# EXTRACT, MATCH, AND SCORE: AN EVALUATION PARADIGM FOR LONG QUESTION-CONTEXT-ANSWER TRIPLETS IN FINANCIAL ANALYSIS

**Bo Hu<sup>†</sup> & Han Yuan<sup>†</sup> & Vlad Pandelea<sup>†</sup> & Wuqiong Luo<sup>†</sup> & Yingzhu Zhao<sup>†</sup> & Zheng Ma<sup>⊠</sup>** American Express Global Decision Science

{Bo.Hu, Han.Yuan1, Vlad.A.Pandelea, Wuqiong.Luo,Yingzhu.Zhao, Zheng.Ma2}@aexp.com

#### ABSTRACT

The rapid advancement of large language models (LLMs) has sparked widespread adoption across diverse applications, making robust evaluation frameworks crucial for assessing their performance. While conventional evaluation metrics remain applicable for shorter texts, their efficacy diminishes when evaluating the quality of long-form answers. This limitation is particularly critical in real-world scenarios involving extended questions, extensive context, and long-form answers, such as financial analysis or regulatory compliance. In this paper, we use a practical financial use case to illustrate applications that handle "long question-context-answer triplets". We construct a real-world financial dataset comprising long triplets and demonstrate the inadequacies of traditional metrics. To address this, we propose an effective Extract, Match, and Score (EMS) evaluation approach tailored to the complexities of long-form LLMs' outputs, providing practitioners with a reliable methodology for assessing LLMs' performance in complex real-world scenarios.

# **1** INTRODUCTION

The advancement of large language models (LLMs) has revolutionized automated model-based solutions to an unprecedentedly accuracy that rivals human performance across a series of natural language processing (NLP) tasks (Yang et al., 2024). Despite their impressive performance on short-form tasks (Chang et al., 2024), the applicability and effectiveness of LLMs in long-form scenarios warrant in-depth exploration, particularly those where short-form summaries lead to substantial information loss (Shaham et al., 2023; Wijesiriwardene et al., 2023). Preliminary studies have investigated long-form challenges in the general domain using either standalone LLMs (Han et al., 2023; Liu et al., 2024; Li et al., 2024; Wang et al., 2024; Laban et al., 2024) or retrievalaugmented generation (RAG) systems (Su et al., 2022; Yu et al., 2023; Zhao et al., 2024a; Qi et al., 2024; Han et al., 2024; Laban et al., 2024).

Compared to the general domain, long-form scenarios are particularly prevalent in the financial sector, where documents such as financial statements and compliance reports often exceed hundreds of pages (Liu et al., 2022; Reddy et al., 2024; Zhao et al., 2024b). However, the absence of effective evaluation metrics poses a significant challenge in financial NLP. In particular, traditional metrics lose their efficacy as evaluation tasks increase in length and complexity. Consequently, many studies in financial NLP simplify complex, open-ended analytical tasks into well-defined, closed-ended tasks such as news sentiment analysis (Guo et al., 2023; Chen et al., 2024b; Kirtac & Germano, 2024; Shah et al., 2023), concept classification (Fonseca & Cohen, 2024), name entity recognition (Shah et al., 2022; Guo et al., 2023; Bhatia et al., 2024), structure boundary detection (Shah et al., 2022), event extraction (Zhou et al., 2021), relation identification (Li et al., 2023; Reddy et al., 2024; Krumdick et al., 2024; Chen et al., 2024a) and abstractive summary (Zhu et al., 2024), they typically

<sup>&</sup>lt;sup>†</sup> These authors contributed equally to this work.

<sup>&</sup>lt;sup>™</sup> Correspondence: Zheng Ma, Singapore Decision Science Center of Excellence, American Express, 1 Marina Boulevard, 018989, Singapore.



Figure 1: An example of saliency point extraction from long-form answer.

address only one or two of the three key aspects that possess the long-form property: context (Reddy et al., 2024; Srivastava et al., 2024; Zhu et al., 2024; Xie et al., 2023; Zhao et al., 2024b; Gupta et al., 2024), question (Krumdick et al., 2024; Chen et al., 2022), and answer (Band et al., 2024; Qi et al., 2024). To the best of our knowledge, while a few efforts have been made in the general domain (Tan et al., 2024), no existing work directly addresses the evaluation challenge of long triplets in finance.

To systematically investigate the current landscape and advance the research on financial long triplets, we (1) Benchmarked the state-of-the-art LLMs on long-form analysis of earning call transcripts from 10 largest constituents in S&P 500 index; (2) Identified the inefficiency of conventional evaluation methods, such as ROUGE, in differentiating generation quality within the context of long triplets (Xu et al., 2023; Tan et al., 2024); (3) Proposed a generalizable evaluation paradigm named EMS (Extract, Match, and Score) to provide fine-grained evaluations and demonstrated its competitive advantages over RAGChecker (Ru et al., 2024) in real-world scenarios of long triplets.

#### 2 Methodology

To address the inefficiency of conventional evaluation approaches in the context of long triplets, we propose a three-stage methodology EMS : (1) Extracts salient points from both the reference and candidate answers; (2) Identifies potential alignments between these points; (3) Scores each matched pair based on predefined criteria. The long-form generated answer is denoted as ans, and the corresponding reference is denoted as ref. Our evaluation pipeline takes ans and ref as input, and provides three saliency points level metrics: EMS-Recall, EMS-Precision, EMS-F1.

#### 2.1 STAGE 1: SALIENCY POINT EXTRACTION

Given a long-form reference answer, we define a function  $f^{(E)}(\cdot)$  that extracts a list of saliency points from the text. Formally,

$$P^{(\text{ref})} = f^{(\text{E})}(ref) = [p_1^{(\text{ref})}, \dots, p_N^{(\text{ref})}],$$
(1)

where  $P^{(\text{ref})}$  denotes the set of saliency points extracted from the reference. We apply the same function  $f(\cdot)$  to the candidate answer:

$$P^{(\text{ans})} = f^{(\text{E})}(ans) = [p_1^{(\text{ans})}, \dots, p_M^{(\text{ans})}],$$
(2)

yielding the set  $P^{(ans)}$ , which contains the saliency points extracted from the candidate answer.

The purpose of saliency points extraction is to decompose the long-form answer into key claims, ensuring that the evaluation score of these claims accurately reflects the quality of the original response. To achieve this, the saliency point extractor should follow well-crafted instructions.

**Including details.** Saliency points are not brief summaries from the original answer, since the detailed information is useful to users and shows as a sign of good quality. Hence, saliency points should maintain all the details from the original long-form answer, so that the scoring over saliency points can represent the quality of original answers. Figure 1 shows an example that the detailed numbers related to the revenue is critical to users with finance background. The shorter extracted point is factually accurate, however, it filters out indispensable information.

**Retaining repeated points.** Low-quality answers are not solely the result of missing or incorrect claims, which can be penalized during scoring, but also stem from poor organization of those claims. Specifically, a long-form answer should be assigned with a low score if it contains repeated content, even when the saliency points from those repetitions are accurate. Therefore, all repeated points are retained during extraction, leading to a saliency point list that includes all duplicates.

**Removing summary.** An exception to the treatment of repeated content is with the summaries from the long-form answers. High-quality long-form answers typically include a structured response with an introduction, detailed body, and conclusion. Since the content in the introduction and conclusion often overlaps with the body, including it as duplicated points would be unnecessary. To address this, we exclude summaries, such as overviews and conclusions, before extracting saliency points.

The extractor can be implemented with several solutions, *e.g.*, fuzzy search for summary removal and retaining all details in the saliency points by directly splitting the long-form answer into sentences. This approach works well for answers where the middle section is already well-organized with bullet points. However, it struggles with answers that contain extensive details and examples. To address this, we utilize an LLM-based pipeline to remove summaries from the beginning and the end, then extract saliency points following the specified instructions discussed above. To control the level of detail in the saliency points, we employ in-context learning (ICL) with a few-shot example setup. The prompt used for saliency point extraction is available in Appendix A.3.

#### 2.2 STAGE 2: MATCHING

In the second stage of EMS, we apply a function  $f^{(M)}$  to establish correspondences between points in  $P^{(\text{ref})}$  and  $P^{(\text{ans})}$ . The output of the matching stage is

$$A = [a_1, a_2, \dots, a_N],\tag{3}$$

where each index  $a_i \in \{-1, 1, 2, ..., M\}$ , is computed with the *i*-th reference saliency point and the candidate answer saliency point list by:

$$a_i = f^{(M)}\left(p_i^{(\text{ref})}, P^{(\text{ans})}\right).$$
(4)

In the function  $f^{(M)}$ , the reference saliency point  $p_i^{(\text{ref})}$  is compared against all answer saliency points in the list  $P^{(\text{ans})}$  to evaluate the semantic matching. If the *i*-th reference saliency point  $p_i^{(\text{ref})}$  finds the best match from  $P^{(\text{ans})}$ , the output  $a_i = j$ , where *j* is the index of the best-matched answer saliency point. If no match is found, the output is  $a_i = -1$ . As illustrated in Figure 2, there are three matched pair of saliency points, therefore, the output of matching stage for this example is [4, 2, 3, -1].

To achieve this, an LLM-based matcher is proposed, leveraging the chain-of-thought (CoT) technique to guide the LLM to evaluate each answer saliency point in  $P^{(ans)}$  sequentially and to identify the best match (if one exists). Moreover, the prompt for matcher is highly flexible, allowing for specific instructions depending on the task requirements, *e.g.*, a good answer should focus more on numerical details such as the revenue or cost during the matching stage. Saliency point matching prompts used by LLM are listed in Appendix A.4.

## 2.3 STAGE 3: SCORING

A scoring function  $f^{(S)}$  is introduced in the last stage of EMS evaluation pipeline. Given  $p_i^{(\text{ref})}$  and matched answer point  $p_{a_i}^{(\text{ans})}$ , the alignment score is computed as

$$s_{i}^{(\text{ref})} = \begin{cases} f^{(\text{S})}\left(p_{i}^{(\text{ref})}, p_{a_{i}}^{(\text{ans})}\right) & a_{i} \neq -1, \\ 0 & a_{i} = -1, \end{cases}$$
(5)

where  $s_i^{(\text{ref})} \in [0, 1]$  and i = 1, ..., N.

An LLM is employed for this stage, which estimates the score for a pair of saliency points based on the extent to which the details of two saliency points are aligned. As in the previous stages, the LLM provides granular flexibility to accommodate task-specific criteria in addition to the general ones, which are defined in our pipeline as follows:



Figure 2: Illustration of matching and scoring procedure in EMS evaluation pipeline. Then EMS-Recall and EMS-Precision are computed by aggregating saliency point level scores.

- **Completeness**: Does the candidate answer point capture all the essential aspects of the reference point?
- Accuracy: Is the information in the candidate answer point contextually correct?
- Level of Detail: Does the candidate answer point provide sufficient details, matching the depth of the reference point?

While the matching stage primarily focuses on ranking and assigning binary scores, *i.e.*, 1 for matched pairs and 0 for unmatched pairs, the scoring stage differs by assigning a soft score (continuous from 0 to 1) to each matched pair of saliency points. With soft scores, the detailed level of semantic alignment between two saliency points can be accurately reflected in the metrics. It is more useful than a binary score as a reference for fine-tuning the LLM or adjusting the prompt to have more accurate LLM-generated answers. The details scoring prompt is in Appendix A.5.

Although the matching and scoring stages can be handled by a single LLM for both tasks, using one LLM for a combined and complex task may result in poorer performance and higher costs compared to using multiple LLMs, each handling a simpler sub-task (Yue et al., 2024). Furthermore, having a separate scoring stage allows the pipeline to integrate with conventional NLP metrics. For example, if ROUGE is chosen as the scoring mechanism, output  $s_i^{(ref)}$  becomes the ROUGE score between two saliency points. Since long-form answers are decomposed into saliency points, which are typically just one or two sentences long, conventional metrics perform better in this scenario than when comparing long-form answers directly.

#### 2.4 EMS METRICS

After the Extract-Match-Score pipeline, the long-form answer and reference are converted to a matching list with the corresponding alignment scores assigned for the matched pairs. In our EMS pipeline, EMS-Recall, EMS-Precision, and EMS-F1 are computed as the metrics.

**EMS-Recall** is the "proportion" of the long-form reference which are accurately captured in the semantic level, which is computed by

$$\text{EMS-Recall} = \frac{\sum_{i=1}^{N} s_i^{(\text{ref})}}{N}$$
(6)

As  $s_i^{(\text{ret})}$  is the soft score, EMS-Recall represents the recall calculated over the level of detailed information, instead of a simple recall over the sentences or claims.

**EMS-Precision** is the "proportion" of the long-form answer that is aligned with the reference. To calculate EMS-Precision, the alignment score for reference saliency points  $s_i^{(\text{ref})}$  are first mapped to

answer saliency points by

$$s_j^{(\text{ans})} = \begin{cases} \max\left(\{s_i^{(\text{ref})}|a_i=j\}\right) & j \in A, \\ 0 & j \notin A. \end{cases}$$
(7)

It is important to note that in the implementation, multiple reference saliency points may match to one single answer saliency point. It may be caused by the repetition of similar information within the long-form reference. In this case, the maximum score is selected to represent the alignment of this answer saliency point. Subsequently, EMS-Precision is computed as

$$\text{EMS-Precision} = \frac{\sum_{j=1}^{M} s_j^{(\text{ans})}}{M}$$
(8)

Similarly, EMS-Precision refers to information-level precision of the LLM-generated long-form answer.

#### **EMS-F1** is computed by

$$EMS-F1 = \frac{2 \times EMS-Precision \times EMS-Recall}{EMS-Precision + EMS-Recall}.$$
(9)

As discussed in Section 2.3, the score  $s_i$  can be derived from either LLM-based scorer or conventional NLP metrics. With the ROUGE score example, ROUGE-F1 is employed for saliency point-level scoring. Subsequently, EMS-F1 is calculated to evaluate the entire long-form answer. As several scorers are implemented and tested, to distinguish pipelines between different scorers, in the following sections:

- EMS<sup>(LLM)</sup> denotes the EMS evaluation pipeline with LLM-based scorer;
- EMS<sup>(ROUGE)</sup> denotes the pipeline where the scorer is implemented by ROUGE-F1;
- EMS<sup>(BERTScore)</sup> denotes the pipeline where the scorer is implemented by BERT score.

The long-form answer level recall, precision, and F1 are then calculated with the outputs from different scorers, *e.g.*, EMS<sup>(LLM)</sup>-F1.

# **3** EXPERIMENT RESULTS

#### 3.1 DATASET

We have constructed a financial Document-Based QA dataset specifically focusing on long-context and long-form answers. This dataset uses the earnings call transcripts from the third quarter of 2024 for the 10 largest S&P 500 constituents at the time of writing this paper, namely, *Apple, NVIDIA, Microsoft, Alphabet, Amazon, Meta Platforms, Tesla, Broadcom, JPMorgan Chase*, and *Eli Lilly.* The dataset captures a rich set of financial discussions, encompassing detailed business strategies, performance analyses, and forward-looking statements, which is a useful resource for financial NLP research and applications.

We analyze the Q&A sections of several earnings call transcripts to identify common themes and question formats used by analysts. We then use these insights to design five questions that cover commonly referred topics. The questions are further refined iteratively with the assistance of LLMs in order to produce high-quality, contextually relevant answers that reflect real-world financial analysis. Each question is designed to address specific aspects of company performance, such as revenue growth, operational challenges, strategic initiatives, financial forecasts, and competitive positioning. The five questions are listed in Appendix A.1.

To construct high-quality answers that serve as a reference for the evaluation of other LLMgenerated answers, we leverage the state-of-the-art LLMs (GPT-4o and Mistral Large). We prompt each LLM to answer each question individually and then generate a consolidated answer by instructing GPT-4o to combine and refine the answers, taking into account their strengths and resolving any discrepancy by referring to the original document. The complete prompts are shown in Appendix A.2.

#### 3.2 **BASELINE METHODS**

We consider both classical numerical metrics and more recent LLM-based evaluation frameworks.

- **BLEU:** BLEU (Papineni et al., 2002) measures the match between a candidate answer and a reference answer by looking at the percentage of overlapping n-grams.
- **ROUGE:** While BLEU focuses on precision, measuring how many n-grams in the candidate answer appear in the reference answer, ROUGE (Lin, 2004) emphasizes recall, evaluating how much of the reference answer is captured by the candidate.
- **BERTScore:** BERTScore (Zhang et al., 2020) uses the pre-trained contextualized token embeddings from BERT to calculate the similarity between candidate answer and reference sentences.
- **RAGChecker:** RAGChecker (Ru et al., 2024) is a latest LLM-based evaluation framework. Although designed for RAG systems, it can assess end-to-end performance via its precision and recall metrics by extracting claims from the model responses and then verifying their entailment against the provided ground truth.

#### 3.3 RESULTS

We have shown the evaluation results averaged across ten companies and all the questions in Table 1. Three LLM models are selected from the Llama3.2 family with the number of model parameters ranging from 1 billion to 90 billion. By this means, we avoid the performance nuances brought by different LLM variations and focus on validating the scaling effect of LLMs with different evaluation metrics.

Table 1 presents that both RAGChecker and our proposed EMS metrics show more nuanced assessments of the response quality. We could clearly see the upward trend of the evaluation figures as the LLM model size increases from 1 bilion to 90 billion. In contrast, conventional evaluation metrics such as BLEU, ROUGE and BERTScore fail to capture this trend with different model sizes. Based on the well-established understanding that larger models offer superior performance, our proposed EMS evaluation is more effective in long-form financial analyses.

# 3.4 DISCUSSION

As mentioned in Section 3.3, compared to conventional metrics such as BLEU, ROUGE and BERTScore, both our EMS approach and RAGChecker offer significantly more nuanced assessments of answer quality. BLEU and ROUGE compute n-grams level matching, which may not capture the overall meaning or the entire context of the response. In addition, the F1 scores produced by RAGChecker and our method EMS<sup>(LLM)</sup> both range between 0.2 and 0.5, indicating that a naive question-answer approach for long triplets in financial analysis may not yield satisfactory results. Practitioners should therefore expect to invest substantial effort into refining their solutions. Our proposed EMS evaluation metric, at the same time, could be helpful in evaluating different solutions effectively, aiding practitioners in selecting suitable backbones for financial analysis.

Compared to the competing approach of RAGChecker, our proposed method offers three key advantages. First, EMS extracts the saliency points, which are carefully designed to maintain all the details from the original long-form answer and could provide practitioners tangible insights on the overall financial analysis.

Second, our extensible EMS method opens the door to multiple scoring mechanisms, including conventional methods, such as BLEU and ROUGE, non-LLMs model-based scoring like BERTScore, and LLM-based scoring. This flexibility allows practitioners to incorporate the advantages of different evaluation metrics and to tailor their scoring strategy based on specific requirements and computation resources. In contrast, RAGChecker is constrained by its configuration reliance on a LLM to generate scores.

Third, rather than the binary (hit-or-miss) outcome used by RAGChecker, our approach employs a continuous scale (0-1) scoring mechanism, enabling finer-grained distinctions in answer quality. This is particularly valuable in pinpointing specific areas of improvement for long-form financial analyses, offering practitioners clearer guidance on how to enhance system performance.

Metric/LLM		llama3.2_1b_ instruct	llama3.2_11b_ vision_instruct	llama3.2_90b_ vision_instruct
BLEU		0.02	0.04	0.03
ROUGE	Precision	0.23	0.21	0.22
	Recall	0.14	0.21	0.19
	F1	0.17	0.20	0.20
BERTScore		0.84	0.85	0.85
RAGChecker	Precision	0.45	0.63	0.63
	Recall	0.15	0.31	0.35
	F1	0.21	0.40	0.44
EMS <sup>(ROUGE)</sup>	Precision	0.08	0.17	0.17
	Recall	0.07	0.12	0.13
	F1	0.07	0.13	0.15
EMS <sup>(BERTScore)</sup>	Precision	0.38	0.63	0.65
	Recall	0.34	0.42	0.47
	F1	0.35	0.49	0.54
EMS <sup>(LLM)</sup>	Precision	0.23	0.45	0.47
	Recall	0.21	0.30	0.36
	F1	0.21	0.35	0.40

 Table 1: Evaluation Results of Ten Companies over All Questions

#### 4 CONCLUSION

In this study, we first highlight the limitations of conventional evaluation approaches in distinguishing generation quality for financial long triplets. To address this issue, we propose an extensible framework, EMS, that assesses answer quality by extracting saliency points, matching with answer saliency points, and scoring the alignment. We conducted experiments on a self-constructed financial Document-Based QA dataset and provided detailed analysis on EMS over both conventional and latest approaches. We envisage that EMS could aid practitioners in answer assessment within the context of long triplets, and facilitate researchers to explore this complex yet essential scenario beyond financial NLP.

#### LIMITATION

This study has three main limitations. First, we primarily evaluated the inference capabilities of LLMs on qualitative information of earning call transcripts and did not dive into the mathematical reasoning, which is crucial in financial NLP (Chen et al., 2021; 2022; Krumdick et al., 2024; Srivastava et al., 2024). Second, LLMs were employed as the primary annotator to generate silver answers, which may introduce biases and errors compared to panel discussions by human experts (Felkner et al., 2024; Rønningstad et al., 2024). We view our effort as an initial study and future studies should recruit financial professionals to develop the golden answer and validate the effectiveness of the proposed EMS evaluation. Third, the current evaluation mainly focused on the schematic comparison and future work will develop a more systematic (Potluri et al., 2023; Koh et al., 2022) and multifaceted approach to consider additional factors such as logical consistency and factuality (Cho et al., 2019; Min et al., 2023; Song et al., 2024).

#### ETHICS STATEMENT

To facilitate reproducibility and ensure the consistency of our findings, we have included the detailed prompt in Appendix. The original earnings call transcripts are publicly available. We acknowledge

the potential for geographic bias in our dataset, as all samples were derived from constituents of the S&P index. Publicly accessible language models were employed at various stages of our experiments. Given their inherent stochasticity, these models may generate unexpected outputs. To mitigate risks to model safety, we refrained from fine-tuning the models.

#### DISCLAIMER

This paper is intended solely for informational purposes and the technique is not a business practice of American Express. The earnings call transcripts is used for demonstration of our evaluation framework and analysis on them is not intended to constitute investment research, advice, recommendations, or an offer to buy or sell securities. The opinions, findings and conclusions of this paper are those of the authors alone and do not reflect the views of American Express.

#### REFERENCES

- Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. Linguistic calibration of long-form generations. In *Proceedings of the International Conference on Machine Learning*, pp. 2732–2778, 2024.
- Gagan Bhatia, El Moatez Billah Nagoudi, Hasan Cavusoglu, and Muhammad Abdul-Mageed. Fin-Tral: A family of GPT-4 level multimodal financial large language models. In *Findings of the Annual Meeting of the Association for Computational Linguistics*, pp. 13064–13087, 2024.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. A survey on evaluation of large language models. ACM Transactions on Intelligent Systems and Technology, 15(3):45, 2024.
- Jian Chen, Peilin Zhou, Yining Hua, Loh Xin, Kehui Chen, Ziyuan Li, Bing Zhu, and Junwei Liang. FinTextQA: A dataset for long-form financial question answering. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 6025–6047, 2024a.
- Tianyu Chen, Yiming Zhang, Guoxin Yu, Dapeng Zhang, Li Zeng, Qing He, and Xiang Ao. EFSA: Towards event-level financial sentiment analysis. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 7455–7467, 2024b.
- Zhiyu Chen, Wenhu 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 Conference on Empirical Methods in Natural Language Processing*, pp. 3697–3711, 2021.
- Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah, and William Yang Wang. ConvFinQA: Exploring the chain of numerical reasoning in conversational finance question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 6279–6292, 2022.
- Woon Sang Cho, Pengchuan Zhang, Yizhe Zhang, Xiujun Li, Michel Galley, Chris Brockett, Mengdi Wang, and Jianfeng Gao. Towards coherent and cohesive long-form text generation. In *Proceed*ings of the Workshop on Narrative Understanding, pp. 1–11, 2019.
- Virginia Felkner, Jennifer Thompson, and Jonathan May. GPT is not an annotator: The necessity of human annotation in fairness benchmark construction. In *Proceedings of the Annual Meeting of* the Association for Computational Linguistics, pp. 14104–14115, 2024.
- Marcio Fonseca and Shay Cohen. Can large language models follow concept annotation guidelines? a case study on scientific and financial domains. In *Findings of the Annual Meeting of the Association for Computational Linguistics*, pp. 8027–8042, 2024.
- Yue Guo, Zian Xu, and Yi Yang. Is ChatGPT a financial expert? evaluating language models on financial natural language processing. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pp. 815–821, 2023.

- Lavanya Gupta, Saket Sharma, and Yiyun Zhao. Systematic evaluation of long-context LLMs on financial concepts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 1163–1175, 2024.
- Rujun Han, Yuhao Zhang, Peng Qi, Yumo Xu, Jenyuan Wang, Lan Liu, William Yang Wang, Bonan Min, and Vittorio Castelli. RAG-QA arena: Evaluating domain robustness for long-form retrieval augmented question answering. In *Proceedings of the Conference on Empirical Methods* in Natural Language Processing, pp. 4354–4374, 2024.
- Wookje Han, Jinsol Park, and Kyungjae Lee. PreWoMe: Exploiting presuppositions as working memory for long form question answering. In *Proceedings of the Conference on Empirical Meth*ods in Natural Language Processing, pp. 8312–8322, 2023.
- Kemal Kirtac and Guido Germano. Enhanced financial sentiment analysis and trading strategy development using large language models. In *Proceedings of the Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pp. 1–10, 2024.
- Huan Yee Koh, Jiaxin Ju, He Zhang, Ming Liu, and Shirui Pan. How far are we from robust long abstractive summarization? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 2682–2698, 2022.
- Michael Krumdick, Rik Koncel-Kedziorski, Viet Dac Lai, Varshini Reddy, Charles Lovering, and Chris Tanner. BizBench: A quantitative reasoning benchmark for business and finance. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 8309–8332, 2024.
- Philippe Laban, Alexander Fabbri, Caiming Xiong, and Chien-Sheng Wu. Summary of a haystack: A challenge to long-context LLMs and RAG systems. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 9885–9903, 2024.
- Xianzhi Li, Samuel Chan, Xiaodan Zhu, Yulong Pei, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah. Are ChatGPT and GPT-4 general-purpose solvers for financial text analytics? a study on several typical tasks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 408–422, 2023.
- Zhuowan Li, Cheng Li, Mingyang Zhang, Qiaozhu Mei, and Michael Bendersky. Retrieval augmented generation or long-context LLMs? a comprehensive study and hybrid approach. In Proceedings of the Conference on Empirical Methods in Natural Language Processing: Industry Track, pp. 881–893, 2024.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization* branches out, pp. 74–81, 2004.
- Shuaiqi Liu, Jiannong Cao, Ruosong Yang, and Zhiyuan Wen. Long text and multi-table summarization: Dataset and method. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pp. 1995–2010, 2022.
- Xiang Liu, Peijie Dong, Xuming Hu, and Xiaowen Chu. LongGenBench: Long-context generation benchmark. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pp. 865–883, 2024.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the Conference on Empirical Methods* in Natural Language Processing, pp. 12076–12100, 2023.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- Abhilash Potluri, Fangyuan Xu, and Eunsol Choi. Concise answers to complex questions: Summarization of long-form answers. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 9709–9728, 2023.

- Zehan Qi, Rongwu Xu, Zhijiang Guo, Cunxiang Wang, Hao Zhang, and Wei Xu. *long*<sup>2</sup>*rag*: Evaluating long-context & long-form retrieval-augmented generation with key point recall. In *Findings* of the Conference on Empirical Methods in Natural Language Processing, pp. 4852–4872, 2024.
- Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai, Michael Krumdick, Charles Lovering, and Chris Tanner. DocFinQA: A long-context financial reasoning dataset. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 445–458, 2024.
- Egil Rønningstad, Erik Velldal, and Lilja Øvrelid. A GPT among annotators: LLM-based entitylevel sentiment annotation. In *Proceedings of The Linguistic Annotation Workshop*, pp. 133–139, 2024.
- Dongyu Ru, Lin Qiu, Xiangkun Hu, Tianhang Zhang, Peng Shi, Shuaichen Chang, Cheng Jiayang, Cunxiang Wang, Shichao Sun, Huanyu Li, et al. Ragchecker: A fine-grained framework for diagnosing retrieval-augmented generation. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, 2024.
- Agam Shah, Suvan Paturi, and Sudheer Chava. Trillion dollar words: A new financial dataset, task & market analysis. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 6664–6679, 2023.
- Raj Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Raman, Charese Smiley, Jiaao Chen, and Diyi Yang. When FLUE meets FLANG: Benchmarks and large pretrained language model for financial domain. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 2322–2335, 2022.
- Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. ZeroSCROLLS: A zero-shot benchmark for long text understanding. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pp. 7977–7989, 2023.
- Yixiao Song, Yekyung Kim, and Mohit Iyyer. VeriScore: Evaluating the factuality of verifiable claims in long-form text generation. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pp. 9447–9474, 2024.
- Pragya Srivastava, Manuj Malik, Vivek Gupta, Tanuja Ganu, and Dan Roth. Evaluating LLMs' mathematical reasoning in financial document question answering. In *Findings of the Annual Meeting of the Association for Computational Linguistics*, pp. 3853–3878, 2024.
- Dan Su, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. Read before generate! faithful long form question answering with machine reading. In *Findings of the Annual Meeting of the Association for Computational Linguistics*, pp. 744–756, 2022.
- Haochen Tan, Zhijiang Guo, Zhan Shi, Lu Xu, Zhili Liu, Yunlong Feng, Xiaoguang Li, Yasheng Wang, Lifeng Shang, Qun Liu, and Linqi Song. ProxyQA: An alternative framework for evaluating long-form text generation with large language models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 6806–6827, 2024.
- Yuqi Wang, Lyuhao Chen, Songcheng Cai, Zhijian Xu, and Yilun Zhao. Revisiting automated evaluation for long-form table question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 14696–14706, 2024.
- Thilini Wijesiriwardene, Ruwan Wickramarachchi, Bimal Gajera, Shreeyash Gowaikar, Chandan Gupta, Aman Chadha, Aishwarya Naresh Reganti, Amit Sheth, and Amitava Das. ANALOGI-CAL a novel benchmark for long text analogy evaluation in large language models. In *Findings* of the Annual Meeting of the Association for Computational Linguistics, pp. 3534–3549, 2023.
- Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. Pixiu: A comprehensive benchmark, instruction dataset and large language model for finance. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, pp. 33469–33484, 2023.
- Fangyuan Xu, Yixiao Song, Mohit Iyyer, and Eunsol Choi. A critical evaluation of evaluations for long-form question answering. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 3225–3245, 2023.

- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *ACM Transactions on Knowledge Discovery from Data*, 18(6):32, 2024.
- Mengxia Yu, Sara Rosenthal, Mihaela Bornea, and Avi Sil. Towards effective long-form QA with evidence augmentation. In *Proceedings of the Workshop on Natural Language Generation, Evaluation, and Metrics*, pp. 155–164, 2023.
- Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. Large language model cascades with mixture of thought representations for cost-efficient reasoning. In *Proceedings of the International Conference on Learning Representations*, 2024.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *Proceedings of the International Conference on Learning Representations*, 2020.
- Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong, and Jie Tang. LongRAG: A dual-perspective retrieval-augmented generation paradigm for long-context question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 22600–22632, 2024a.
- Yilun Zhao, Yitao Long, Tintin Jiang, Chengye Wang, Weiyuan Chen, Hongjun Liu, Xiangru Tang, Yiming Zhang, Chen Zhao, and Arman Cohan. FinDVer: Explainable claim verification over long and hybrid-content financial documents. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 14739–14752, 2024b.
- Zhihan Zhou, Liqian Ma, and Han Liu. Trade the event: Corporate events detection for news-based event-driven trading. In *Findings of the Annual Meeting of the Association for Computational Linguistics*, pp. 2114–2124, 2021.
- Jie Zhu, Junhui Li, Yalong Wen, and Lifan Guo. Benchmarking large language models on CFLUE a Chinese financial language understanding evaluation dataset. In *Findings of the Annual Meeting* of the Association for Computational Linguistics, pp. 5673–5693, 2024.

# A APPENDIX

A.1 FIVE QUESTIONS DESIGNED FOR THE FINANCIAL QA DATASET

#### **QUESTION 1**

Can you provide a detailed analysis of the key themes, topics, and forward-looking guidance presented during the earnings call? Please include specific figures, strategic initiatives, and performance metrics highlighted by the executives, as well as any macroeconomic factors that may influence the business. Additionally, how do these elements shape the company's future trajectory and overall market position?

#### **QUESTION 2**

In the context of an earnings call, could you outline the specific risks and uncertainties that management highlights regarding future performance? Please include insights into challenges related to investments, cybersecurity, and operational strategies, as well as any external factors that could significantly impact the company's operations and financial outlook. Additionally, please reference relevant information on risk factors from the company's recent filings or official communications and discuss how these challenges might affect the company's financial performance and future growth.

#### QUESTION 3

To gain insights into the company's strategic direction and priorities, what key elements and initiatives highlighted in the earnings call transcript reflect their approach to growth, innovation, and market positioning? Please elaborate on specific strategies related to product categories, research and development, capital expenditures, partnerships, and how these initiatives aim to enhance customer value and operational efficiency, ultimately influencing the company's trajectory moving forward.

#### QUESTION 4

In the context of an earnings call transcript, what significant industry trends and current macroeconomic conditions are addressed by company executives that may affect both the company's performance and the broader industry landscape? Please highlight insights related to competitive dynamics, commodity markets, and specific challenges faced by the company, including relevant data points and strategic considerations that underscore the influence of these trends and conditions.

#### QUESTION 5

In the context of an earnings call, how can we gain insights into the company's strategies and priorities regarding cash flow management? Specifically, what key components should we examine, such as cash flow from operations, capital expenditures, and free cash flow trends? Additionally, what factors contribute to the resilience and growth of free cash flow margins, and how does the organization balance investment opportunities with shareholder returns amidst market fluctuations? What does this reveal about the company's approach to financial performance and resource allocation for future growth initiatives?

## A.2 ANSWER GENERATION PROMPTS

## Individual LLM Answer Generation Prompt

Instructions:

1. Review the Question: Carefully read the provided question and ensure that you fully understand what is being asked.

2. Make sure that the answer is of high quality with the following consideration:

- Clarity: Is the answer easy to understand? Are any terms or concepts unclear?

- Completeness: Does the answer fully address the question, or are there missing aspects that need to be covered?

- Accuracy: Were there factual errors or outdated information in the previous answer? Is the answer well-supported by evidence?

- Conciseness: Is the answer concise without sacrificing important details? Avoid unnecessary verbosity.

3. Final Review: Once the draft of the answer is complete, re-read the answer to ensure that:

- All aspects of the question are addressed.

- The answer is factually accurate, clear, and well-supported by evidence. Don't include any citation in the final answer.

Question: <Question>

Figure 3: The prompt used to generate the individual answers from GPT-40 and Mistral Large.

# Consolidated Answer Generation Prompt Instructions: You will be provided with: 1. An earnings call transcript. 2. The original question asked. 3. Three versions of answers generated by different LLMs.

Your task is to combine these answers into a single, high-quality final response. Follow the instructions below carefully:

#### 1. Review the Original Question:

Carefully read the provided question to ensure you fully understand its requirements.

#### 2. Analyze the Provided Answers:

- Identify the strengths of each answer. Consider clarity, completeness, accuracy, and conciseness.
- Look for overlap, unique points, and discrepancies among the answers.

#### 3. Combine and Refine:

- Synthesize the best elements of all three answers into a cohesive, well-structured response.
- Ensure the final answer is clear, concise, and comprehensive, addressing all aspects of the original question.
- Resolve any discrepancies or conflicting information by considering the context of the earnings call transcript.

#### 4. Final Quality Check:

- Ensure the final response is factually accurate, clear, and free of redundancy.
- Avoid unnecessary verbosity while retaining all essential details.
- Do not include citations or direct references to the original sources.

#### **Output Format:**

Provide a single, polished answer that effectively combines the key points, figures, strategic insights, and forward-looking guidance from the provided answers. Ensure it fully addresses the original question.

# **Original Question:**

<Question>

# **Earnings Call Transcript (for reference):** <Attached pdf file>

# Answer Versions to Combine:

Answer Version 1: <Answer\_Version\_1> Answer Version 2: <Answer\_Version\_2>

Figure 4: The prompt used to form the final answer by combining the answers from GPT-40 and Mistral Large.

#### A.3 SALIENCY POINTS EXTRACTION PROMPT

#### Saliency Points Extraction Prompt

#### Task Description:

You are tasked with extracting bullet points (key points) from the provided Candidate Answer.

Guidelines:

Extraction Task: Focus on extracting the content without evaluating its factual accuracy. Conciseness: Each bullet point must be no longer than two sentences.

Include Overlap and Repetition: Ensure that the extracted key points retain all overlapping and repeated information as presented in the original content. Do not remove or merge

redundant details; if similar information appears multiple times, include it in the extracted key points as is.

Format: The output should be in a List of Strings, where each string is a distinct bullet point.

In-Context Examples:

Example 1

Candidate Answer:

Tesla Bot, also known as Optimus, is a humanoid robot developed by Tesla Inc. It is designed to perform repetitive or unsafe tasks, leveraging Tesla's advancements in AI and robotics. Standing approximately 5'8" tall and weighing around 125 pounds, Optimus features human-like proportions and a sleek design. It uses Tesla's AI technology, including computer vision and self-learning algorithms, to navigate and interact with its environment. The bot is envisioned as a tool to enhance productivity and safety in industrial, household, and other labor-intensive settings.

Bullet Points:

"Tesla Bot is a humanoid robot developed by Tesla Inc.",

"Tesla Bot is also known as Optimus.",

"Tesla Bot is designed to perform repetitive or unsafe tasks in industrial, household, and other labor-intensive settings.",

"Tesla Bot is approximately 5'8" tall and weighs around 125 pounds, with human-like proportions and a sleek design.",

"Tesla Bot uses AI technology to navigate and interact with its environment."

]

Example 2

Candidate Answer:

Lenovo is a global technology company headquartered in Beijing, China, and Morrisville, North Carolina, USA. Founded in 1984, it is renowned for designing, manufacturing, and selling computers, smartphones, servers, and other technology products. Lenovo is a market leader in PCs, offering popular ThinkPad and Yoga series laptops, alongside gaming-focused Legion products. Lenovo is a market leader in PCs, offering popular ThinkPad and Yoga series laptops, alongside gaming-focused Legion products.

**Bullet Points:** 

"Lenovo is a global technology company with headquarters in Beijing, China, and Morrisville, North Carolina, USA.",

"Founded in 1984, Lenovo specializes in designing, manufacturing, and selling computers, smartphones, servers, and other tech products.",

"Lenovo is a PC market leader, known for ThinkPad and Yoga laptops, as well as gaming-focused Legion products.",

"Lenovo is a PC market leader, known for ThinkPad and Yoga laptops, as well as gaming-focused Legion products."

1

Your Task:

Generate the bullet points for the following candidate answer based on the guidelines above.

Candidate Answer: <ans>

**Bullet Points:** 

Figure 5: The prompt used to extract saliency points from candidate answers.

# A.4 SALIENT POINTS MATCHING PROMPT

#### Saliency Points Matching Prompt

Task Description:

Your task is to identify the most matched keypoint in the provided Candidate Keypoint List for a given Reference Keypoint. Each item in the candidate keypoint list contains one index and one string representin the keypoint.

Matching Criteria:

The matched keypoint should represent the same information as the reference keypoint. If no keypoint in the candidate list matches the reference, return -1.

Pay special attention to matching numerical details between the two key points. If both key points mention the same specific numbers (e.g., "growth of 30%"), this significantly increases their likelihood of being a match. However, ensure the context of the numbers aligns. For example, "35% growth" should not be matched to "35 people" as the numbers pertain to entirely different contexts. Always evaluate numerical information within the broader context of the key points.

Input Format:

Reference Keypoint: "..." Candidate Keypoint List: [ 0: "...", 1: "...", 2: "..." ] Output Format:

The output should be a single integer index of the best-matched candidate keypoint. If no match is found, output -1. Do not include any explanations or reasoning in the output. In-Context Examples:

Example 1 Reference Keypoint: "Lenovo is a global technology company with headquarters in different countries." Candidate Keypoint List:

"Lenovo is a global technology company with headquarters in Beijing, China, and Morrisville, North Carolina, USA.",

"Founded in 1984, Lenovo specializes in designing, manufacturing, and selling computers, smartphones, servers, and other tech products.",

"Lenovo is a PC market leader, known for ThinkPad and Yoga laptops, as well as gaming-focused Legion products."

Matched Index: 0

Example 2 Reference Keypoint: "The alias of Tesla Bot is Optimus." Candidate Keypoint List:

"Tesla Bot is a humanoid robot developed by Tesla Inc.",

"Tesla Bot is also known as Optimus.",

"Tesla Bot is designed to perform repetitive or unsafe tasks in industrial, household, and other labor-intensive settings.",

"Tesla Bot is approximately 5'8" tall and weighs around 125 pounds, with human-like proportions and a sleek design.", 'Tesla Bot uses AI technology to navigate and interact with its environment." Matched Index: 1 Example 3 Reference Keypoint: "PCIe bifurcation allows a single PCIe slot to be divided into multiple lanes." Candidate Keypoint List: "Flexible Leasing Options: Offers a variety of leasing solutions, including wet lease, dry lease, and lease-purchase agreements tailored to meet airline requirements.", "Comprehensive Fleet Management: Provides access to a wide range of aircraft models, ensuring compatibility with operational needs and passenger demands.", "Cost-Effective Solutions: Reduces the financial burden of aircraft ownership through competitive lease terms and efficient asset utilization.", "Global Network: Connects airlines with a diverse pool of lessors and aircraft owners across the world." Matched Index: -1 Your Task: Find the best matched keypoint index from the candidate keypoint list for the given reference keypoint. **Reference Keypoint:** <ref> Candidate Keypoint List: <candid> Matched Index:

Figure 6: The prompt used to match candidate saliency points with reference saliency points.

# A.5 ALIGNMENT SCORING PROMPT

#### Alignment Scoring Prompt

Task Description
------------------

Evaluate the similarity between two key points provided by a language model and assign a score between 0 and {max\_score} based on the following criteria:

Score {max\_score}: If the two key points are exactly matched, including all details.

Intermediate Score: If the two key points share some information but are not fully matched. The more they match, the closer the score should be to  $\{max\_score\}$ .

Score 0: If the two key points are entirely irrelevant to each other.

Guidelines:

The higher the similarity, the higher the score (closer to {max\_score}). The lower the similarity, the lower the score (closer to 0). Output must be a single integer from 0 to {max\_score}. Do not provide any explanations or reasoning in your response. In-Context Examples: Example 1: Input: Keypoint 1: The Gemini model supports over 2 billion monthly users across products like Search, Google Cloud, YouTube, and Google Maps, with API calls increasing 14x in six months. Keypoint 2: Sundar Pichai mentions that Gemini is now available on GitHub Copilot, with over 2 billion monthly users across all seven products. Output: 7 Example 2: Input: Keypoint 1: Google Cloud revenue increased 35% YoY to \$11.4 billion, with a 17% operating margin, driven by AI solutions like Vertex AI and BigOuery. Keypoint 2: Alphabet's Cloud revenue grew 35% year-over-year, with operating margins increasing to 17%. Output: 10 Example 3: Input: Keypoint 1: AI-driven features like AI Overviews and Google Lens are transforming user experiences and increasing engagement. Keypoint 2: The executives discuss the benefits of GenAI, including reduced costs, greater customer engagement, and faster response times. Output: 5 Example 4: Input: Keypoint 1: Alphabet plans to advance its AI portfolio with the next-generation Gemini model and broader enterprise integrations. Keypoint 2: The executives discuss various AI-powered products and services, including Gemini, Google Cloud AI, and Google DeepMind. Output: 1 Input: Keypoint 1: <kp1> Keypoint 2: <kp2> Output: Matching Score:

Figure 7: The prompt used to determine the alignment score between two saliency points.