
Benchmarking Table Comprehension in the Wild

Yikang Pan*
Boson AI
University of Toronto
yikang@boson.ai

Yi Zhu
Boson AI
yi@boson.ai

Rand Xie
Boson AI
rand@boson.ai

Yizhi Liu
Boson AI
yizhiliu@boson.ai

Abstract

Large Language Models (LLMs), while being increasingly dominant on a myriad of knowledge-intensive activities, have only had limited success understanding lengthy table-text mixtures, such as academic papers [1] and financial reports [2]. Recent advances of long-context LLMs have opened up new possibilities for this field. Nonetheless, we identify two roadblocks: (1) Prior benchmarks of table question answering (TableQA) have focused on isolated tables without context [3], making it hard to evaluate models in real-world scenarios. (2) Prior benchmarks have focused on some narrow skill sets of table comprehension such as table recognition [4], data manipulation/calculation [5], table summarization [6] etc., while a skilled human employs those skills collectively. In this work, we introduce TableQuest, a new benchmark designed to evaluate the holistic table comprehension capabilities of LLMs in the natural table-rich context of financial reports. We employ a rigorous data processing and filtering procedure to ensure that the question-answer pairs are logical, reasonable, and diverse. We experiment with 7 state-of-the-art models, and find that despite reasonable accuracy in locating facts, they often falter when required to execute more sophisticated reasoning or multi-step calculations. We conclude with a qualitative study of the failure modes and discuss the challenges of constructing a challenging benchmark. We make the evaluation data, judging procedure and results of this study publicly available to facilitate research in this field.

1 Introduction

In many fields, ranging from finance to healthcare, tables play a crucial role in organizing and interpreting vast amounts of data. They allow individuals to discern patterns, extract meaningful insights, and make informed decisions by presenting information in a concise, structured format. However, the tabular format poses visible challenges for large language models (LLMs), which are fundamentally designed to process data in sequential, text-based formats rather than navigating the complex spatial relationships inherent in tables. Despite growing interest and progress in natural language processing, LLMs still struggle with extracting, understanding and reasoning about information presented in tabular form, thereby limiting their use in industrial applications.

While numerous approaches have been proposed to enhance LLMs' performance on tasks involving tabular data, there is a significant gap in comprehensive evaluation frameworks that capture the breadth and complexity of real-world use cases. Existing benchmarks tend to be narrowly focused,

*Work completed during an internship at Boson AI.

Note 7 - Solar Energy Systems, Net
Our solar energy systems, net, consisted of the following (in millions):

	December 31, 2022	December 31, 2021
Solar energy systems in service	\$ 6,795	\$ 6,795
Initial direct costs related to customer solar energy system lease acquisition costs	154	206
Less: accumulated depreciation and amortization (1)	(6,469)	(6,888)
	5,216	(4,488)
Solar energy systems under construction	1	2
Solar energy systems pending interconnection	12	1
Solar energy systems, net (2)	\$ 5,229	\$ 5,409

(1) Depreciation and amortization expense during the years ended December 31, 2022, 2021 and 2020 was \$ 220 million, \$220 million and \$ 220 million, respectively.

(2) As of December 31, 2022 and 2021, there were \$ 740 million and \$402 million, respectively, of gross solar energy systems under lease pass-through lease arrangements with accumulated depreciation of \$221 million and \$143 million, respectively.

Note 8 - Property, Plant and Equipment, Net
Our property, plant and equipment, net, consisted of the following (in millions):

	December 31, 2022	December 31, 2021
Machinery, equipment, vehicles and office furniture	\$ 22,312	\$ 23,228
Tooling	3,147	2,579
Leasehold improvements	3,166	2,268
Land and buildings	9,505	7,751
Computer equipment, hardware and software	3,799	2,073
Construction in progress	5,781	4,262
Less: Accumulated depreciation	(41,767)	(32,989)
	(12,297)	(10,442)
Total	\$ 29,725	\$ 23,548

Construction in progress is primarily comprised of ongoing construction and expansion of our facilities, and equipment and tooling related to the manufacturing of our products. Completed assets are transferred to their respective asset classes and depreciation begins when an asset is ready for its intended use.

Depreciation expense during the years ended December 31, 2022, 2021 and 2020 was \$ 3.31 billion, \$2.42 billion and \$1.91 billion, respectively.

Phonics has entered into an agreement with us on Capillary Heats with investments in the production equipment that covers its manufacture and supply of heat exchangers. Under our arrangements with Phonics, we plan to purchase the full value from their production equipment at negotiated prices. As the result of the arrangement with Phonics, we will be able to purchase the full value from their production equipment at negotiated prices. We account for the arrangement as a lease under ASC 842. We account for the production equipment as intangible assets when production commences. We account for each lease and any non-lease components associated with that lease as a single lease component for all asset classes, except production equipment classes embedded in supply agreements. This results in us recording the cost of these production equipment when property, plant and equipment, net, on the consolidated balance sheets with a corresponding liability recorded to debt and finance leases. Depreciation on Phonics production equipment is computed using the useful production method whereby capitalized costs are amortized over the total expected productive life of the respective assets. As of December 31, 2022 and 2021, we had cumulatively capitalized gross costs of \$2.02 billion and \$1.55 billion, respectively, on the consolidated balance sheets in relation to the production equipment under our Phonics arrangement.



Extract



Calculate



Analyze

Question	Answer	Reference CoT steps
How much are the total commitments expiring within 1 year as listed in the table?	\$721 million.	-
Calculate the difference in the adjusted operating margin between fiscal 2022 and fiscal 2023.	0.2%	... Calculate difference: 15.4% - 15.2% = 0.2%
Assess the impact of capital expenditures on the trend of free cash flow over the period from 2021 to 2023.	Increased capital expenditures from 2021 to 2023 do not strictly correlate with ...	1. Define relationship between free cash flow and capital expenditures ...

Figure 1: A sample of our TableQuest benchmark showcasing the different aspects of table comprehension. The sample image on the left was for demonstration, which was rendered from the source HTML file.

often concentrating on specific aspects like table-based question answering [7, 3, 8, 9, 10, 11] or information retrieval [12, 13, 1], and fail to represent the full spectrum of challenges that professionals encounter when working with tabular data.

To address this gap, we introduce *TableQuest*, a novel benchmark specifically designed to evaluate the proficiency of LLMs in understanding and reasoning with tabular data across a diverse set of scenarios. TableQuest is a question-answering dataset over public traded companies' annual filings which contain ample text and tables. We carefully construct the questions such that on answering the questions correctly, the models demonstrate the following capabilities:

Extraction (Easy level). Given the popularity of long-context retrieval benchmarks (e.g., needle-in-a-haystack [14]), we adapt this format by focusing on retrieving information from table cells within real-world documents, rather than from randomly injected facts. It not only enables evaluating models to distinguish similar pieces of information under an ecologically valid condition, but also requires models to extract insights of two dimensions (rows and columns) from linear inputs.

Calculations (Medium level). Finance reports, the essential tools to make informed financial decisions and strategize for future growth, provide the perfect ground for a wide range of numerical tasks, including data analysis, forecasting, budgeting, and performance tracking. Prior works have explored this area in a wide range of heavily-engineered systems [9, 15], breaking down the problem into multiple steps for multiple models specializing in each. In our benchmark, we present this challenge in a monolithic fashion, pressuring-testing the models' ability to spot evidence in tables as well as conduct accurate calculations.

Analytics (Hard level). One step further from accurate calculations, we ask models to uncover insights from numbers they extracted or calculated and output in words. The aim is that models could digest the reports and output summaries in a query-focused manner, effectively aiding humans at document comprehension.

In our evaluation of TableQuest, we explored the performance of a wide array of both proprietary and open-source LLMs, employing an automated evaluation system to assess the quality and correctness of the models' responses. Our initial results indicate that while some models can handle straightforward data extraction tasks with reasonable accuracy, they often falter when required to execute more sophisticated reasoning or multi-step calculations. Common issues include difficulty maintaining context over longer sequences, following instructions (particularly format), inconsistencies in applying domain-specific knowledge and reasoning, highlighting areas where current models fall short and where future improvements are necessary.

2 TableQuest

2.1 Data Preparation

Source. We collect the 10-K reports of S&P 500 companies directly from the SEC’s EDGAR database², which offers comprehensive HTML annotations for each table, capturing hierarchical structures and detailed cell information, including text content and formatting. This structured format allows for the extraction and post-processing of data using HTML tags, enabling a thorough representation of the complex relationships and structures present within the tables. Despite a lack of consensus on the best format for table serialization [16, 17], we employed HTML due to two reasons: 1. it can preserve table structures better than other languages (e.g. natural text, markdown, json etc.), and 2. the internet being a critical source for web-scale LLM pre-training corpora [18, 19].

Question-Answer Synthesis. Existing question-answering datasets often exhibit internal homogeneity, such as focusing solely on tabular answers derived from multiple tables [20]. To address this limitation, we implemented a progressive three-round chat format, where each round builds upon the previous one, resulting in a sequence of question-answer pairs. In each round, we introduce a new question that demands an additional capability, expanding the model’s range of responses. This iterative process ensures that each prior question-answer pair informs and enhances the subsequent one, fostering greater diversity in the questions generated. Additionally, we incorporated explicit chain-of-thought (CoT) prompting to guide the model in articulating its reasoning steps, which facilitated more thorough quality inspection. We present an validation study on the quality of the question-answer generation pipeline with a self-supervised data source, which we detail in Appendix F. The multi-round chat prompts can be found in Appendix A.

Quality filtering. Given the difficulty of questions, we employ a hybrid human-machine verification pipeline to validate the questions and answers generated were relevant, reasonable and valid. We collaborated closely with a team of finance experts to develop the initial guidelines for constructing the questions. In Appendix B, you can find the data processing and quality assurance specifics.

Metric	Easy	Medium	Hard
Avg. Question Length	90.19	135.04	162.79
Avg. Answer Length	27.70	15.61	202.61
Table Statistics			
Avg. Tables per file	1.82	1.90	1.86
Avg. Rows per Table	11.31	10.87	10.49
Avg. Columns per Table	13.94	13.32	13.44

Table 1: Summary Statistics for TableQuest Dataset (Eval) by Difficulty. The context for these questions are sampled randomly, such that the complexity of the tables should be on the same level.

3 Experiments

Inputs and Outputs. We sample 240 questions from our synthesized dataset, 80 of each difficulty level, as the evaluation subset of our benchmark. Each question comes with the context on which it was generated, which was included into the prompt for generation, along with instructions and a one-shot prompt for answer formatting purposes, the details of which can be found in Appendix A. The metadata of the sample set can be found in Table 1.

Configurations. We experiment with 7 models, 4 proprietary and 3 open-source, as detailed in table 2. We employ the following configuration for the models included in our test:

We use a vLLM[21] backend for serving the open-sourced models (including the Qwen, and the Llama families). We use a temperature of 0.05, with a token limit of 16,384 or 8,192 as determined by the underlying LLM³ For proprietary models, we use their official API services with a temperature of 0.05, with their maximum input token limit provided respectively. We truncate the input exceeding the input limits, which accounts for only <2% of all samples.

²<https://www.sec.gov/search-filings>

³in our tested models, only Meta-Llama-3-70B-Instruct uses an 8,192 limit.

Evaluation Framework. In traditional question-answering (QA) tasks, a range of rule-based and word-matching evaluation metrics have been commonly used to assess model performance. Some of the most prevalent metrics include Exact Match (EM), F1 Score, BLEU (Bilingual Evaluation Understudy) [22], ROUGE (Recall-Oriented Understudy for Gisting Evaluation) [23], METEOR (Metric for Evaluation of Translation with Explicit ORdering) [24], and BERTScore [25].

Despite their widespread usage, these metrics exhibit significant limitations, particularly when applied to analytical or open-ended questions. Therefore, in addition to word-matching based metrics, we employ a machine judge for a more interpretable and flexible judging of long-form answers. Specifically, for analytics tasks, we compare each model’s response to a baseline (in our case, gpt-4o using text input, with a 1000 base score), ask the judging model (gpt-4o-turbo) to reason about which model is more accurate, logical, and clear in its output. The prompts we used are detailed in Appendix A. For extraction and calculation tasks, we also provide accuracy as a metric in Appendix D.

4 Results

Table 2 demonstrate model performance using ELO scores, with shades of blue and red demonstrating better and worse scoring.

Closed-source models like gpt-4-turbo and claude-3-5-sonnet clearly outperform open-source models, particularly on harder tasks, as evidenced by their significantly higher ELO ratings. GPT-4-turbo leads with an overall ELO rating of 1164.35, excelling in medium difficulty tasks while maintaining strong performance in easier tasks. Claude-3-5-sonnet also performs well across all difficulty levels, showing a more balanced profile with a particularly strong performance in easy tasks. We observe that among open-source model, However, Meta-Llama-3.1-70B-Instruct, an open-source model, shows promise by delivering competitive performance in medium tasks (ELO 1008.80) and a decent overall ELO rating (869.27), standing out as a notable open-source option despite trailing behind the closed-source leaders. We present a more in-depth analysis on what causes the performance gap in Appendix C with the statistics in Appendix D.

Models can excel at easier tasks, but be significantly weaker at hard tasks. Notably, models like gemini-1.5-pro and gemini-1.5-flash excel in simpler tasks, with Easy task ELO ratings reaching up to 967.86 and 880.19 respectively, but struggle with harder tasks, as indicated by their lower Hard ELO ratings (both at 779.98). Similarly, Meta-Llama-3.1-70B-Instruct demonstrates moderate performance in Easy and Medium tasks (ELO 880.19 and 1008.80 respectively) but falls significantly short in Hard tasks, with an ELO rating of 703.98. These patterns highlight the specialization of some models for simple tasks while exposing weaknesses in their ability to handle more complex scenarios. As a side note, we found certain multi-modal models exhibit this pattern as well, which we demonstrate in the complete evaluation results in Appendix D, but we call for future work for further investigation in multi-modal table comprehension benchmarking.

Model	Overall ELO	Easy ELO	Medium ELO	Hard ELO
gpt-4-turbo	1164.35	1104.56	1372.43	1045.06
claude-3-5-sonnet-20240620	1081.39	1250.12	1219.87	832.23
gemini-1.5-pro	896.81	967.86	930.28	779.98
gemini-1.5-flash	880.88	880.19	969.48	779.98
Meta-Llama-3.1-70B-Instruct	869.27	880.19	1008.80	703.98
Qwen2-72B-Instruct	867.81	884.54	912.63	793.03
Meta-Llama-3-70B-Instruct	772.05	853.85	833.20	606.52

Table 2: ELO ratings across difficulties. Baseline is GPT4o, and we start with a base score of 1000.

5 Conclusion

In this paper, we introduce TableQuest, a new benchmark, that addresses the gaps by evaluating LLMs’ ability to reason with tables in natural, table-rich context of financial reports. Our results with 7 state-of-the-art models reveal that, while LLMs can locate facts with reasonable accuracy, they struggle with complex reasoning and multi-step calculations. This highlights the need for future improvements in LLMs to enhance their proficiency in tasks requiring comprehensive table comprehension.

References

- [1] Minzheng Wang, Longze Chen, Cheng Fu, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang, and Yongbin Li. Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA, June 2024. arXiv:2406.17419 [cs].
- [2] Pranab Islam, Anand Kannappan, Douwe Kiela, Rebecca Qian, Nino Scherrer, and Bertie Vidgen. FinanceBench: A New Benchmark for Financial Question Answering, November 2023. arXiv:2311.11944 [cs, stat].
- [3] 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. 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, 2021. Association for Computational Linguistics.
- [4] Xinyi Zheng, Doug Burdick, Lucian Popa, Xu Zhong, and Nancy Xin Ru Wang. Global Table Extractor (GTE): A Framework for Joint Table Identification and Cell Structure Recognition Using Visual Context, December 2020. arXiv:2005.00589 [cs].
- [5] Xiaokang Zhang, Jing Zhang, Zeyao Ma, Yang Li, Bohan Zhang, Guanlin Li, Zijun Yao, Kangli Xu, Jinchang Zhou, Daniel Zhang-Li, Jifan Yu, Shu Zhao, Juanzi Li, and Jie Tang. TableLLM: Enabling Tabular Data Manipulation by LLMs in Real Office Usage Scenarios, April 2024. arXiv:2403.19318 [cs].
- [6] Ankur P. Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. ToTTo: A Controlled Table-To-Text Generation Dataset, October 2020. arXiv:2004.14373 [cs].
- [7] Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Hailey Schoelkopf, Riley Kong, Xiangru Tang, Mutethia Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, Dragomir Radev, and Dragomir Radev. FeTaQA: Free-form Table Question Answering. *Transactions of the Association for Computational Linguistics*, 10:35–49, January 2022.
- [8] Naihao Deng, Zhenjie Sun, Ruiqi He, Aman Sikka, Yulong Chen, Lin Ma, Yue Zhang, and Rada Mihalcea. Tables as Texts or Images: Evaluating the Table Reasoning Ability of LLMs and MLLMs. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Findings of the Association for Computational Linguistics ACL 2024*, pages 407–426, Bangkok, Thailand and virtual meeting, August 2024. Association for Computational Linguistics.
- [9] Yilun Zhao, Hongjun Liu, Yitao Long, Rui Zhang, Chen Zhao, and Arman Cohan. FinanceMath: Knowledge-Intensive Math Reasoning in Finance Domains, November 2023.
- [10] Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 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, 2021. Association for Computational Linguistics.
- [11] Sunjun Kweon, Yeonsu Kwon, Seonhee Cho, Yohan Jo, and Edward Choi. Open-WikiTable: Dataset for Open Domain Question Answering with Complex Reasoning over Table, May 2023. arXiv:2305.07288 [cs].
- [12] Bogdan Kostić, Julian Risch, and Timo Möller. Multi-modal Retrieval of Tables and Texts Using Tri-encoder Models, October 2021. arXiv:2108.04049 [cs].
- [13] Junjie Huang, Wanjun Zhong, Qian Liu, Ming Gong, Daxin Jiang, and Nan Duan. Mixed-modality Representation Learning and Pre-training for Joint Table-and-Text Retrieval in OpenQA, October 2022. arXiv:2210.05197 [cs].

- [14] Kamradt G. Needle in a haystack. https://github.com/gkamradt/LLMTest_NeedleInAHaystack, 2024.
- [15] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaying Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. InternLM2 Technical Report, March 2024. arXiv:2403.17297 [cs].
- [16] Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, Yanjun Qi, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, and Christos Faloutsos. Large Language Models (LLMs) on Tabular Data: Prediction, Generation, and Understanding – A Survey, June 2024. arXiv:2402.17944 [cs].
- [17] Weizheng Lu, Jiaming Zhang, Jing Zhang, and Yueguo Chen. Large Language Model for Table Processing: A Survey, February 2024. arXiv:2402.05121 [cs].
- [18] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer, September 2023. arXiv:1910.10683 [cs, stat].
- [19] Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The FineWeb Datasets: Decanting the Web for the Finest Text Data at Scale, June 2024. arXiv:2406.17557 [cs].
- [20] Vaishali Pal, Evangelos Kanoulas, Andrew Yates, and Maarten de Rijke. MultiTabQA: Generating Tabular Answers for Multi-Table Question Answering.
- [21] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- [22] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a Method for Automatic Evaluation of Machine Translation. In Pierre Isabelle, Eugene Charniak, and Dekang Lin, editors, *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics.
- [23] Chin-Yew Lin. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
- [24] Satyanjeev Banerjee and Alon Lavie. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss, editors, *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics.
- [25] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. BERTScore: Evaluating Text Generation with BERT, February 2020. arXiv:1904.09675 [cs].
- [26] Minzheng_Wang. MozerWang/Loong, August 2024. original-date: 2024-06-16T12:21:44Z.

A Prompts

Below is the inference prompt which we use on the 7 models tested.

```
[System]

You are a helpful question-answering chatbot.

[User]

You are asked to answer questions based on provided material containing tables and text in HTML. You should layout your reasoning steps (if necessary) to obtain the correct answer (marked by "steps"), then provide a succinct answer (marked by "answer") in JSON. If units can be known, the answer should include units such as $, %, million and etc.

<one-shot example>

## Context:

<context>

## Question:

<question>
```

Below is the prompt we use for our judging model, gpt-4-turbo.

```
[User]

You will be presented with a question, the context, a reference answer marked with "gt", and two model predictions marked by "pred_1" and "pred_2" in JSON format. Compare the two predictions and evaluate which one better answers the question based on the reference answer. Output your rationale (marked by "rationale" in JSON) for your decision, and then your decision (marked by "better_model" in JSON, with possible values "model_1", "model_2", or "tie") in JSON format. Please use the following listed aspects and their descriptions as evaluation criteria:



- Accuracy: The assistant's answer is semantically consistent with the gold answer, preferably with the reference answer verbatim;
- Clarity: The assistant's answer is clear and easy to understand;
- Logical: The assistant's reasoning steps are clear and reasonable, does not make gigantic leaps between steps and makes assumptions without clear evidence;



## Context in JSON:
{ question: <question>
  gt: <answer>
  pred_1: <LLM1's response>
  pred_2: <LLM2's response>
  ... }
## Output:
```

Below is the synthesis prompts we used for our multi-round data generation. Each generated question-answer pair from a lower difficulty are passed into the prompt serving as a negative example, encouraging the question-answer pair to be accurate to the intended difficulty level.

```
Prompt 1

You are an underwriter examining a company's performance to make a critical decision which has a critical aspect to your company and your job. Generate a question that you have for the snippet annual report from {company} {year}. Then, present your answer.
```

Prompt 2

You are designing a test question for middle schoolers' ability to interpret information from the table. Generate a question that asks for the information from a cell. Then provide the correct answer.

...

Output:

Prompt 3

You are now designing a test question for high schoolers' ability to calculate based on given instructions. Generate a question which involves numbers not directly mentioned in the table. Then, provide the steps to calculate. Finally, present the correct answer (a number or a few words, NOT a sentence).

...

Output:

Prompt 4

You're an experienced university professor designing test questions for corporate finance. Based on the material provided, generate an analytical question that focuses on a fraction of the content. Then, provide a detailed step-by-step reference reasoning steps. Finally, provide a concise reference answer which should be less than 3 sentences long for the ease of grading.

...

Output:

B Quality assurance on the data synthesis pipeline

Our process involved generating an initial set of questions, followed by iterative feedback sessions with these experts to refine and enhance the quality of the questions. To ensure a thorough evaluation, we selected a representative sample of 50 questions from each difficulty level for human annotation. For this task, we engaged four STEM graduates who were responsible for locating the sources of answers, verifying the alignment of human-provided answers with the synthetic ones, and assessing whether the quality of the questions appropriately matched their designated difficulty levels. Throughout this iterative process, we maintained a high standard of accuracy. We identified the most common error cases, which were then used as in-context examples to further fine-tune the performance of GPT-4-Turbo, ensuring it could accurately re-evaluate the question-answer pairs. This approach allowed us to systematically address and rectify discrepancies, enhancing the overall robustness and validity of our benchmark. Finally, the refined set of question-answer pairs demonstrated high fidelity to real-world scenarios, reflecting the domain knowledge and quality standards initially set out.

C Case Study

Given the gap in open-source/proprietary model performance, we conduct a qualitative analysis with one model from each world as representatives, llama-3.1-70B-Instruct and GPT-4-Turbo. We employ the most recent reasoning models to date, o1-preview, to judge and critique both responses on the same question for comparison. On a 3-point scale scoring system, GPT-4-Turbo leads by 0.34 on average. We identified 12.5% of the evaluation samples on which GPT-4-Turbo scored 2 or higher than Llama-3.1-70B-Instruct.

The performance disparity between Llama and GPT-4-Turbo in the financial analysis benchmark reveals limitations in Llama's capabilities. Firstly, the prevalence of empty responses, particularly in questions requiring straightforward fact retrieval, suggests deficiencies in Llama's long-context information retrieval and instruction following output generation mechanisms. Secondly, the cases of

calculation errors and misinterpretation of numerical data, indicates weaknesses in its mathematical reasoning abilities and financial data comprehension. These errors ranged from failing to calculate specific percentages to misinterpreting financial trends, highlighting a need for enhanced numerical processing capabilities. Lastly, Llama demonstrated difficulties with complex questions requiring multi-step reasoning. In these instances, Llama often provided incomplete analyses or failed to address key aspects of multi-faceted queries, suggesting limitations in its ability to decompose and address complex problems. These findings collectively point to areas of improvement for all open sourced models, particularly in domains requiring precise numerical analysis and multi-step reasoning. Below we showcase five representative examples of the differences in the judgement from the judge to the two models.

C.1 Case 1: Empty Response vs Detailed Answer (ID: 93)

Llama Response

The model prediction is completely empty, providing no answer to the question asked.

GPT-4-Turbo Response

The prediction accurately identifies the number of units in the property 'The Grand' as 243, which matches the reference answer verbatim. The prediction also provides a clear and logical step-by-step explanation of how the answer was derived, including locating the property in a table and identifying the relevant column for unit numbers.

C.2 Case 2: Calculation Error vs Accurate Computation (ID: 186)

Llama Response

The reference answer states a specific increase of approximately 1115.38% in non-subscription revenue due to the merger. However, the model prediction does not mention this percentage increase and instead concludes that the exact impact cannot be quantified, which contradicts the reference answer.

GPT-4-Turbo Response

The prediction provides a detailed and logical step-by-step analysis of the non-subscription revenue growth from 2021 to 2023, including calculations of the absolute and percentage increase. It also correctly identifies the impact of the IHS Markit merger as a substantial positive influence on this growth, which aligns with the reference answer.

C.3 Case 3: Misinterpretation of Question (ID: 991)

Below we have the question, ground truth answer, and the context document in Figure 2. The model outputs and the judgements are laid out in the following.

Question and Ground Truth

Question: Analyze the effectiveness of the cash flow hedge strategy for the cross-currency swaps by comparing the fair values of the contracts between 2022 and 2023.

Ground Truth Answer: The fair value of the cross-currency swaps decreased by \$13.2 million from 2022 to 2023, indicating reduced effectiveness or negative economic impact on the hedge during this period.

NOTE 5. FINANCIAL INSTRUMENTS

As of December 30, 2023, the aggregate U.S. dollar equivalent notional value of our outstanding commodity contracts and foreign exchange contracts was \$5.8 million and \$1.34 billion, respectively. Our outstanding foreign exchange contracts as of December 30, 2023 were recorded in various currencies, primarily the U.S. dollar, euro, Chinese renminbi, British pound sterling and Hong Kong dollar.

We recognize derivative instruments as either assets or liabilities at fair value in the Consolidated Balance Sheets. We designate commodity forward contracts on forecasted purchases of commodities and foreign exchange contracts on forecasted transactions as cash flow hedges. We also enter into foreign exchange contracts to offset certain of our economic exposures arising from foreign exchange rate fluctuations.

Cash Flow Hedges

For derivative instruments that are designated and qualify as cash flow hedges, the effective portion of the gain or loss on the derivative is reported as a component of "Accumulated other comprehensive loss" and reclassified into earnings in the same period(s) during which the hedged transaction impacts earnings. Gains and losses on these derivatives, representing either hedge ineffectiveness or hedge components excluded from the assessment of effectiveness, are recognized in current earnings. Except for the cross-currency swap discussed below, cash flow hedges were not material in 2023, 2022 or 2021.

Cross-Currency Swap

In March 2020, we entered into U.S. dollar to euro cross-currency swap contracts with a total notional amount of \$250 million to effectively convert our fixed-rate U.S. dollar-denominated debt into euro-denominated debt, including semiannual interest payments and the payment of principal at maturity. During the term of the contract, which ends on April 30, 2030, we pay fixed-rate interest in euros and receive fixed-rate interest in U.S. dollars. These contracts have been designated as cash flow hedges. The fair value of these contracts was \$2.3 million as of December 30, 2023 and \$15.5 million as of December 31, 2022, which were included in "Other Assets" in the Consolidated Balance Sheets. Refer to Note 9, "Fair Value Measurements," to the Consolidated Financial Statements for more information.

We recorded no ineffectiveness from our cross-currency swap to earnings during 2023, 2022 or 2021.

Other Derivatives

The following table shows the fair value and balance sheet locations of other derivatives as of December 30, 2023 and December 31, 2022:

(In millions)	Asset				Liability			
	Balance Sheet Location	2023	2022	Balance Sheet Location	2023	2022		
Foreign exchange contracts	Other current assets	\$ 6.3	\$ 4.3	Other current liabilities	\$ 6.0	\$ 9.6		
Commodity contracts	Other current assets	—	—	Other current liabilities	—	.2		
		\$ 6.3	\$ 4.3		\$ 6.0	\$ 9.8		

For other derivative instruments not designated as hedging instruments, the gain or loss is recognized in current earnings.

The following table shows the components of the net gains (losses) recognized in income related to these derivative instruments:

(In millions)	Statements of Income Location	2023	2022
Foreign exchange contracts	Cost of products sold	\$ 3.4	\$ 5.6
Foreign exchange contracts	Marketing, general and administrative expense	5.5	(4.3)
		\$ 8.9	\$ 1.3

Figure 2: Excerpt from the company's financial report

LLaMA Prediction

```
{ "steps": "To analyze the effectiveness of the cash flow hedge strategy for the cross-currency swaps, we need to calculate the change in fair value of the contracts between 2022 and 2023. The change in fair value is calculated as the difference between the fair value in 2023 and the fair value in 2022. Then, we calculate the percentage change in fair value by dividing the change in fair value by the original fair value in 2022 and multiplying by 100.", "answer": { "change_in_fair_value": "-$13.2 million", "percentage_change_in_fair_value": "-85.16%" } }
```

GPT Prediction

```
{ "2023": "$2.3 million", "2022": "$15.5 million", "Analysis": "The fair value of the cross-currency swaps decreased from $15.5 million in 2022 to $2.3 million in 2023, indicating a significant reduction in the value of these hedges over the year. This could suggest a decrease in effectiveness or adverse market conditions impacting the value of these financial instruments." }
```

LLaMA Analysis	GPT Analysis
<ul style="list-style-type: none"> • Quantitative approach with specific calculations • Calculated change in fair value: - \$13.2 million • Calculated percentage change: - 85.16% • Interpreted decrease as indicating effectiveness in reducing foreign exchange rate exposure • Provided a step-by-step analysis with precise calculations • Misinterpreted the decrease as a positive outcome • Lacked broader context in the analysis 	<ul style="list-style-type: none"> • Qualitative approach focusing on overall trends • Did not provide specific calculations • Interpreted decrease in fair value as suggesting reduced effectiveness or adverse market conditions • Offered a concise summary highlighting potential implications • Contextualized the change within broader market conditions • Lacked detailed quantitative analysis • Provided a more accurate assessment of the hedge’s effectiveness

D Complete Result

We layout the complete evaluation statistics of the previous 7 models, and 2 multi-modal variants using images as the context (gpt-4o, and gemini-1.5-pro) under different evaluation frameworks in Table 3.

E Dataset Details

We detail the input length by token count in Figure 3. As demonstrated by the green line, the inputs exceeding the 16k token limits are the outliers. Given that all but one model tested in our experiments accept 16k tokens, we believe that the truncation employed to fit context limit does not pose a significant disadvantage to open-source models with a smaller context limit.

Model	ELO Rating		Accuracy (%)		
	Overall	Hard	Overall	Easy	Medium
gpt-4o-vision	892.46	797.37	70.0	72.5	67.5
gemini-1.5-pro	896.81	779.98	66.25	70.0	62.5
gpt-4-turbo	1164.35	1045.06	65.625	63.75	67.5
claude-3-5-sonnet-20240620	1081.39	832.23	65.625	67.5	63.75
Qwen2-72B-Instruct	867.81	793.03	65.0	65.0	65.0
gemini-1.5-flash	880.88	779.98	64.375	68.75	60.0
gemini-1.5-pro-vision	760.52	639.71	54.375	62.5	46.25
Meta-Llama-3-70B-Instruct	772.05	606.52	51.875	53.75	50.0
Meta-Llama-3.1-70B-Instruct	869.27	703.98	48.75	48.75	48.75

Table 3: Model Performance Comparison: ELO Rating and Accuracy

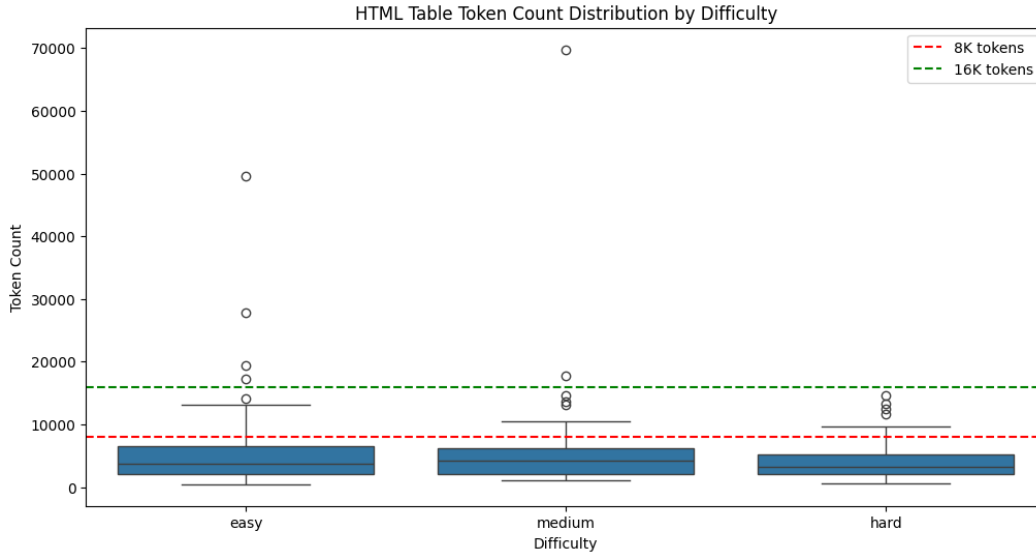


Figure 3: Input lengths statistics.

Length	Distribution (%)	Recall
0-5	26.7%	0.704
5-10	15.0%	0.788
10-15	19.3%	0.816
15-20	15.8%	0.782
20+	23.2%	0.642
Avg.	100%	0.736

Table 4: Metrics by Answer Length. The recall slump within the short-form answer group (0-5 tokens) is primarily attributable to the slight differences in calculated and claimed number strings.

F Self-supervised question-answer synthesis validation

Despite extensive efforts from both human and machine annotators, the complexity and the domain knowledge required to validate whether the questions are meaningful, or whether the answer aptly responds to the question, i.e. whether our external supervision setup was, in fact, factual - remains a challenging point to prove. Therefore, in an attempt to verify the efficacy of the data synthesis pipeline, we adapted it for a new domain, diversifying our data sources. More importantly, this adaptation enabled us to leverage self-supervision signals within the context document to fact-check synthesized answers reliably.

Specifically, we selected recent peer-reviewed papers from reputable artificial intelligence conferences as source material for our data synthesis pipeline. Our selection was based on three reasons: (1) Conference papers in this domain often contain rich tabular content, which is frequently referenced in the main text with detailed analysis; (2) The expert peer-review process ensures a certain quality of content; and (3) The AI field has a high volume of contributions, providing a solid foundation for scaled analysis.

We collected 936 academic papers from January 2022 to May 2024, utilizing a subset of a recently released dataset [26] focused on long-context reasoning. Since the data was obtained in markdown format through OCR from the original PDFs, it was incompatible with our existing synthesis pipeline, which was designed for HTML-format handling. We employed several preprocessing steps: (a) splitting each markdown paper into chunks according to its native formatting and extracting tables and headers using regular expressions; (b) gathering all chunks that explicitly reference each table to serve as context for question synthesis; and (c) converting all markdown tables into canonical HTML

form using the *markdownify* package. These steps transformed the original data source into a format compatible with our synthesis pipeline and provided high-quality evidence for verifying synthetic answers. All other configurations, including model settings and prompts, were kept consistent with the original setup to preserve the validity of the validation process. Synthetic questions are generated with a requirement to cite a grounding sentence from the evidence pool, while synthetic answers are created based solely on the corresponding tables and questions, without referencing direct evidence from the paper. Following quality filtering procedures, we obtained 2,100 question-answer pairs, with 700 representing each difficulty level.

To evaluate the faithfulness of synthetic answers to the source, we used recall as our evaluation metric. Specifically, we leveraged the *nlk* package to tokenize each answer string and the corresponding evidence string, which theoretically contains the answer. We then identified the common tokens between the two and calculated the recall rate based on the total tokens in the answer string. The results (as shown in Table 4) indicate an average recall rate of 0.74, demonstrating our synthesis model’s (GPT-4-turbo) strong capability to produce accurate answers that align with the source.

In conclusion, our adaptable data synthesis pipeline demonstrates robust performance in generating factually grounded question-answer pairs, particularly through its effective handling of domain-specific content like academic papers with tabular data leveraging a self-supervision signal. This versatility highlights its potential for applications across diverse domains that require structured and unstructured data synthesis, making it a valuable tool for generating and verifying knowledge in areas such as long-context reasoning and fact-checking. Future extensions of this pipeline could further enhance its applicability by incorporating additional domains and refining its adaptability to varying data formats.