	<b>32.39</b>	<b>41.50</b>	<b>45.35</b>	<b>42.36</b>	3.7
Gold evidence $\diamond$	100	64.96	100	69.23	100	63.50	100	65.90	1.3

Table 2: Main evaluation results on LOFin-1.4k . Methods marked with  $\diamond$  use GPT-4o for generation. Here,  $k$  denotes the average number of input passages used during generation, and Gold evidence shows only the correct page. **Bold** indicates the highest performance.

- **Numeric (Text):** The answer is a number derived by extracting and combining numerical information from text.
- **Textual:** The answer is a textual explanation.

Examples and statistics for these categories are in Table 12.

**Implementation details.** Our system leverages several pre-trained and fine-tuned models across its components. For the passage retriever, we fine-tune a DeBERTa-v3 model (He et al., 2021) with  $n_{neg} = 8$  for negative sampling, a batch size of 128, 3 epochs, and a learning rate of  $2 \times 10^{-7}$  on a single GeForce RTX 4090 GPU. For document retrieval, we employ the E5 model (Wang et al., 2024a) as the bi-encoder and use DeBERTa-v3 as the reranker. The answer generator is powered by OpenAI’s GPT-4o, while other LLM-based tasks (query transformation, document summarization, and evidence curation) are handled by Qwen-2.5-7B-Instruct (Yang et al., 2024).

Our framework runs for a maximum of  $i_{max} = 3$  iterations. The document retriever retrieves  $k_D = 5$  documents, and the passage retriever extracts  $k_P = 5$  passages, with the passage filter capping the output at  $k'_P = 10$ . Additional hyperparameters are provided in Appendix C.1.

**Baseline methods.** We compare our proposed method with several state-of-the-art RAG approaches, including RQ-RAG (Chan et al., 2024) and Self-RAG (Asai et al., 2023). To ensure a fair comparison with our approach—which employs GPT-4o as the answer generator—we use the same answer generator for the latest retrieval algorithms: Dense, HybridSearch (Wang, 2024), HHR (Ari-

vazhagan et al., 2023), and IRCoT (Trivedi et al., 2022). In addition, we evaluate against Perplexity (Perplexity, 2023), a commercial LLM service that combines web search with LLMs. RQ-RAG, Self-RAG, and IRCoT employ iterative LLM-based retrieval, while Dense serves as a strong baseline, it employs OpenAI’s text-embedding-3-small as its encoder for DPR, with results reranked by DeBERTa-v3. All methods construct passages by concatenating the title and content during retrieval. Detailed configurations for each baseline are provided in Appendix C.2.

**Metrics.** We evaluate numeric answers for accuracy by considering rounding and truncation, while textual answers are evaluated using GPT-4o and FAMMA prompts (Xue et al., 2024) (see Appendix C.3). We measure retrieval performance at the page level, using recall and precision against ground-truth evidence pages as metrics. In particular, since chunk locations or units can vary for Page recall, we standardize them at the page level to ensure consistent performance measurement.

## 5.2 Main Result

Table 2 shows the retrieval performance (page recall) and final answer accuracy for HiREC and the baselines. Our approach outperforms all baselines, achieving at least 10% higher page recall and 13% higher answer accuracy than the second-best model, Dense. The result validates the effectiveness of our method in retrieval for standardized documents. Furthermore, HiREC retrieves an average of only 3.7 passages, demonstrating its efficiency in selecting high-quality evidence through evidence curation.

Method	Page Precision	Page Recall	Answer Accuracy
HiREC	21.79	45.35	42.36
w/o HR	14.75	34.16	32.76
w/o EC	4.70	41.41	36.70
w/o Fine-tuning	21.07	42.77	40.13
w/o Filter	8.43	50.19	42.08

Table 3: Ablation study




Figure 3: Comparison of company, document, page error rates for HiREC and baselines.

Notably, in the textual category, the answer accuracies are higher than the page recalls for all methods. This suggests that LLMs can often answer text-based questions correctly even when retrieval is incomplete. Numeric questions require more precise reasoning and the table category remains especially challenging.

### 5.3 Analysis

**Ablation Study** Table 3 shows page precision, page recall and answer accuracy for HiREC when each component is removed. Hierarchical retrieval (HR) and evidence curation (EC) denote the two components. The setting w/o HR corresponds to the outcome of applying the Dense method in conjunction with EC, while the setting w/o EC represents the initial retrieval performance of HR ( $k_P = 10$ ).

The w/o Fine-tuning setting uses an unfine-tuned reranker to address fairness concerns related to table-specific prior knowledge. Despite a slight drop in performance, HiREC still outperforms the Dense baseline by over 10% in accuracy. Even when Dense is paired with a fine-tuned reranker, it only reaches an average AnswerAcc of 30.55%, maintaining a similar performance gap.

HR is critical for enhancing retrieval accuracy because its absence produces the lowest performance. In addition, the initial search performance of HR exceeds that of the dense method combined with EC. The w/o EC setting demonstrates the effectiveness of both passage filter and the complementary




Figure 4: Recall, precision, and passages per query by iteration. EC stands for Evidence curation.




Figure 5: Precision-recall curve (recall on X-axis, precision on Y-axis)

question generator. HR supported by EC delivers enhanced precision, recall and answer accuracy. Furthermore, results from the w/o Filter setting reveal the advantages of the complement and filtering. Although the complementary component attains the highest recall score in the absence of filtering, its accuracy is lower than that of HiREC with filtering. This outcome is attributed to the inclusion of incorrect information that causes conflicts when filtering is not applied (Xu et al., 2024).

**Error type analysis.** Figure 3 presents the error count for cases in which the company indicated in the document is incorrect for the retrieved passage of each method. A company error is defined as an instance in which both the passage and the document contain incorrect company information, and the results are reported based solely on the top-1 retrieved passage. HiREC achieves the fewest errors by correctly identifying companies during document retrieval, which in turn ensures accurate passage retrieval.

**Evidence curation by iteration.** Figure 4 shows that iterative evidence curation (EC) enhances page recall and precision while reducing passages per query compared to the initial hierarchical retrieval (w/o EC). As iterations progress, retrieval performance steadily improves and efficiency increases, demonstrating evidence curation effectiveness.

Method	Retrieval		Generation	
	Tokens	Cost (\$)	Tokens	Cost (\$)
Dense	-	-	2,666	9.3
			/ 177	/ 2.5
IRCoT	9,610	33.4	3,475	12.1
	/ 1,372	/ 19.1	/ 177	/ 2.5
HiREC	4,291	*14.9	1,052	3.7
	/ 313	/ *4.4	/ 159	/ 2.2

Table 4: Cost-efficiency comparison. Each cell shows the values for the input/output. The asterisk (\*) indicates that, although a much smaller open-source LLM was used, the cost is computed using GPT-4o’s pricing.

**Precision-recall.** Figure 5 presents the precision-recall curve, where  $k$  varies from 1 to 50 for baseline methods. As expected, increasing  $k$  improves recall but lowers precision due to the retrieval of less relevant passages. HiREC achieves both higher precision and recall, consistently exceeding the maximum values observed across all  $k$  in the baseline range. This highlights ability of HiREC to retrieve relevant information more effectively while maintaining accuracy.

**Cost-efficiency.** Table 4 presents the average input/output token counts and total API cost incurred during the retrieval and generation processes. HiREC achieves the highest performance while using significantly fewer tokens and lower cost than the baselines by filtering out irrelevant passages before the answer generation stage. Moreover, compared to IRCoT in retrieval, HiREC uses far fewer tokens, demonstrating that its filtering and complementary question generation enable efficient iteration. Finally, our results show that even a relatively small LLM can effectively perform curation without incurring high costs.

#### 5.4 Analysis of Performance Across Various LLM Generators

We thoroughly analyze the effectiveness of the HiREC framework across diverse LLM generator models. For this analysis, additional evaluations were conducted using open-source models such as DeepSeek-R1-Distill-Qwen-14B (Guo et al., 2025) and Qwen-2.5-7B-Instruct (Yang et al., 2024) as generators, in place of GPT-4o.

As shown in Table 5, HiREC consistently demonstrates superior QA performance across various LLM generators, proving its robustness. Notably, even when employing smaller open-source models, HiREC configurations outperform the Dense baseline. HiREC +Deepseek-14B achieved

Method	Numeric (Table)	Numeric (Text)	Textual	Average
<b>Dense</b>				
+ Qwen-2.5-7B	19.76	22.34	29.52	23.87
+ Deepseek-14B	25.44	31.87	35.00	30.77
+ GPT-4o	23.69	26.18	31.00	29.22
<b>HiREC</b>				
+ Qwen-2.5-7B	<u>27.62</u>	<u>33.33</u>	<u>36.00</u>	<u>32.32</u>
+ Deepseek-14B	<u>34.39</u>	<u>41.39</u>	<u>40.53</u>	<u>38.76</u>
+ GPT-4o	<b>37.23</b>	<b>48.35</b>	<b>41.50</b>	<b>42.36</b>

Table 5: Performance comparison across various LLM generators. Underlined values denote HiREC configurations utilizing smaller generators that surpass Dense + GPT-4o performance.

Method	Financebench		FinQA		SEC-QA	
	Page Recall	Answer Acc	Page Recall	Answer Acc	Page Recall	Answer Acc
Self-RAG	9.00	13.33	22.93	3.24	4.59	4.72
RQ-RAG	13.67	18.67	20.32	2.88	9.38	4.72
IRCoT	9.33	20.00	32.24	22.66	4.20	7.09
HybridSearch	10.00	29.33	29.95	23.29	10.42	7.87
HHR	16.44	34.00	41.50	29.50	10.13	5.51
Dense	26.11	33.33	40.38	27.61	15.70	9.45
<b>HiREC</b>	<b>40.00</b>	<b>50.00</b>	<b>52.49</b>	<b>41.61</b>	<b>19.97</b>	<b>14.17</b>

Table 6: Performance results by data source.

over 9% higher average answer accuracy compared to Dense+GPT-4o. Another key implication is that integrating sLLMs in the retrieval stage proves to be an effective strategy for reducing overall inference costs while maintaining strong performance.

#### 5.5 Performance results by data source

Our analysis of various data sources aims to evaluate benchmark data leakage risk and demonstrate the robustness of the HiREC framework.

As shown in Table 6, HiREC consistently demonstrates superior performance across all data sources. Notably, all methods, including HiREC, exhibit relatively lower performance on the SEC-QA subset, as multi-document QA tasks require the synthesis of information from multiple sources for accurate answers. These results suggest the LOFin benchmark is more challenging due to multi-document scenarios and mitigates leakage risk by relying on retrieved evidence.

The final results for the LOFin benchmark, including the expanded SEC-QA subset, can be found in the Appendix.

#### 5.6 Comparison of Commercial LLMs with Web Search

Commercial systems such as SearchGPT and Perplexity combine LLMs with web search to answer



Step	Content
<b>Question</b>	<i>What is Adobe’s year-over-year change in unadjusted operating income from <b>FY2015</b> to <b>FY2016</b> (in units of percents and round to one decimal place)? Give a solution to the question by using the income statement.</i>
<b>Initial Retrieval</b>	<ul style="list-style-type: none"> <li>✓ <b>ADBE_2015_10K (p.59)</b>: 3,112,300 3,045,960 Operating income 903,095 412,685 422,723 Non-operating ...</li> <li>✗ <b>ADBE_2015_10K (p.39)</b>: Net income of \$629.6 million increased by \$361.2 million, or 135%, during fiscal 2015 ...</li> <li>✗ <b>ADBE_2015_10K (p.45)</b>: Amortization of purchased intangibles 68.7 ... Total operating expenses \$3,148.1 ...</li> </ul>
<b>Evidence Curation</b>	<p><b>Relevant passages:</b> ADBE_2015_10K (p.59)    <b>Answerable:</b> False</p> <p><b>Explanation:</b> The context does not provide operating income figures for both <b>FY2015</b> and <b>FY2016</b>. We only have operating income for <b>FY2015</b> and some intermediate values, but no direct comparison between <b>FY2015</b> and <b>FY2016</b>.</p> <p><b>Complementary question:</b> <i>What is Adobe’s operating income for <b>FY2016</b>?</i></p>
<b>Complementary Retrieval</b>	<ul style="list-style-type: none"> <li>✓ <b>ADBE_2016_10K (p.61)</b>: 3,148,099 3,112,300 Operating income 1,493,602 903,095 412,685 Non-operating ...</li> <li>✗ <b>ADBE_2016_10K (p.35)</b>: ITEM 6. SELECTED FINANCIAL DATA ... Net income \$1,168,782 \$ 629,551 ...</li> <li>✗ <b>ADBE_2017_10K (p.57)</b>: CONSOLIDATED STATEMENTS OF INCOME ... Operating income 2,168,095 ...</li> </ul>
<b>Evidence Curation</b>	<b>Relevant passages:</b> ADBE_2015_10K (p.59), ADBE_2016_10K (p.61) <b>Answerable:</b> True
<b>Generation</b>	<b>Ground Truth Answer:</b> 46% <b>Generated Answer:</b> 46 <b>Correct:</b> ✓

Table 7: Case study illustration of HiREC framework effectiveness

Method	Numeric (Table)	Numeric (Text)	Textual	Average
Perplexity	2.5	5.0	37.5	15.0
SearchGPT	15.0	32.5	35.0	27.5
HiREC	30.0	65.0	45.0	46.7

Table 8: Answer accuracy for Perplexity, SearchGPT, and HiREC on 40 samples per category.

questions using financial data. Table 8 compares these systems with our approach where Perplexity uses the llama-3.1-sonar-large-128k-online model and SearchGPT uses GPT-4o. In our evaluation based on 40 questions per category, HiREC consistently outperforms these baselines especially on numeric questions, indicating that although commercial systems effectively retrieve relevant documents they often miss the precise numerical details necessary for accurate computation.

## 5.7 Case Study

We analyze how evidence curation filters out unnecessary information and effectively retrieves missing details, and how this impacts the final results. Table 7 presents the retrieval and evidence curation outcomes over iterations for a given question. The initial pass of hierarchical retrieval retrieves passages containing operating income for Adobe for FY2015, but the question remains unanswerable due to missing FY2016 data. The answerability checker detects this gap and triggers the complementary question generator, leading to a complementary pass that retrieves the missing FY2016 operating income. With complete evidence, the

answer generator successfully computes the 46% year-over-year change, highlighting the effectiveness of iterative refinement in HiREC.

## 6 Conclusion

We introduced HiREC, a retrieval-augmented framework for question answering over standardized financial documents. Our hierarchical retrieval reduces confusion from repeated boilerplate content, while evidence curation filters irrelevant passages and recovers missing information. To evaluate under realistic open-domain conditions, we constructed LOFin, a large-scale benchmark with 1,595 QA pairs across 145,000 SEC filings. The dataset includes multi-document and multi-hop questions that go beyond prior financial QA benchmarks. Experiments show that HiREC consistently outperforms state-of-the-art baselines in both retrieval quality and answer accuracy, while maintaining cost efficiency through selective passage use. Overall, our findings suggest that HiREC provides a scalable and effective solution for open-domain QA over complex, standardized financial documents.

## Limitations

In our study, we use LLMs for query transformation, evidence curation, and answer generation, making our approach dependent on the performance of the LLMs. We utilize a relatively small LLM, Qwen 2.5 7B (Yang et al., 2024), compared to commercial models.

The benchmarks we employ, FinQA and FinanceBench, are publicly available datasets. Therefore, it is possible that pre-trained large language models have already been exposed to these datasets during their training. In Table 2, GPT-4o (Zero-shot) achieves the second-highest textual performance.

## Acknowledgements

This study was supported by the National R&D Program for Cancer Control through the National Cancer Center (NCC), funded by the Ministry of Health & Welfare, Republic of Korea (RS-2025-02264000).

This work was also supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2023-00261068, Development of Lightweight Multimodal Anti-Phishing Models and Split-Learning Techniques for Privacy-Preserving Anti-Phishing), and by the Convergence Security Core Talent Training Program (IITP-2024-RS-2024-00423071), supervised by IITP and also funded by the Ministry of Science and ICT (MSIT), Korea.

## References

- Manoj Ghuhana Arivazhagan, Lan Liu, Peng Qi, Xinchu Chen, William Yang Wang, and Zhiheng Huang. 2023. [Hybrid hierarchical retrieval for open-domain question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10680–10689, Toronto, Canada. Association for Computational Linguistics.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. *arXiv preprint arXiv:2310.11511*.
- Mariam Barry, Gaëtan Caillaut, Pierre Halftermeyer, Raheel Qader, Mehdi Mouayad, Fabrice Le Deit, Dimritri Cariolaro, and Joseph Gesnouin. 2025. Graphrag: Leveraging graph-based efficiency to minimize hallucinations in llm-driven rag for finance data. In *Proceedings of the Workshop on Generative AI and Knowledge Graphs (GenAIK)*, pages 54–65.
- Chi-Min Chan, Chunpu Xu, Ruibin Yuan, Hongyin Luo, Wei Xue, Yike Guo, and Jie Fu. 2024. Rq-rag: Learning to refine queries for retrieval augmented generation. *arXiv preprint arXiv:2404.00610*.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.
- Xinyue Chen, Pengyu Gao, Jiangjiang Song, and Xiaoyang Tan. 2024. Hiqa: A hierarchical contextual augmentation rag for massive documents qa. *arXiv preprint arXiv:2402.01767*.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, et al. 2021. Finqa: A dataset of numerical reasoning over financial data. *arXiv preprint arXiv:2109.00122*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- Pranab Islam, Anand Kannappan, Douwe Kiela, Rebecca Qian, Nino Scherrer, and Bertie Vidgen. 2023. Financebench: A new benchmark for financial question answering. *arXiv preprint arXiv:2311.11944*.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.
- Viet Dac Lai, Michael Krumbick, Charles Lovering, Varshini Reddy, Craig Schmidt, and Chris Tanner. 2024. Sec-qa: A systematic evaluation corpus for financial qa. *arXiv preprint arXiv:2406.14394*.
- 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*, 33:9459–9474.
- 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*, 12:157–173.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*.

- OpenAI. 2025. searchgpt: A comprehensive guide. <https://www.openai.com/searchgpt> [Accessed: 2025-02-10].
- Perplexity. 2023. Perplexity.ai. <https://www.perplexity.ai/>. [Large language model].
- Nicholas Pipitone and Ghita Hourir Alami. 2024. Legalbench-rag: A benchmark for retrieval-augmented generation in the legal domain. *arXiv preprint arXiv:2408.10343*.
- Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai, Michael Krumbick, Charles Lovering, and Chris Tanner. 2024. Docfinqa: A long-context financial reasoning dataset. *arXiv preprint arXiv:2401.06915*.
- Bhaskarjit Sarmah, Dhagash Mehta, Benika Hall, Rohan Rao, Sunil Patel, and Stefano Pasquali. 2024. Hybridrag: Integrating knowledge graphs and vector retrieval augmented generation for efficient information extraction. In *Proceedings of the 5th ACM International Conference on AI in Finance*, pages 608–616.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy. *arXiv preprint arXiv:2305.15294*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. *arXiv preprint arXiv:2212.10509*.
- Jing Wang. 2024. *Financerag with hybrid search and reranking*.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024a. Multilingual e5 text embeddings: A technical report. *arXiv preprint arXiv:2402.05672*.
- Yuhao Wang, Ruiyang Ren, Junyi Li, Wayne Xin Zhao, Jing Liu, and Ji-Rong Wen. 2024b. Rear: A relevance-aware retrieval-augmented framework for open-domain question answering. *arXiv preprint arXiv:2402.17497*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024. Knowledge conflicts for llms: A survey. *arXiv preprint arXiv:2403.08319*.
- Siqiao Xue, Tingting Chen, Fan Zhou, Qingyang Dai, Zhixuan Chu, and Hongyuan Mei. 2024. Famma: A benchmark for financial domain multilingual multimodal question answering. *arXiv preprint arXiv:2410.04526*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Antonio Jimeno Yepes, Yao You, Jan Milczek, Sebastian Laverde, and Renyu Li. 2024. Financial report chunking for effective retrieval augmented generation. *arXiv preprint arXiv:2402.05131*.
- Qinggong Zhang, Shengyuan Chen, Yuanchen Bei, Zheng Yuan, Huachi Zhou, Zijin Hong, Junnan Dong, Hao Chen, Yi Chang, and Xiao Huang. 2025. A survey of graph retrieval-augmented generation for customized large language models. *arXiv preprint arXiv:2501.13958*.
- Yilun Zhao, Yitao Long, Hongjun Liu, Ryo Kamoi, Linyong Nan, Lyuhao Chen, Yixin Liu, Xiangru Tang, Rui Zhang, and Arman Cohan. 2024. DocMath-eval: Evaluating math reasoning capabilities of LLMs in understanding long and specialized documents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16103–16120, Bangkok, Thailand. Association for Computational Linguistics.
- Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. Tat-qa: A question answering benchmark on a hybrid of tabular and textual content in finance. *arXiv preprint arXiv:2105.07624*.
- Ziyuan Zhuang, Zhiyang Zhang, Sitao Cheng, Fangkai Yang, Jia Liu, Shujian Huang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2024. EfficientRAG: Efficient retriever for multi-hop question answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3392–3411, Miami, Florida, USA. Association for Computational Linguistics.

## Appendix

### A Experiments on the Expanded LOFin-1.6k Benchmark

#### A.1 Motivation and Dataset Extension

The initial version of the LOFin-1.4k benchmark was heavily biased toward single-document QA, particularly from FinQA, which limited its coverage of complex financial analysis scenarios. To address this limitation and incorporate reviewer feedback, we expanded the SEC-QA subset from 127 to 333 question-answer pairs by adding 206 new examples. All newly added QA pairs were constructed from recent SEC filings collected up to April 2025, and are designed to require multi-document and multi-page reasoning. The questions include inter-company comparisons, trend analysis, and temporal reasoning. All examples were manually created, which prevents the possibility of data leakage.

Dataset Source	FinQA	SEC-QA	Financebench	Total
# QAs	1,112	333	150	1,595
Ratio (%)	69.72	20.88	9.40	100

Table 9: Distribution of expanded LOFin-1.6k QA pairs by dataset source.

##### A.1.1 Experiments

This section reports the results on the expanded benchmark LOFin-1.6k, following the same settings as in Section 5. The evaluation compares our proposed model HiREC against the second-best baseline, Dense. Table 10 presents the main results by answer type. HiREC consistently outperforms Dense across all categories. Table 11 shows performance by data source. Although SEC-QA is a more challenging task, HiREC still demonstrates superior results compared to the baseline. The category distribution by answer type is summarized in Table 13, and the dataset composition by source is reported in Table 9.

##### A.1.2 Discussion and Conclusion

The experimental results on the extended SEC-QA subset in expanded LOFin-1.6k validate the robustness and effectiveness of HiREC under realistic multi-document and multi-page QA settings. With this expansion, the LOFin-1.6k benchmark now represents a significantly more challenging and realistic task compared to existing financial QA datasets. Overall, these results demonstrate that

HiREC effectively performs evidence-based reasoning and provide strong empirical support that addresses reviewer concerns regarding dataset diversity, integrity, and model generalization.

### B Supplementary Details on the LOFin Benchmark

#### B.1 Categorization Details

Tables 12 and 13 summarize the distribution of QA pairs in LOFin by answer type before and after the SEC-QA expansion. The original dataset is heavily skewed toward numeric-table questions, while the expanded version introduces a larger portion of textual and multi-hop reasoning questions, illustrating the increased complexity and diversity of the benchmark.

#### B.2 FinQA Open-Domain Conversion Prompt

You are a financial AI Assistant. The following financial questions are provided without including the company names and document period. Rewrite the questions to include the given information.

- Maintain the original meaning of the question while allowing for varied expressions that enable accurate and open-ended responses without altering the factual content.
- The question must include the company name and the report year.
- Ticker is used as information about the company name and should not be included in the question.
- While rephrasing the question, integrate additional information naturally into the sentence without using simple connectors like 'In', 'for', or 'according to'.

Document information:

- company ticker: {ticker}
- document\_period: {period}

Question: {question}

Output format is ##new\_question:

Table 14: This prompt template is used to convert closed-domain FinQA questions into an open-domain format by seamlessly integrating company names and report years into the questions, while preserving their original meaning.

#### B.3 FinQA Gold Page Selection

Our evidence page annotation for FinQA relies on a dual-method approach. Using FinQA’s qid, we first identify candidate documents and discard



Model	Numeric Table			Numeric Text			Textual			Average		
	Precision	Recall	Accuracy	Precision	Recall	Accuracy	Precision	Recall	Accuracy	Precision	Recall	Accuracy
Dense	3.47	33.88	22.92	3.77	37.73	31.50	2.90	16.18	23.93	3.38	29.27	26.12
HiREC	<b>25.63</b>	<b>48.76</b>	<b>35.57</b>	<b>21.64</b>	<b>52.75</b>	<b>44.32</b>	<b>13.54</b>	<b>23.59</b>	<b>26.95</b>	<b>20.27</b>	<b>41.70</b>	<b>35.61</b>

Table 10: Performance by answer type on the expanded LOFin-1.6k benchmark.

Model	FinanceBench			FinQA			SEC-QA			Average		
	Precision	Recall	Accuracy	Precision	Recall	Accuracy	Precision	Recall	Accuracy	Precision	Recall	Accuracy
Dense	1.80	16.00	36.00	4.00	37.28	26.44	2.88	12.66	13.51	2.81	21.98	25.00
HiREC	<b>15.67</b>	<b>37.33</b>	<b>47.33</b>	<b>25.44</b>	<b>51.18</b>	<b>39.44</b>	<b>12.74</b>	<b>18.67</b>	<b>14.85</b>	<b>17.95</b>	<b>35.73</b>	<b>33.63</b>

Table 11: Performance by data source on the expanded LOFin-1.6k benchmark.

Category	Numeric (Text)	Numeric (Table)	Textual	Total
# QAs	273	916	200	1389
Ratio (%)	19.65	65.95	14.40	100

Table 12: Distribution of initial LOFin QA pairs by category.

Category	Numeric (Text)	Numeric (Table)	Textual	Total
# QAs	273	925	397	1595
Ratio (%)	17.11	58.00	24.89	100

Table 13: Distribution of expanded LOFin QA pairs by category.

any questions associated with documents that have page separation issues. Within each candidate document, we compute BM25 similarity between the FinQA gold table context (which emphasizes distinct numerical values) and each page’s content to select candidate evidence pages. Additionally, an NLI model evaluates the semantic similarity between the concatenated pre\_text and post\_text (i.e., the context surrounding the table) and the content of each page. If both methods select the same top candidate, that page is accepted and subsequently verified; otherwise, manual annotation is performed to ensure accuracy. Out of 1147 instances, 894 (78%) were automatically accepted. However, among these accepted pages, 9 instances (approximately 1% of 894) contained errors due to OCR issues that merged content from two pages into one. In the remaining 218 instances (19%), the top pages selected by the two methods differed, necessitating manual annotation, and an additional 35 cases (3%) were discarded during preprocessing. This combined approach ensures robust and accurate identification of the correct evidence pages for FinQA.

## B.4 Question Templates of SEC-QA

- How much common dividends did {company} pay in the last {num\_year} years in US dollars?
- What is the percentage difference of {company1}’s {metric} compared to that of {company2}?
- What is {company}’s overall revenue growth over the last {num\_year}-year period?
- Among {company\_names}, what is the {metric2} of the company that has the highest {metric1}?

## C Experimental Settings

### C.1 Hyperparameters

- Candidate document count  $k'_D$  : 100
- Final document count  $k_D$  : 5
- Final passage count  $k_P$  : 5
- Maximum relevant passages  $k'_P$  : 10
- Maximum iterations  $i_{max}$  : 3
- LLM temperature: 0.01
- Chunking tool: Langchain<sup>5</sup>’s RecursiveCharacterTextSplitter
- Chunk size: 1024
- Overlap between chunks: 30

### C.2 Baselines

**Perplexity:** The Perplexity baseline employs the llama-3.1-sonar-large-128k-online model, integrating web search to supplement the context used for answer generation.

**Self-RAG:** Self-RAG utilizes a 13B model in a Long-form setting. It leverages the Contriever model to retrieve 10 context passages for each query, ensuring comprehensive coverage of the relevant information.

<sup>5</sup><https://www.langchain.com/>

**RQ-RAG:** `text-embedding-3-small` is used in RQ-RAG to retrieve three context passages at each retrieval step. Configured for multi-hop QA with an exploration depth of two, it employs iterative retrieval to continuously refine the contextual information used for answer generation.

**HybridSearch:** HybridSearch combines BM25-based scores with dense scores computed using the `text-embedding-3-small` model. This hybrid approach adheres to its original configuration to effectively balance lexical and semantic matching.

**IRCoT:** IRCoT relies on a BM25-based retriever and employs GPT-4o for chain-of-thought (CoT) reasoning. It performs iterative retrieval by executing up to 3 iterations, with 5 passages being retrieved during each iteration.

**HHR:** HHR implements a hybrid retrieval strategy that combines dense and sparse retrieval methods. In this approach, the `text-embedding-3-small` model is used for dense retrieval, and the system first retrieves the top 5 documents; within each document, the top 10 passages are then selected.

**Dense:** The Dense method uses the `text-embedding-3-small` model to initially retrieve the top 50 relevant passages. These passages are subsequently reranked using a DeBERTa-v3 model to determine their final order.

### C.3 Textual Evaluation Prompt

This prompt is used to evaluate the factual correctness of generated textual answers by instructing the LLM to compare them against the ground truth within the provided context.

```
You are a highly knowledgeable expert and teacher in the finance domain.
You are reviewing a student's answers to financial questions.
You are given the context, the question, the student's answer and the student's explanation and the ground-truth answer.
Please use the given information and refer to the ground-truth answer to determine if the student's answer is correct.
The input information is as follows:
context: '{gold context}'
question: '{question}'
ground-truth answer: '{answer}'
student's answer: '{generated answer}'
Please respond directly as either 'correct' or 'incorrect'.
```

Table 15: GPT-4o LM Evaluation Prompt Template

## D LLM Prompts

This section presents a detailed overview of the prompting strategies employed in our framework. Our prompt templates are designed based on common instruction formats introduced in the Prompt Engineering Guide<sup>6</sup>, and were initially generated using GPT-4o in an instruction-following mode. To ensure consistency with each module's output format, we manually refined the initial prompts through a lightweight post-editing process.

No explicit constraints were imposed during prompt construction. We did not apply few-shot prompting due to practical limitations such as prompt length and inference cost. The final prompts are tailored to the specific needs of each module—query rewriting, document summarization, evidence curation, and answer generation—and are provided in this section.

These prompt designs play a central role in facilitating efficient information processing and enabling accurate responses to complex financial questions. By disclosing our prompting approach in detail, we aim to enhance the reproducibility and transparency of our overall system.

### D.1 Summarization Prompt

To assist document indexing during hierarchical retrieval (section 4.1.1), we use a concise LLM-based summarization prompt as follows Table 16:

```
You are a helpful assistant. Summarize the following text:
```

Table 16: Summarization prompt

<sup>6</sup><https://www.promptingguide.ai/>

## D.2 Query Transformation Prompt

To reduce noise from overly specific or ambiguous financial terms in user queries, we rewrite them into focused retrieval queries using the prompt below [Table 17](#):

```
You are an AI that rewrites user questions about financial topics into concise meta-focused queries.

1) Identify the key financial terms or metrics in the question.
2) Determine which type of documents typically contain those terms.
3) Transform the user’s question into a short query referencing the financial terms and the relevant documents.
4) Do not reveal the transformation process or provide examples.
5) Output only the final rewritten query.

## Question: {Question}

### Output format
## Query: {Rewritten query}
```

Table 17: Financial query transformation instructions

## D.3 Evidence Curation Prompt

The evidence curation module jointly performs three tasks—passage filtering, answerability checking, and complementary question generation—using a single prompt as follows [Table 18](#):

```
### Instruction
You are a financial expert. Evaluate the provided context to determine if it contains enough information to answer the given question.
1. Read the context carefully and decide if it contains enough information to answer the question.
2. If it is answerable, set 'is_answerable: answerable' and provide the answer in 'answer'.
3. If it is not answerable, set 'is_answerable: unanswerable'. Then:
- List the relevant document IDs in 'answerable_doc_ids' in order of relevance (from most to least relevant).
- Explain what specific information is missing in 'missing_information'.
- Provide a concise question in 'refined_query' to search for exactly that missing information.
4. Output your result strictly in the specified format below using '##' headers.

###Inputs
Context:
Context1 (ID: 1): Title is {title1}. Content is {content1}
Context2 (ID: 2): Title is {title2}. Content is {content2}
...
Question: {question}

### Output format
## is_answerable: answerable or unanswerable
## missing_information: If 'unanswerable', specify the details or data needed; if 'answerable', None
## answer: If 'answerable', provide the answer; if 'unanswerable', then None
## answerable_doc_ids: Provide a list of document IDs that contain relevant information (e.g., [1, 2]). If none, use []
## refined_query: If 'unanswerable', provide a refined question to obtain the missing information;
```

Table 18: Evidence curation instructions

## D.4 Generation Prompt

We adopt two generation styles depending on the question type. For numerical reasoning tasks, we use Program-of-Thoughts (PoT) prompting to generate Python code (see [Table 19](#)); for textual reasoning, we apply Chain-of-Thought (CoT) prompting (see [Table 20](#)).

**[System Input]**

You are a financial expert, you are supposed to generate a Python program to answer the given question based on the provided financial document context. The returned value of the program is supposed to be the answer.

```
```python
def solution():
    # Define variables name and value based on
    the given context
    guarantees = 210
    total_exposure = 716

    # Do math calculation to get the answer
    answer = (guarantees / total_exposure) *
    100

    # return answer
    return answer
```
```

**[User Input]**

Context:

Sources: {title1} - {content1}

Sources: {title2} - {content2}

...

Question: {question}

Please generate a Python program to answer the given question. The format of the program should be the following:

```
```python
def solution():
    # Define variables name and value based on
    the given context
    ...
    # Do math calculation to get the answer
    ...
    # return answer
    return answer
```
```

Continue the program to answer the question. The returned value of the program is supposed to be the answer:

```
```python
def solution():
    # Define variables name and value based on
    the given context
    ```
```

Table 19: The prompt for generating a Python program to answer a financial question.

**[System Input]**

You are a financial expert, and you are supposed to answer the given question based on the provided financial document context. You need to first think through the problem step by step, documenting each necessary step. Then, you are required to conclude your response with the final answer in your last sentence as "Therefore, the answer is final answer."

**[User Input]**

Context:

Sources: {title1} - {content1}

Sources: {title2} - {content2}

...

Question: {question}

Let's think step by step to answer the given question.

Table 20: Chain-of-Thought prompt for generating a Python program to answer a financial question.

## E Generalization to Other Domains

To evaluate the effectiveness of our method beyond the financial domain, we assess retrieval performance on the LegalBench-RAG (Pipitone and Alami, 2024) dataset, a legal-domain benchmark where retrieval is particularly challenging due to domain-specific terminology and structured documents. Using a mini question set, we conduct retrieval over the full document corpus to simulate realistic conditions.

As shown in Table 21, HiREC achieves the highest Precision@k (41.42) and Recall@k (48.50) while using the lowest average number of retrieved passages per query ( $k = 3.2$ ). This result highlights the framework's ability to efficiently retrieve essential content with minimal redundancy. In particular, HiREC demonstrates strong performance on structured datasets such as MAUD and CUAD, outperforming other models by a large margin in both P@k and R@k.

ContractNLI is a document-level natural language inference task in which each contract is evaluated against a fixed set of hypotheses. Unlike passage-based QA tasks, it requires global reasoning over the full document. As a result, Dense may benefit from broad recall, potentially retrieving passages that touch on relevant hypotheses by chance. In contrast, HiREC follows a structured evidence curation pipeline that emphasizes precision, often selecting fewer but more targeted documents. This structural difference helps explain the smaller performance gap between HiREC and Dense on ContractNLI compared to other tasks.



Model	PrivacyQA		ContractNLI		MAUD		CUAD		Average		Avg. $k$
	P@k	R@k	P@k	R@k	P@k	R@k	P@k	R@k	P@k	R@k	
Hybrid	21.13	35.99	12.06	34.36	1.34	2.05	7.73	21.97	10.57	23.59	5.0
HHR	20.31	35.57	19.28	54.25	8.04	14.20	8.66	25.49	14.07	32.38	5.0
Dense	<u>31.55</u>	<b>52.21</b>	<u>28.56</u>	<b>78.69</b>	<u>11.03</u>	<u>20.87</u>	<u>10.82</u>	<u>29.92</u>	<u>20.49</u>	<u>45.42</u>	5.0
HiREC	<b>55.79</b>	<u>49.94</u>	<b>35.46</b>	<u>53.31</u>	<b>33.70</b>	<b>36.80</b>	<b>40.74</b>	<b>53.95</b>	<b>41.42</b>	<b>48.50</b>	<b>3.2</b>

Table 21: Average Precision and Recall for the LegalBench-RAG dataset. The best performance values are highlighted in bold and the second-best are underlined.

Overall, these results confirm the generalizability of HiREC to non-financial domains, demonstrating robustness across diverse legal document types and retrieval tasks.

## F Failure Case Study

**Retrieval Failure Case** In this case, the model fails to retrieve all relevant documents across multiple fiscal years. The question explicitly asks for the total common dividends paid over the past three years (as of 2023), which requires retrieving information from multiple annual filings. However, the model only retrieves passages from the 2023 report. This indicates difficulty in reasoning over temporal scope and aggregating evidence when relevant information is distributed across documents.

**Generation Failure Case** In this case, the model retrieves a passage containing sufficient information to answer the question. However, the generated answer is incorrect due to faulty arithmetic reasoning. Specifically, the free cash flow should be computed as the difference between cash from operations and capital expenditures. Although both values are retrieved, the model produces an incorrect value, indicating limitations in table-grounded numerical reasoning and code execution.

Step	Content
<b>Question</b>	<i>How much common dividends did Bank of America pay in the last 3 years, as of 2023, in millions of US dollars?</i>
<b>Gold Evidences</b>	BAC_2023_10K (p.94), BAC_2022_10K (p.94), BAC_2021_10K (p.94)
<b>Retrieval</b>	<p>✗ <b>BAC_2023_10K (p.141)</b>: NOTE 13 Shareholders' Equity Common Stock ...</p> <p>✗ <b>BAC_2023_10K (p.150)</b>: Defined Contribution Plans The Corporation...</p>
<b>Generation</b>	<b>Ground Truth Answer:</b> 25,718 <b>Generated Answer:</b> 2,560

Table 22: Retrieval failure: model fails to retrieve documents from multiple years required to answer a multi-year aggregation question.

Step	Content
<b>Question</b>	<i>According to the information provided in the statement of cash flows, what is the FY2020 free cash flow (FCF) for General Mills? Answer in USD millions.</i>
<b>Gold Evidences</b>	GIS_2020_10K (p.51)
<b>Retrieval</b>	<p>✓ <b>GIS_2020_10K (p.51)</b>: ... \$3,676.2... Capital expenditures \$460.8 ...</p> <p>✗ <b>GIS_2020_10K (p.17)</b>: Adjusted diluted EPS of \$3.61 increased 12 percent on a constant-currency basis ...</p> <p>✗ <b>GIS_2020_10K (p.36)</b>: Free cash flow \$3,215.4 ... Net cash provided by operating activities conversion ...</p>
<b>LLM Output</b>	cash_from_operations = 3700 # in USD millions capex = 461 # in USD millions
<b>Generation</b>	<b>Ground Truth Answer:</b> 3,215 <b>Generated Answer:</b> 3,239

Table 23: Generation failure: despite correct retrieval, the model fails to compute free cash flow accurately from table values.

Category	Examples
Numeric (Table)	<p><b>Question:</b> What percentage of the estimated purchase price for Hologic in 2007 is attributed to the net tangible assets?</p> <p><b>Answer:</b> 3.8%</p> <p><b>Evidence:</b> HOLX_2007_10K (p.110)  The components and initial allocation of the purchase price consist of:  Net tangible assets acquired as of September 18, 2007: \$2,800  Developed technology and know how: \$12,300  Customer relationship: \$17,000  Trade name: \$2,800  ...  Estimated Purchase Price: \$73,200.</p>
Numeric (Text)	<p><b>Question:</b> During the years 2004 to 2006, what was the average impairment on construction in progress, expressed in millions, for American Tower Corporation as reported in their 2006 financial documents?</p> <p><b>Answer:</b> 2.63</p> <p><b>Evidence:</b> AMT_2006_10K (p.106)  Construction-In-Progress Impairment Charges—For the years ended December 31, 2006, 2005 and 2004, the Company wrote-off approximately \$1.0 million, \$2.3 million and \$4.6 million, respectively, of construction-in-progress costs, primarily associated with sites that it no longer planned to build.</p>
Textual	<p><b>Question:</b> What are the geographies that American Express primarily operates in as of 2022?</p> <p><b>Answer:</b> United States, EMEA, APAC, and LACC</p> <p><b>Evidence:</b> AXP_2022_10K (p.154)  Effective for the first quarter of 2022, we changed the way in which we allocate certain overhead expenses by geographic region. As a result, prior period pretax income (loss) from continuing operations by geography has been recast to conform to current period presentation there was no impact at a consolidated level.</p> <p>(Millions) United States EMEA APAC LACC Other Unallocated Consolidated  2022  Total revenues net of interest expense 41,396 4,871 3,835 2,917 (157) 52,862</p>

Table 24: Examples of financial QA pairs categorized by question type. The highlighted segments in the evidence indicate the spans that are most relevant to the question.