

DocFinQA: A Long-Context Financial Reasoning Dataset

Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai
Michael Krumdick, Charles Lovering, Chris Tanner
Kensho Technologies
varshini.bogolu@kensho.com

Abstract

For large language models (LLMs) to be effective in the financial domain – where each decision can have a significant impact – it is necessary to investigate realistic tasks and data. Financial professionals often interact with documents spanning hundreds of pages, but most financial research datasets only deal with short excerpts from these documents. To address this, we introduce a long-document financial QA task. We augment 7,437 questions from the existing FinQA dataset with full-document context, extending the average context length from under 700 words in FinQA to 123k words in DocFinQA. We conduct extensive experiments over retrieval-based QA pipelines and long-context language models. Based on our experiments, DocFinQA proves a significant challenge for even state-of-the-art systems. We also provide a case study on a subset of the longest documents in DocFinQA and find that models particularly struggle with these documents. Addressing these challenges may have a wide-reaching impact across applications where specificity and long-range contexts are critical, like gene sequences and legal document contract analysis. DocFinQA dataset is publicly accessible¹.

1 Introduction

The frequent need to reason over large volumes of textual and tabular data makes financial analysis particularly challenging for LLMs (Azzi et al., 2019). Existing work on automating financial numerical reasoning focuses on unrealistically specific document snippets (Chen et al., 2021; Zhu et al., 2021). Datasets are often limited to pre-selected document sections, failing to reflect the broader and more realistic scenarios faced by analysts (Masson and Montariol, 2020). Financial professionals usually sift through hundreds of pages per document, requiring a deep understanding of

¹<https://huggingface.co/datasets/kensho/DocFinQA>

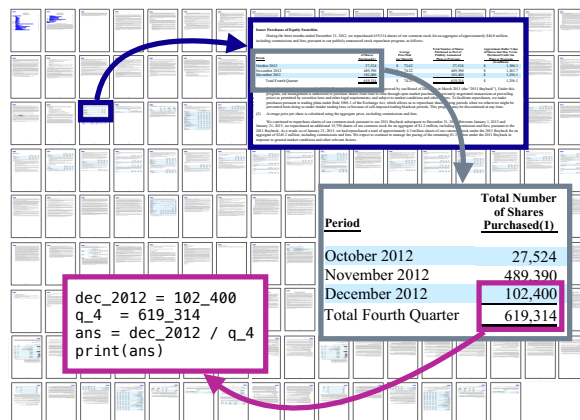


Figure 1: DocFinQA extends FinQA to documents often over 150 pages long (100K+ tokens), so it is difficult to find the pertinent information. The question for the example above is: “For the quarter December 31, 2012 what was the percent of the total number of shares purchased in December?” The correct answer is 16.5%.

both content and structure to navigate and extract pertinent information effectively. Current long-document QA datasets such as NarrativeQA Kočíský et al. (2018) do not test the quantitative reasoning skills needed in the financial domain.

In this work, we introduce DocFinQA, a long-document financial question-answering task. We extend the FinQA dataset of expert annotated questions and answers (Chen et al., 2021) with full Securities and Exchange Commission (SEC) reports. This results in a significantly longer context in the DocFinQA dataset – by a factor of 175 – than the FinQA dataset. Additionally, we manually verified and annotated questions of the test set. The resulting long-document QA task offers a more realistic evaluation of a model’s reasoning capabilities over financial documents. In line with recent work on program synthesis for financial QA (Koncel-Kedziorski et al., 2023), the questions in DocFinQA are appended with Python programs to generate the answers, allowing for training and evaluating program synthesis models for use in

realistic financial workflows.

Using this setup, we evaluate retrieval-based and long-context LLM systems. We study a typical retrieval pipeline that chunks and encodes the document, searching for the best chunks given a question, and passing the question and top- k chunks to a generative QA model (Hsu et al., 2021).

We also evaluate retrieval-free approaches using long-context LLMs (Weston and Sukhbaatar, 2023). Our results show that the successful employment of LLMs in financial settings requires further study of the specific nuances of the financial domain, such as context disambiguation. Our dataset represents a step towards better capturing these nuances.

2 Related Work

Prior studies in financial question answering focus on non-numerical reasoning (Day and Lee, 2016; Jørgensen et al., 2023; Maia et al., 2018). Short-context grounded numerical reasoning tasks were introduced with datasets such as FinQA (Chen et al., 2021) and TAT-QA (Zhu et al., 2021). Recently, understanding long documents has attracted more attention for tasks involving events (Yang et al., 2018), table of contents (Bentabet et al., 2020), and causal relations (Mariko et al., 2022). However, to the best of our knowledge, this is the first attempt to address financial numerical QA grounded in long documents with upwards of hundreds of pages of context for each question.

Long-document QA has been studied in NLP with the introduction of datasets such as SearchQA (Dunn et al., 2017), NarrativeQA (Kočíský et al., 2018), QuALITY (Pang et al., 2022), and PDF-Triage (Saad-Falcon et al., 2023). Due to the limited context size of LLMs, retrieval-based models are commonly used to filter irrelevant text (Izacard et al., 2022; Lewis et al., 2020). Recently, advances in attention mechanisms (Beltagy et al., 2020; Dao et al., 2022) and positional embeddings (Press et al., 2021; Su et al., 2023) allow for end-to-end grounded QA with context windows of more than 100k tokens. However, these methods suffer from loss of important context (Zhang et al., 2023) and often fail to make full use of longer inputs (Liu et al., 2023). Our work studies the intersection of numerical reasoning and long-document processing, and our results demonstrate that there is still ample room for improvement in this domain.

3 DocFinQA Dataset

Dataset Representation: Each question in FinQA is a triplet (c^{golden}, q, a) composed of a golden context c^{golden} , a question q , and an answer a written in human language. An example of FinQA is shown in Table 5 (See Appendix A). We extend the dataset in two ways: (1) context c^{golden} is extended to the full document context D , and (2) we added a Python program p that produces the answer a . Each final sample in DocFinQA is a quartet (D, q, p, a) . An example of DocFinQA is shown in Table 6 (See Appendix A).

Filings Collection: For each question of the FinQA dataset, we identify the corresponding SEC filing from which it was created. We retrieve the filing in HTML/XML format from SEC’s EDGAR service and parse the text and table into clean markdown format (Wang et al., 2023). The collection and parsing processes are presented in more detail in Appendix A and Appendix B, respectively. Figure 2 shows the distribution of document lengths in DocFinQA.




Figure 2: Histogram of document length (#words) in DocFinQA dataset with dash line representing the average length of the documents. The purple line depicts the proportion of documents where the question context is within the current number of words.

Chunking and Alignment: To study retrieval-based QA systems, we split each document D into a set of chunks $C = \{c_1, \dots, c_n\}$. Each chunk consists of 2,750 characters (~ 509 tokens) with a 20% overlap to avoid missing context at the edges. To compute the performance, we identify the best context chunk, c^\star , from the chunk set C associated with each document D that includes the information to answer question

Dataset	#Docs	#QAs	#Words	Multi-page	Numeric	Tabular
NarrativeQA	1,572	46,765	63,000	✓	-	-
QuALITY	381	6,737	5,159	✓	-	-
PDFTriage	82	908	12,000	✓	✓	✓
TAT-QA	2,757	16,552	260	-	✓	✓
FinQA	2,789	8,281	687	-	✓	✓
DocFinQA	801	7,437	123,453	✓	✓	✓

Table 1: Comparison of DocFinQA and existing Finance QA and Long Document QA dataset. DocFinQA includes **multi-page** documents with both **numeric** and **tabular** data.

q . Since FinQA already provides c^{golden} , we compute a pair-wise score (c_i, c^{golden}) , for all chunks, $c_i \in C$, including the golden chunk. We find that four-gram-based similarity score offers the sharpest matching signal among tri-gram, four-gram, and fuzzy matching. The chunk with the highest score is selected as the target context chunk for retrieval. We verify that this process results in good c^\star chunks through manual inspection and by substituting c^\star for c^{golden} in a few-shot QA evaluation with GPT-3.5.

Code Generation: The FinQA dataset provides solutions in a “program” syntax that, when executed, yields the answer (e.g., in Figure 1 the solution is `divide(102400, 619314)`). However, this derivation does not provide meaningful context of what is being calculated. In our running example, 102400 is not semantically grounded to the document. [Koncel-Kedziorski et al. \(2023\)](#) augments FinQA with readable Python code (including named variables like, `dec_shares = 102_400`) that can be executed to derive the answer, providing a layer of interpretability. Thus, we use the code-enhanced version of DocFinQA (See Appendix C).

Statistics: The resultant DocFinQA dataset comprises of 5,735 training, 780 development, and 922 test samples, derived from 801 unique SEC filings. Table 1 shows the statistics and characteristics of DocFinQA in comparison with other finance numerical reasoning and long-document QA datasets.

Impact of Data Selection - DocFinQA vs FinQA: Due to the limited availability of complete SEC filings (refer Appendix A) and imperfections in the code generation process, DocFinQA encompasses 7,437 out of 8,281 of FinQA questions. This process may filter out a collection of question types

that the LLM did not answer due to its limited capability. We investigate the impact of this process by comparing the distribution of the question types in FinQA and DocFinQA. To do this, we show the distribution of questions grouped by their first 2 non-stop words in Figure 6 (Appendix E). The most important observation is that, overall, the distribution of the question set in DocFinQA and FinQA are very similar. No major groups are being filtered out by our data selection process. The dominant questions (above 1% in FinQA) remain dominant and no major impact on the percentages of those questions is observed. The mid-group (above 0.2% in FinQA) question sets see a mixed effect. A large portion of these questions are increased in percentage while some experience significant loss (e.g., “*what percentual*” and “*what decrease*”). Lastly, the long tail group (under 0.2% in FinQA) either remains the same (e.g., “*percent total*” and “*what greatest*”) or is completely wiped out due to a small population (e.g., “*was average*”, and “*what return*”).

4 Retrieval-based QA Evaluation

Retrieval Task: We test three models for context retrieval: ColBERT (**ColB**) ([Khattab and Zaharia, 2020](#)), Sentence-BERT (**SentB**) ([Reimers and Gurevych, 2019](#)), and OpenAI’s **Ada** ([Greene et al., 2022](#)). Further, we finetune the ColBERT model (**FT ColB**) on the training set of DocFinQA to evaluate an in-domain model. More details on the fine-tuning process are given in Appendix G. We also test a matching-based model, BM25 ([Robertson et al., 1995](#)), but observe poor performance (See Appendix F for details).

To retrieve context for a question q over chunk set C , we encode both q and C with the encoding models mentioned above. This results in an embedding, v_q , for the question and chunk embeddings $V_C = \{v_{c_i} | c_i \in C\}$. We compute the

Model	Size	Upper Bound		Original ColBERT			Finetuned ColBERT			Sentence-BERT			OpenAI ADA		
		*	*	Top 1	Top 3	Top 3	Top 1	Top 3	Top 3	Top 1	Top 3	Top 3	Top 1	Top 3	Top 3
		1 shot	3 shot	3 shot	1 shot	3 shot	3 shot	1 shot	3 shot	3 shot	3 shot	1 shot	3 shot	3 shot	1 shot
Falcon	7B	2.0	2.0	<u>1.9</u>	0.0	0.0	<u>1.9</u>	1.3	0.0	<u>1.2</u>	0.1	1.3	<u>2.0</u>	0.1	0.0
MPT	7B	6.8	6.6	<u>4.5</u>	0.8	0.2	<u>4.9</u>	1.0	1.2	<u>3.9</u>	0.6	0.8	<u>4.3</u>	1.6	2.0
MPT	30B	27.1	31.0	<u>15.3</u>	2.2	1.7	<u>16.8</u>	3.2	<u>3.8</u>	1.1	3.8	2.7	<u>15.7</u>	10.4	5.1
Llama 2	7B	17.3	22.0	<u>12.8</u>	5.8	8.0	<u>14.0</u>	6.0	10.3	<u>8.9</u>	2.7	6.5	<u>11.2</u>	4.0	11.0
Llama 2 + SFT	7B	67.1	69.7	30.0	<u>32.6</u>	31.3	32.2	<u>35.3</u>	33.9	19.9	24.1	<u>24.3</u>	28.7	<u>29.4</u>	27.7
Llama 2	13B	30.0	33.4	<u>14.4</u>	10.4	14.1	<u>19.1</u>	11.9	14.5	<u>14.9</u>	7.9	10.2	<u>18.3</u>	9.8	13.7
CodeLlama	7B	26.9	34.0	12.6	11.4	<u>16.1</u>	15.7	12.3	<u>16.8</u>	11.9	8.9	<u>13.2</u>	15.4	14.2	<u>17.5</u>
CodeLlama	13B	32.1	39.0	19.5	14.8	<u>21.5</u>	21.2	15.7	<u>22.5</u>	13.2	8.5	<u>16.0</u>	18.3	14.4	<u>20.9</u>
Mistral	7B	39.7	48.8	<u>23.0</u>	18.8	21.3	<u>25.9</u>	16.8	25.2	<u>19.0</u>	13.6	17.6	20.9	18.8	<u>22.1</u>
GPT 3.5	-	67.3	67.5	36.0	<u>39.0</u>	38.8	38.8	<u>40.7</u>	40.2	24.8	30.1	<u>36.3</u>	35.0	36.5	<u>36.9</u>

Table 2: Performance of the models on DocFinQA in one-shot and few-shot in-context learning settings for the top 1 and top 3 retrieved chunk contexts on the development set. For each model, the best performance among all configurations is in **bold**. For each model, the best performance among different configurations for the same retrieval model is underlined. Top 1 and Top 3 indicate the number of retrieved chunks used as context for a configuration. *The single original context chunk from the FinQA test set is used to estimate the upper bound.

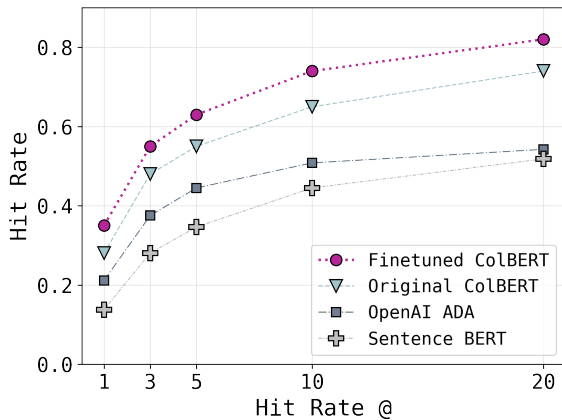


Figure 3: Hit rate of retrieval models.

cosine similarity between v_q and each vector in V_C to retrieve the top- k most similar chunks. We evaluate these models using HR@ k on the test set of DocFinQA using the target c^* . Results are shown in Figure 3. The **FT ColB** yields the highest HR, followed by **ColB**. **FT ColB** yields an average improvement of 91% HR over **SentB** and obtains a 0.35 (HR@1) and 0.55 (HR@3).

4.1 Question Answering Task

We formulate the QA task as a few-shot in-context learning task (Brown et al., 2020). For each in-context example, we only provide the relevant chunk and the answer. For the actual query, we provide k chunks. More details of the few-shot settings are provided in Appendix H.

We evaluate Falcon (Penedo et al., 2023), MPT (MosaicML, 2023), LLaMa 2 and CodeLlama (Touvron et al., 2023), Mistral (Jiang et al., 2023),

GPT3.5 (Brown et al., 2020; OpenAI, 2023) models. We weren't able to evaluate proprietary models such as GPT3 (Brown et al., 2020), BloombergGPT (Wu et al., 2023) due to their inaccessibility. We skipped models that were not finetuned for code generation such as PIXIU (Xie et al., 2023) and FinGPT (Yang et al., 2023) due to their poor performance. We also skipped models trained for other languages such as BBT-Fin (Lu et al., 2023) and XuanYuan 2.0 (Zhang and Yang, 2023).

Performance on the development set: Table 2 reports the full performance of the development set with four retrieval models and three few-shot settings. This results in a total of twelve unique configurations. For both fine-tuned and pre-trained models, we use greedy decoding whenever applicable. One trend noted was that all generic LLMs showed higher accuracy with shorter context and more few-shot examples i.e. top chunk with 3 shots. While code-based LLMs such as Starcoder and CodeLLama showed higher accuracy with longer context i.e. top 3 chunks with 3 shots. This trend is also depicted in Figure 10.

Performance on the test set: Table 3 reports the performance of the same 10 state-of-the-art models on the test set of DocFinQA. The few-shot setting and retrieval model configuration for each LLM are treated as hyperparameters and are picked based on the performance of the development set. We observe that larger models outperform smaller models (e.g., MPT 30B vs MPT 7B). Models trained on code yield higher accuracy than non-

Model/Size	ColB	SentB	ADA	FT ColB
Falcon/7B	2.3	0.3	1.2	1.8
MPT/7B	4.6	2.7	3.8	4.8
MPT/30B	17.3	11.1	12.1	18.1
Llama 2/7B	13.5	8.9	11.1	13.5
Llama 2/13B	18.7	12.7	14.9	19.1
CodeLlama/7B	15.6	12.2	15.2	16.8
CodeLlama/13B	19.1	13.8	18.8	21.0
Mistral/7B	23.2	14.9	21.5	25.0
Llama 2/7B+SFT	32.9	24.8	34.3	36.1
GPT-3.5/-	41.6	33.8	36.4	42.6

Table 3: Performance on DocFinQA test set. For each row, the best performance among all retrieval models is in **bold**. The fewshot setting is selected based on the best performance on the development set (See Table 2).

code models (e.g., CodeLlama vs Llama). Models with additional supervised finetuning (e.g., Llama 2/7B+SFT) and instruction tuning (e.g., GPT-3.5) are among the best examined. Notably, Mistral 7B outperforms several larger models, although it lags behind Llama 2/7B+SFT and GPT-3.5.

The **FT ColB** model is the best retrieval model in all but one setting. It yields a marginal but consistent improvement over the **ColB**, and a large improvement over **SentB** and **Ada**.

5 Case Study w/ 100K+ Token Documents

Recent LLMs can handle context lengths of 128K tokens, but more than 40% of the documents in DocFinQA remain unanswerable even at this content length (see Figure 2). Here, we evaluate performance on a **test subsample** of 200 randomly selected documents, each of which has 100K or more tokens due to the monetary and temporal costs of human evaluation and GPT4.

We explore two retrieval-free options - System 2 Attention (**S2A**) and **Iterative** method. S2A extracts relevant information from each 100K-token chunk of a document before answering the question using the combined extracted information as context (Weston and Sukhbaatar, 2023). The Iterative method produces the output program iteratively as the LLM processes each 100k section of the document. A temporary answer program (initially “None”) is input with each section to the LLM. We also report the performance of the best retrieval-based model (**Retrieval**) based on the experiment in Section 4.

We conducted human evaluations on these 200

Model/Size + Method	w/ Retrieval	Test Subsample
Human	No	41.0
Mistral/7B + Iterative	No	11.5
Mistral/7B + S2A	No	15.5
Mistral/7B + Retrieval	Yes	20.0
GPT-4 + Iterative	No	20.0
GPT-4 + S2A	No	23.0
GPT-4 + Retrieval	Yes	47.5

Table 4: Retrieval-free performance on a case-study of 100K+ token documents.

questions highlighting the challenging nature of this dataset with experienced but non-expert human participants (See Appendix J for details). Non-expert human performance on DocFinQA is lower than human performance reported in FinQA (Chen et al., 2021) (41% versus 50.7%). This can be attributed to the difficulty of finding the golden page, compared to the golden page being given in FinQA. Notably, the expert performance reported in FinQA is 91.2%.

Nonetheless, the non-expert human performance is double that of retrieval-free GPT-4 on these long documents, and roughly triple that of retrieval-free Mistral models. The performance of the iterative method was worse than S2A for both GPT-4 and Mistral with a reduced accuracy of 3% and 4%, respectively. With retrieval, both Mistral and GPT-4 outperform their retrieval-free counterparts, with the assisted GPT-4 now on par with the human cohort. Together, these results highlight that DocFinQA is a difficult test for long-document QA and that there is still room for significant improvement in this domain. For instance, further exploration into methods that combine information across multiple calls to a document-processing LLM is warranted.

6 Conclusion

This paper introduces a realistic document-level question-answering dataset over financial reports. Each question includes a full financial report (averaging 123K words), a far greater challenge than previous work that hones in on pre-specified content. Our findings reveal that this more realistic setting presents a significantly more difficult challenge, thereby opening new avenues for research in quantitative financial question answering.

Acknowledgement

We express our gratitude to our colleagues at Ken-sho Technologies and S&P Global for their invaluable contributions to data annotation, which greatly enhanced the completion of this project. We extend our appreciation to the anonymous reviewers for their supportive input.

Limitation

This work introduced an extension of the existing FinQA dataset. Due to limited human resources, we only validated the test set while the training and the development set were not fully validated. As a result, we can not make any claim of bias and question quality in the not-yet-validated data points offered in this paper. Additionally, as discussed in section 3, the code provided in this work was generated by WizardCoder LLMs. We assume that the code is correct if it produces correct or approximately close to the golden answer. This method may generate both false positive codes (the code that generates the correct answer with incorrect rationales) and false negative codes (the correct code that fails the approximation test).

Broader Impact and Ethical Considerations

We do not foresee any considerable risks associated with our work given that it is an extension of an open-source dataset and uses publicly available documents. To uphold transparency, the paper provides detailed documentation of the dataset creation process, including the sources of data and annotation details. Our dataset serves as a resource to underscore the need for longer context-oriented benchmarks both within and outside the financial domain and does not intend to criticize one or more LLMs.

The annotation in this work is done automatically or in-house, so no crowd-sourced or contract annotators were hired throughout the process. The human evaluation in this study was done by full-time paid coworkers known to the authors.

References

Abderrahim Ait Azzi, Houda Bouamor, and Sira Ferradans. 2019. [The FinSBD-2019 shared task: Sentence boundary detection in PDF noisy text in the financial domain](#). In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing*, pages 74–80, Macao, China.

Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. [LongFormer: The long-document transformer](#). *arXiv preprint arXiv:2004.05150*.

Najah-Imane Bentabet, Rémi Juge, Ismail El Maarouf, Virginie Moulleron, Dialekti Valsamou-Stanislawski, and Mahmoud El-Haj. 2020. [The financial document structure extraction shared task \(FinToc 2020\)](#). In *Proceedings of the 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarisation*, pages 13–22, Barcelona, Spain (Online). COLING.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021. [FinQA: A dataset of numerical reasoning over financial data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3697–3711, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [FlashAttention: Fast and memory-efficient exact attention with io-awareness](#). *Advances in Neural Information Processing Systems*, 35:16344–16359.

Min-Yuh Day and Chia-Chou Lee. 2016. [Deep learning for financial sentiment analysis on finance news providers](#). In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1127–1134. IEEE.

Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. [SearchQA: A new Q&A dataset augmented with context from a search engine](#). *arXiv preprint arXiv:1704.05179*.

Ryan Greene, Ted Sanders, Lilian Weng, and Arvind Neelakantan. 2022. [New and improved embedding model](#).

Chao-Chun Hsu, Eric Lind, Luca Soldaini, and Alessandro Moschitti. 2021. [Answer generation for retrieval-based question answering systems](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4276–4282, Online. Association for Computational Linguistics.

Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. [Atlas: Few-shot learning with retrieval augmented language models](#). *arXiv preprint arXiv:2208.03299*.

- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Rasmus Jørgensen, Oliver Brandt, Mareike Hartmann, Xiang Dai, Christian Igel, and Desmond Elliott. 2023. [MultiFin: A dataset for multilingual financial NLP](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 894–909, Dubrovnik, Croatia. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. [Colbert: Efficient and effective passage search via contextualized late interaction over bert](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Rik Koncel-Kedziorski, Michael Krumdtick, Viet Lai, Varshini Reddy, Charles Lovering, and Chris Tanner. 2023. [Bizbench: A quantitative reasoning benchmark for business and finance](#). *arXiv preprint arXiv:2311.06602*.
- 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. 2023. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*.
- Dakuan Lu, Jiaqing Liang, Yipei Xu, Qianyu He, Yipeng Geng, Mengkun Han, Yingsi Xin, Hengkui Wu, and Yanghua Xiao. 2023. [Bbt-fin: Comprehensive construction of chinese financial domain pre-trained language model, corpus and benchmark](#). *arXiv preprint arXiv:2302.09432*.
- Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. 2018. WWW’18 open challenge: financial opinion mining and question answering. In *Companion proceedings of the the web conference 2018*, pages 1941–1942.
- Dominique Mariko, Hanna Abi-Akl, Kim Trottier, and Mahmoud El-Haj. 2022. [The financial causality extraction shared task \(FinCausal 2022\)](#). In *Proceedings of the 4th Financial Narrative Processing Workshop @LREC2022*, pages 105–107, Marseille, France. European Language Resources Association.
- Corentin Masson and Syrielle Montariol. 2020. [Detecting omissions of risk factors in company annual reports](#). In *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing*, pages 15–21, Kyoto, Japan. -.
- MosaicML. 2023. [MPT-30B: Raising the bar for open-source foundation models](#).
- OpenAI. 2023. [GPT-4 technical report](#).
- Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel Bowman. 2022. [QuALITY: Question answering with long input texts, yes!](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5336–5358, Seattle, United States. Association for Computational Linguistics.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for ccon llm: outperforming curated corpora with web data, and web data only](#). *arXiv preprint arXiv:2306.01116*.
- Ofir Press, Noah A Smith, and Mike Lewis. 2021. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *Proceedings of the 2022 International Conference on Learning Representations (ICLR)*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. [Okapi at TREC-3](#). *Nist Special Publication Sp*, 109:109.
- Jon Saad-Falcon, Joe Barrow, Alexa Siu, Ani Nenkova, Ryan A Rossi, and Franck Dernoncourt. 2023. [Pdf-triage: Question answering over long, structured documents](#). *arXiv preprint arXiv:2309.08872*.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2023. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*, page 127063.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*.

- Jilin Wang, Michael Krumdick, Baojia Tong, Hamima Halim, Maxim Sokolov, Vadym Barda, Delphine Vendryes, and Chris Tanner. 2023. A graphical approach to document layout analysis. In *International Conference on Document Analysis and Recognition*, pages 53–69. Springer.
- Jason Weston and Sainbayar Sukhbaatar. 2023. System 2 attention (is something you might need too). *arXiv preprint arXiv:2311.11829*.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kam-badur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564*.
- Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. 2023. Pixiu: A comprehensive benchmark, instruction dataset and large language model for finance. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Hang Yang, Yubo Chen, Kang Liu, Yang Xiao, and Jun Zhao. 2018. DCFEE: A document-level Chinese financial event extraction system based on automatically labeled training data. In *Proceedings of ACL 2018, System Demonstrations*, pages 50–55, Melbourne, Australia. Association for Computational Linguistics.
- Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023. Fingpt: Open-source financial large language models. *arXiv preprint arXiv:2306.06031*.
- Shiyue Zhang, David Wan, and Mohit Bansal. 2023. Extractive is not faithful: An investigation of broad unfaithfulness problems in extractive summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2153–2174, Toronto, Canada. Association for Computational Linguistics.
- Xuanyu Zhang and Qing Yang. 2023. Xuanyuan 2.0: A large chinese financial chat model with hundreds of billions parameters. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4435–4439.
- 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. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.

A SEC Filing Collection

Each data point in the FinQA dataset consists of a document identification field as shown in Table 5. This field is made up of 3 sections separated by a forward slash. The first is a string called company ticker symbol, the second refers to the year in which this document was filed and the third is the page number in the document where the answer can be found.

Downloading the right 10-K filing from the SEC begins with identifying the company code from the company ticker symbol. For example, C/2017/page_328.pdf-1 in FinQA maps to the CITIGROUP INC with company code 831001. This mapping is obtained from the official file released by SEC which can be found here <https://www.sec.gov/file/company-tickers>. We automatically generate a URL using the company code obtained. From the SEC website, either filings are downloaded as TXT, HTML, or XBRL using the generated URL. At this stage, approximately 6.5% (or 543) data points corresponding to approximately 9.4% (or 17) documents were dropped, either due to lack of mapping or non-availability of older documents. Further, the conversion of the downloaded files to PDF caused a loss of 117 data points (19 unique documents) due to formatting issues.

ID: C/2017/page_328.pdf-1

Context:

Performance graph comparison of five-year cumulative total return the following graph and table compare the cumulative total return on Citi 2019s common stock, which is listed on the NYSE under the ticker symbol 201cc 201d and held by 65691 common stockholders of record as of January 31, 2018, with the cumulative total return of the S&P 500 index and the S&P financial index over the five-year period through December 31, 2017. The graph and table assume that \$ 100 was invested on December 31, 2012 in Citi 2019s common stock, the S&P 500 index and the S&P financial index, and that all dividends were reinvested . comparison of five-year cumulative total return for the years ended date Citi S&P 500 financials.

DATE CITI S&P 500 S&P FINANCIALS
:— :— :— :—
31-Dec-2012 100.0 100.0 100.0
31-Dec-2013 131.8 132.4 135.6
31-Dec-2014 137.0 150.5 156.2
31-Dec-2015 131.4 152.6 153.9
31-Dec-2016 152.3 170.8 188.9
31-Dec-2017 193.5 208.1 230.9

Question:

What was the percentage cumulative total return for the five year period ended 31-dec-2017 of citi common stock?

Answer:

93.5%

Table 5: Example from **FinQA** dataset. The context provided here has been formatted from the original dataset values.

Context:

Table of Contents

UNITED STATES SECURITIES AND EXCHANGE COMMISSION

ANNUAL REPORT PURSUANT TO SECTION 13 OR 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934

For the Fiscal Year Ended December 30, 2006

Commission file number 1-4171

Kellogg Company

(Exact Name of Registrant as Specified in its Charter)

Delaware (State of Incorporation) (I.R.S. Employer Identification No.) One Kellogg Square (Address of Principal Executive Offices) Securities registered pursuant to Section 12(b) of the Securities Act: Title of each class: Name of each exchange on which registered:

...

The Consolidated Financial Statements and related Notes, together with Management's Report on Internal Control over Financial Reporting, and the Report thereon of Pricewaterhouse Coopers LLP dated February 23, 2007, are included herein in Part II, Item 8.

(a) 1. Consolidated Financial Statements Consolidated Statement of Earnings for the years ended December 30, 2006, December 31, 2005 and January 1, 2005. Consolidated Statement of Shareholders' Equity for the years ended December 30, 2006, December 31, 2005 and January 1, 2005. Notes to Consolidated Financial Statements.

(a) 2. Consolidated Financial Statement Schedule All financial statement schedules are omitted because they are not applicable or the required information is shown in the financial statements or the notes thereto.

(a) 3. Exhibits required to be filed by Item 601 of Regulation S-K The information called for by this Item is incorporated herein by reference from the Exhibit Index on pages 61 through 64 of this Report. Pursuant to the requirements of Section 13 or 15(d) of the Securities Exchange Act of 1934, the Registrant has duly caused this Report to be signed on its behalf by the undersigned, thereunto duly authorized, this 23rd day of February, 2007. Pursuant to the requirements of the Securities Exchange Act of 1934, this Report has been signed below by the following persons on behalf of the Registrant and in the capacities and on the dates indicated. Electronic(E), | 10.48 | | IBRF |

| :— | :— | :— |

| | Commission file number 1-4171.* | |

| 21.01 | Domestic and Foreign Subsidiaries of Kellogg. | E |

| 23.01 | Consent of Independent Registered Public Accounting Firm. | E |

| 24.01 | Powers of Attorney authorizing Gary H. Pilnick to execute our Annual Report on Form 10-K for the fiscal year ended December 30, 2006, on behalf of the Board of Directors, and each of them. | E |

| 31.1 | Rule 13a-14(a)/15d-14(a) Certification by A.D. David Mackay. | E |

| 31.2 | Rule 13a-14(a)/15d-14(a) Certification by John A. Bryant. | E |

| 32.1 | Section 1350 Certification by A.D. David Mackay. | E |

| 32.2 | Section 1350 Certification by John A. Bryant. | E |

Question:

What was the average cash flow from 2004 to 2006?

Program:

```
net_cash_2006 = 957.4
net_cash_2005 = 769.1
net_cash_2004 = 950.4
total_net_cash = net_cash_2006 + net_cash_2005 + net_cash_2004
average_net_cash = total_net_cash / 3
answer = average_net_cash
```

Answer:

892.3

Table 6: Examples from **DocFinQA** dataset with text and tables from entire SEC document as context (truncated for legibility), question, associated program and answer. A full report can be founded here https://www.annualreports.com/HostedData/AnnualReportArchive/k/NYSE_K_2006.pdf

B Parsing SEC Filings

Since each filing contains many tables, maintaining the structure and order during extraction is critical for numerical reasoning. We convert each HTML-formatted filing to PDF format and use a finance-specific PDF extractor to parse the filing into markdown format. This process ensures that: (i) our dataset is grounded in the relevant financial documentation and (ii) all the tables in the filings are parsed with high precision into a consistent format without any HTML-tag noise.

We explore different methods for parsing SEC filings consisting of HTML and XML markup into text and markdown tables for use in our QA systems. To evaluate parsing strategies, we measure HR@k (Hit Rate @ k) when searching for the gold chunk among all document chunks for a single document using the FinQA question as the search query. Queries and document chunks are encoded with OpenAI’s ADA model. We compare BeautifulSoup, a standard library for manipulating HTML and XML formatted data, and Kensho Extract, a finance-specific text and table extraction model.² Figure 4 shows the performance of these two methods.

Additionally, we note a better downstream performance of finance-specific models with Kensho Extract retrieved-context compared to that of BeautifulSoup. Qualitative analysis of the different parsers reveals that Kensho Extract is better at structuring the tables used in financial documents, resulting in better readability which seems to extend to the encodings.




Figure 4: Accuracy for varying HR@ for two context extraction methods.

²Passing the raw HTML/XML to the language model produces near-zero performance.

C Code Conversion

Figure 5 shows the steps of converting (a) derivation of the result in FinQA into (b) dummy Python code with dummy variable names, and finally transforming it to (c) a meaningful Python program in DocFinQA following the work by Koncel-Kedziorski et al. (2023).

(a)
`subtract(34.8, 1.2), divide(#0, 34.8)`

(b)
`a = 34.8 - 1.2`
`b = a/34.8`
`c = b * 100`

(c)
`payments_decrease = 34.8 - 1.2`
`change = payments_decrease/34.8`
`answer = change * 100`

Figure 5: Example of code conversion. (a) Original FinQA’s derivation. (b) Dummy Python Program (c) Meaningful Python Code in DocFinQA.

D Model Details

In this work, we used the base models of Falcon, MPT, Llama 2, CodeLlama, and Mistral throughout our work. These models were not trained with supervised finetuning or reinforcement learning human feedback. The GPT-3.5 model employed in this study is `gpt-3.5-turbo-0613` while the GPT-4 model used is `gpt-4-1106-preview`.

We also included the Llama 2/7B + SFT that was finetuned on the training set of DocFinQA with golden chunk from FinQA (*c^{golden}*). The finetuning process takes 3 epochs with a batch size of 32. We use the context provided by the FinQA dataset as the input due to the limited maximum token length of the model. The maximum token length is set to 2048. The model is finetuned on 8 x Nvidia A100-80GB GPUs. We use AdamW optimizer with learning rate of $2e-6$. The training process takes 4 hours to complete.

E Distribution of question by question types

Figure 6 shows the distribution by question types of the dataset before (FinQA, green) and after (DocFinQA, purple) the automatic data selection.




Figure 6: Distribution of questions grouped by question types in the original FinQA and DocFinQA. The x-axis (percentage) is presented in log scale to magnify the differences between the two sets.

F Performance of retrieval methods

Figure 7 shows a pilot study comparing dense retrieval with OpenAI ADA and Sentence BERT versus sparse retrieval (BM 25) on the development set. We can see that the dense retrieval model offers a much higher hit ratio.




Figure 7: Accuracy for varying HR@ for three search methods on the development set

Figure 8 shows the prompt template with in-context learning that we used.

```

Context: {golden chunk}
Question: {question}
Python Program: {program}
Answer: {answer}

Context: {golden chunk}
Question: {question}
Python Program: {program}
Answer: {answer}

Context: {golden chunk}
Question: {question}
Python Program: {program}
Answer: {answer}

Context: {first chunk}
{second chunk}
{third chunk}
Question: {question}
Python Program:
    
```

Figure 8: Prompt template with Top-3 context and 3-shot In-Context Learning.

G ColBERT Finetuning

We finetune the original ColBERT v1 model on the train set of DocFinQA. For each data point, we perform chunking and alignment to generate one golden chunk and $n - 1$ negative chunks. For training, we generate a list of tuples (qid, pid+, pid-), where qid refers to the question, pid+ refers to the golden chunk and pid- refers to each of the negative chunks in that document. We train the model for a total of 3 epochs and store the checkpoints at the end of each epoch. The hit rate of the Finetuned ColBERT model after each epoch on the development set is shown in Figure 9. We observe that after the first epoch, additional finetuning does not show any performance improvement. The Finetuned ColBERT model referred to in this study thus uses the weights after the first epoch of training.




Figure 9: Hit rate of different ColBERT variants on the development set of DocFinQA.

H Few-shot Settings

Due to the limited context length of the LLMs, the number of few-shot demonstrations and the number of chunks fed into the In-Context Learning must be optimized. We explore 3 settings of the number of few-shot examples and 4 settings of the number of chunks used as context in the query. Figure 10 shows the performance of these settings in retrieval and answered by Llama 13B and CodeLlama 13B on the development set. We see that a higher number of few-shot examples (numshot=3) yield consistently better performance compared to a lower one (numshot=1).




Figure 10: A QA performance plot on the development set of DocFinQA for the Llama 2 13B and CodeLlama 2 13B models for each of the 12 configurations

I Golden Chunk Position

Figure 11 shows the distribution of the position golden chunk with the documents. We see that most of the golden chunks appear within the first 250 chunks (approximately 125K tokens which can be fed into the newest generative models). Nonetheless, there are a substantial number of questions that the golden chunk appears beyond this threshold.

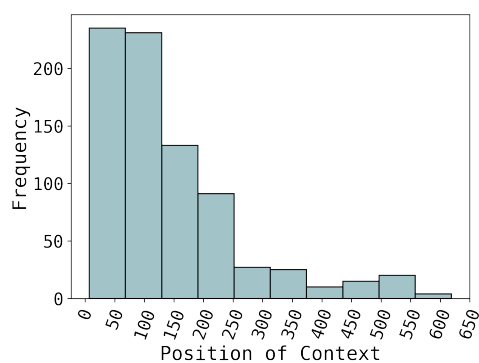


Figure 11: Histogram of the position of the FinQA context in the original SEC filing that is split into chunks of size 2750.

J Human Evaluation Setting

We recruited three data professionals with 4-5 years of experience working with financial documents, including but not limited to 10-K filings, to estimate human evaluation. The professionals were provided with the entire document in PDF format, maintaining the SEC’s original format for ease of reading. They were allowed to use the keyword-search feature of PDF reader applications and a simple calculator for basic arithmetic operations required for this task. On average, the professionals spent 25 minutes per question.