

Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA

Minzheng Wang^{1,2*}, Longze Chen^{1,3*}, Cheng Fu⁴, Shengyi Liao⁴, Xinghua Zhang⁴
Bingli Wu⁴, Haiyang Yu⁴, Nan Xu², Lei Zhang^{1,3}, Run Luo^{1,3}, Yunshui Li^{1,3}
Min Yang^{3†}, Fei Huang⁴, Yongbin Li^{4†}

¹ School of Artificial Intelligence, University of Chinese Academy of Sciences

² MAIS, Institute of Automation, Chinese Academy of Sciences

³ Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

⁴ Alibaba Group

✉: wangminzheng2023@ia.ac.cn; lz.chen2@siat.ac.cn

Abstract

Long-context modeling capabilities have garnered widespread attention, leading to the emergence of Large Language Models (LLMs) with ultra-context windows. Meanwhile, benchmarks for evaluating long-context LLMs are gradually catching up. However, existing benchmarks employ irrelevant noise texts to artificially extend the length of test cases, diverging from the real-world scenarios of long-context applications. To bridge this gap, we propose a novel long-context benchmark, **Loong**, aligning with realistic scenarios through extended multi-document question answering (QA). Unlike typical document QA, in Loong’s test cases, each document is relevant to the final answer, ignoring any document will lead to the failure of the answer. Furthermore, Loong introduces four types of tasks with a range of context lengths: *Spotlight Locating*, *Comparison*, *Clustering*, and *Chain of Reasoning*, to facilitate a more realistic and comprehensive evaluation of long-context understanding. Extensive experiments indicate that existing long-context language models still exhibit considerable potential for enhancement. Retrieval augmented generation (RAG) achieves poor performance, demonstrating that Loong can reliably assess the model’s long-context modeling capabilities.¹

1 Introduction

Large Language Models (LLMs) have exhibited remarkable proficiency in diverse downstream applications (OpenAI, 2023). Recent works focus on scaling up the context window of LLMs (Xiong et al., 2024; Peng et al., 2024; Chen et al., 2024b), which is crucial for LLMs in handling complex

*Equal contribution.

†Corresponding author.

¹The code and benchmark are available at <https://github.com/MozerWang/Loong> and <https://github.com/AlibabaResearch/DAMO-ConvAI/tree/main/Loong>

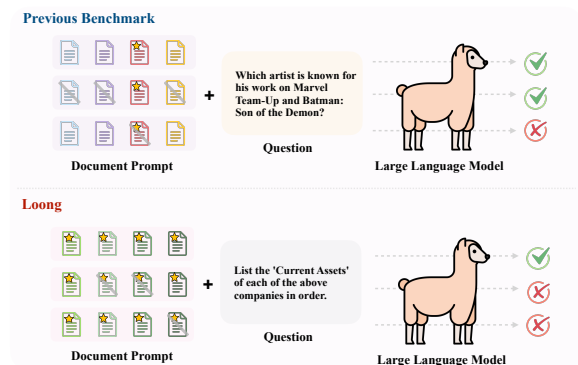


Figure 1: Previous benchmarks vs. Loong. ★ marks the existence of evidence related to the answer in that document. Compared to *centralized* distribution in previous ones, evidence in Loong are *scattered* in different parts across multi-document long contexts, necessitating that no document can be ignored for success.

tasks that require delving deeply into long texts. A few LLM (e.g. GPT-4o, Gemini-Pro) websites have been equipped with the intelligent document analysis function, allowing users to upload documents for answering queries. These demand the model leverage its long-context capability to conduct an in-depth analysis of multiple long documents. Meanwhile, retrieval-augmented generation (RAG) has emerged as a widely adopted framework that prompts LLMs with multiple relevant retrieved contents and can significantly improve model performance (Wu et al., 2024; Chen et al., 2024a).

However, there remains a lack of appropriate benchmarks for evaluating long-context understanding in real-world multi-document scenarios. Multi-document input as long-context modeling possesses extensive application scenarios of LLMs, such as analysis of financial reports over the years. Nevertheless, most existing benchmarks only place emphasis on single-document long contexts (An et al., 2024; Li et al., 2024; Kamradt, 2023) or involve multi-document question answering settings

Benchmark	Multi-doc Tasks	Broad Length Sets	Avoidance of Contamination	Realistic Scenarios	High Evidence Dispersion	Multilingual
L-Eval (An et al., 2024)	✗	✗	✗	✓	✗	✓
LongBench (Bai et al., 2024)	✓	✗	✗	✓	✗	✓
Marathon (Zhang et al., 2024a)	✓	✓	✗	✓	✗	✓
LooGLE (Li et al., 2024)	✗	✗	✗	✓	✗	✓
InfiniteBench (Zhang et al., 2024b)	✓	✗	✗	✓	✗	✓
RULER (Hsieh et al., 2024)	✓	✓	✗	✗	✗	✗
NIAH (Kamradt, 2023)	✗	✓	✓	✗	✗	✗
Loong (Ours)	✓	✓	✓	✓	✓	✓

Table 1: Characteristics of Loong, where the evidences are scattered across multi-document long contexts.

by adding distracting information to the input of existing short-context QA datasets (Hsieh et al., 2024). As shown in Figure 1, evidence supporting the answer in previous benchmarks is relatively centralized, such as being contained within a single document. Yet, such a centralized distribution of evidence may cause the model to overlook certain documents and take shortcuts to formulate an answer, simplifying the modeling of the real context. Moreover, the prevalent evaluation tasks, such as “needle in a haystack” (NIAH) (Kamradt, 2023), only scratch the surface of long-context understanding by searching from context, far from real-world demands.

We commence with “leave no document behind” and scatter the evidence across multi-document long contexts. In this context, bypassing any document will lead to an erroneous answer, which better tests the long-context modeling ability. To this end, this paper develops Loong, an innovative benchmark crafted to evaluate the long-context ability of LLMs across multiple documents in real-world scenarios. Loong typically consists of 11 documents per test instance on average, nearly all of which are selected from the year 2024, spanning three real-world scenarios in English and Chinese: (1) *Financial Reports*, (2) *Legal Cases*, and (3) *Academic Papers*. Meanwhile, Loong introduces new evaluation tasks from the perspectives of *Spotlight Locating*, *Comparison*, *Clustering*, and *Chain of Reasoning*, every test case is newly annotated. Furthermore, Loong features inputs of varying lengths (e.g., 10K-50K, 50K-100K, 100K-200K, >200K) and evaluation tasks of diverse difficulty, enabling fine-grained assessment of LLMs across different context lengths and task complexities.

We conduct extensive experiments on Loong to test the long-context modeling capabilities of several advanced LLMs. The empirical results show

that even the current most powerful LLMs still struggle with the tasks in Loong, suggesting significant room for improvement in current LLMs. Furthermore, this paper conducts in-depth analyses regarding the behavior of long-context LLMs involving RAG and the scaling law of context size. Our main contributions are summarized as follows:

- Loong primarily focuses on testing the long-context modeling ability of LLMs across multiple documents by scattering the evidence to examine the actual length of long contexts.
- Loong offers evaluation sets with varying input lengths and evaluation tasks of differing difficulty, encompassing novel task categories. All test instances are newly annotated and checked to guarantee quality.
- Extensive experiments and analyses deeply unveil the performance of Long-Context LLMs, also offering insights into improving Long-Context modeling abilities.

2 Related Work

2.1 Long-Context Language Models

With support for increasingly larger context windows, closed-source LLMs have taken the lead in the field of long-context modeling. From 128k to 1000k, GPT-4o (OpenAI, 2023), Claude3-200k (Anthropic, 2024a) and Gemini-pro1.5-1000k (Reid et al., 2024) are capable of modeling increasingly longer documents, expanding the new scenarios that LLMs can handle.

Considering the quadratic complexity of Transformer (Vaswani et al., 2017), training LLMs with extensive context windows from scratch necessitates substantial computational resources, exceeding the capabilities of the general researchers. Consequently, recent studies have explored ways to

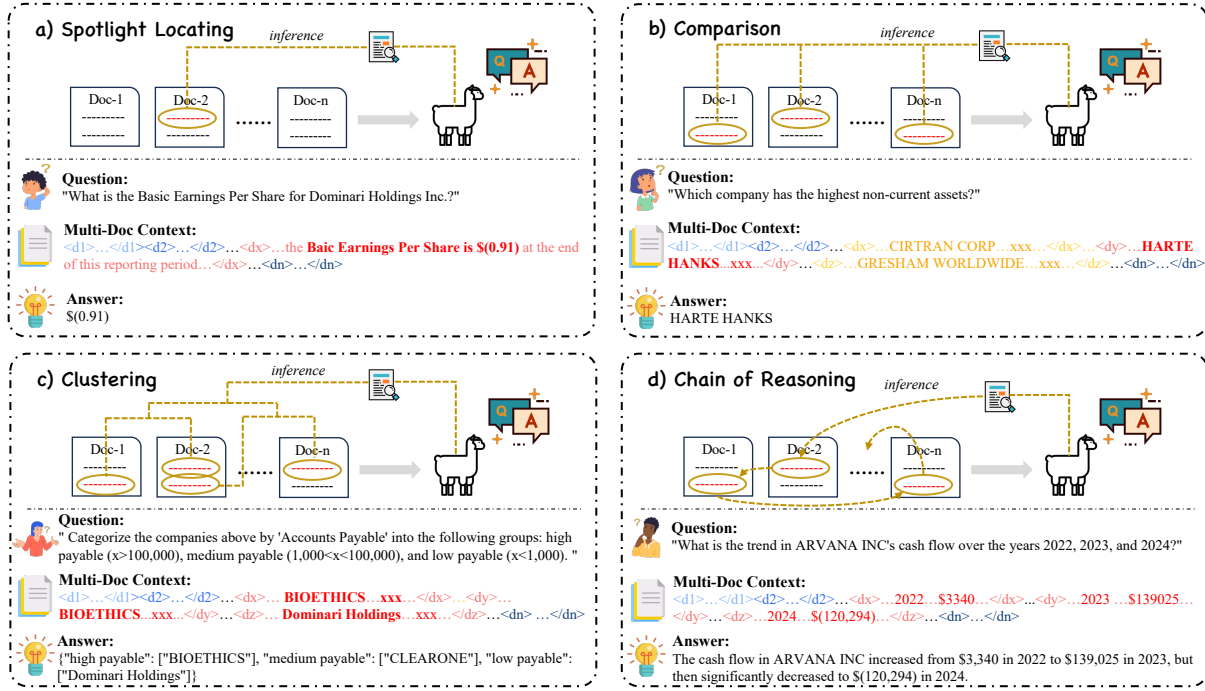


Figure 2: Showcase of four evaluation tasks in Loong (`<di>...</di>` marks the content of the i -th document). a) *Spotlight Locating*: Locate the evidence. b) *Comparison*: Locate and compare the evidence. c) *Clustering*: Locate and cluster the evidence into groups. d) *Chain of Reasoning*: Locate and reasoning along a logical chain.

expand the context length of these models during the fine-tuning stage. For example, PI (Chen et al., 2023), NTK-aware (bloc97, 2023), YaRN (Peng et al., 2024), Giraffe (Pal et al., 2023), Code LLaMA (Roziere et al., 2023), and PoSE (Zhu et al., 2024) adapts position embedding based on the rotary position encoding (RoPE) (Su et al., 2024), with only a few fine-tuning steps, the context length can be efficiently extended.

Another strong baseline for long-context modeling is the sliding window method. Various sliding window-based variants such as ALibi (Press et al., 2022), xPos (Sun et al., 2023), PCW (Ratner et al., 2023), LM-Infinity (Han et al., 2024), StreamingLLM (Xiao et al., 2023) are used to achieve efficient context scaling. Yet they diverge from the global perception characteristic of the Transformer, failing to exploit the entire context.

2.2 Long-Context Benchmarks

Long-context modeling methods are rapidly evolving, yet the quality of existing benchmarks does not align with this progress. Synthetic task such as Needle-in-a-Haystack (NIAH) (Kamradt, 2023) and Counting stars (Song et al., 2024) are initially utilized for evaluating long-context language models (LCLMs) due to their lower construction costs, but they are indicative of only a surface form of

long-context understanding.

Longbench (Bai et al., 2024), LooGLE (Li et al., 2024) and Marathon (Zhang et al., 2024a) are earlier benchmarks for comprehensive assessment of long context. However, the average length for most tasks is between 5k and 25k, far less than the window size of LCLMs. L-Eval (An et al., 2024), BAMBOO (Dong et al., 2024), CLongEval (Qiu et al., 2024) and InfiniteBench (Zhang et al., 2024b) contain sufficiently long evaluation data, and the wide variety of tasks makes the assessment more comprehensive. RULER (Hsieh et al., 2024) creates a comprehensive testing method with flexibly adjustable length and difficulty, yet they only add distracting information to the input of existing short-context QA datasets.

While these long-context benchmarks have their advantages, there still lack a benchmark that is sufficiently long, free from data contamination (Golchin and Surdeanu, 2023), and fully aligned with the real-world multi-document question answering scenario. We conduct a detailed comparison with existing works in Table 1.

2.3 Retrieval Augmented Language Models

Leveraging long documents as external knowledge, Retrieval Augmented Language Models (RALMs) have achieved comparable or even better perfor-

mance than LCLMs fine-tuned for specific tasks with long documents. In previous study, RALMs could directly utilize the content retrieved during the inference phase. REPLUG (Shi et al., 2024) treats the language model as a black box and the retrieval component as an adjustable plug-and-play module. RETRO (Borgeaud et al., 2022) use a chunked cross-attention module to incorporate the retrieved text. Additionally, Xu et al. (2024) explored whether RALMs or LCLMs are more suitable for long-context tasks under a larger parameter setting. However, there is currently a lack of analysis on what tasks RALMs and LCLMs each excel at, thus making it difficult to determine which type a black box model belongs to.

3 Loong: A Long-Context Benchmark

3.1 Overview

The Loong benchmark comprises tasks across four categories: *Spotlight Locating*, *Comparison*, *Clustering*, and *Chain of reasoning*. To align with realistic scenarios, we collect documents from three domains: *financial reports*, *academic papers*, and *legal cases*. Furthermore, all tasks are presented in the question-answering format, which are all newly annotated by GPT-4o and humans. Totally, Loong includes 1600 test instances in both Chinese and English, featuring four sets with different intervals of context size: Set1 (10-50K), Set2 (50-100K), Set3 (100-200K) and Set4 (200-250K). We use tiktoken² tokenizer to tokenize the input and report the number of tokens. Table 2 and Appendix C show the details of data statistics. Appendix E presents a comparison of evidence distribution between Loong and LongBench. The following sections will provide a detailed description of the evaluation task and benchmark construction.

3.2 Evaluation Task

Based on various multi-document semantic relationships and LLMs’ handling of multi-document input, we propose new task categories for multi-document long-context modeling and closer alignment with real-world scenarios. Figure 2 illustrates the evaluation tasks of the Loong benchmark. Appendix B shows the detailed test case and prompt of each task.

²<https://platform.openai.com/tokenizer>

Category	Avg Token	Language	#Test Instance
<i>Task</i>			
Spotlight Locating	119.3K	EN, ZH	250
Comparison	110.6K	EN, ZH	300
Clustering	109.8K	EN, ZH	641
Chain of Reasoning	103.9K	EN, ZH	409
<i>Sub Task</i>			
Sequential Enumeration	103K	EN, ZH	87
Extremum Acquisition	115K	EN, ZH	143
Range Awareness	111K	EN, ZH	70
Report Integration	117K	EN, ZH	250
Citation&Reference	105K	EN	270
Case Classification	106K	ZH	121
Temporal Analysis	112K	EN, ZH	100
Citation Chain	91K	EN	130
Link the Links	117K	ZH	113
Solitaire	94K	ZH	66
<i>Domain</i>			
Financial Reports	117.5K	EN, ZH	700
Legal Cases	107.2K	ZH	500
Academic Papers	100.9K	EN	400
<i>Length Set</i>			
Set1 (10-50K)	37.8K	EN, ZH	323
Set2 (50-100K)	75.6K	EN, ZH	564
Set3 (100-200K)	138.9K	EN, ZH	481
Set4 (200-250K)	233.9K	EN, ZH	232

Table 2: Data statistics of Loong benchmark.

3.2.1 Spotlight Locating

The spotlight locating task is designed to assess the model’s capability for knowledge localization, which constitutes the foundation ability of long-context processing. In this task, the evidence is contained in only one of multiple documents, which is the atomic setting of the key information locating. The spotlight locating task is aimed at examining the LLMs’ ability to search for evidence within one document from multiple ones. Other documents, which are in the same domain and have similar semantics as the document but are unrelated to the question, will serve as noise texts. The upper left of Figure 2 provides an example of the spotlight locating task.

3.2.2 Comparison

The comparison task is primarily aimed at evaluating the model’s ability to compare multi-source information with long contexts. In this event, the evidence supporting the answer are distributed across multiple documents, testing the LLMs’ ability to locate dispersed evidence and to correlate and compare them.

Comparison task includes three sub-tasks: 1) *Sequential Enumeration*: Based on the concrete numerical value of a specific attribute, it requires the model to list all specific values corresponding

to that attribute across multiple documents in a given order. 2) *Extremum Acquisition*: It requires the model to deduce the extremum of all values corresponding to certain attributes in multiple documents. 3) *Range Awareness*: Given a specific numerical or conceptual range, the model should output all objects within multiple documents that meet the condition. The upper right of Figure 2 gives an example of comparison task.

3.2.3 Clustering

The clustering task entails an assessment of the model’s ability to aggregate multi-source information from multi-document long contexts. The task requires the LLM to perform two main functions: first, to extract relevant data from multiple documents, and second, to integrate this information based on specific criteria. Ultimately, the LLM must cluster related evidence spread across these documents into coherent groups.

The clustering task encompasses three sub-tasks: 1) *Report Integration*: This sub-task requires the model to group the evidence existing in the provided financial reports into corresponding sets based on textual or numerical criteria. 2) *Citation&Reference*: For a given paper, the model is tasked with identifying its citations and references from the candidate papers. 3) *Case Classification*: Given the causes of several legal cases, the model is required to accurately categorize judgment documents. The bottom left of Figure 2 depicts an example of the clustering task.

3.2.4 Chain of Reasoning

The chain of reasoning task requires the model to engage in multi-document reasoning along a logical pathway. This task evaluates the model’s proficiency in logical reasoning, which requires LLMs to locate the corresponding evidence within multiple documents and model the logical relationships among them for deducing the answer.

The chain of reasoning task contains four sub-tasks: 1) *Temporal Analysis*: This task requires the model to analyze the changes or trends of a particular attribute based on the temporal relationship, such as taking into account the financial reports of a certain company over consecutive years or multiple quarters. 2) *Citation Chain*: This task requires the model to accurately understand each paper’s content and its interconnections, ultimately inferring the linear citation relationships among them. 3) *Link the Links*: This task involves presenting fact

descriptions and trial results from different judgment documents separately. The model is tasked with accurately pairing each fact description with its corresponding trial result. 4) *Solitaire*: This task first requires the model to match causes of action with judgment documents correctly, and then to sequentially infer multiple judgment documents based on the given sequence of causes of action. The bottom right of Figure 2 gives an example of the chain of reasoning task.

3.3 Benchmark Construction

3.3.1 Data Collection

We established six criteria for the manual collection of the required English and Chinese documents: (1) *Timeliness*: The majority of the documents are the latest ones from the year 2024; (2) *Accessibility*: The data is publicly available and permitted for download and collection; (3) *Appropriate Length*: Collecting longer documents as much as possible and ensure they fit within the four designated length sets; (4) *Parseability*: Chosen documents are easy to process and parse, facilitating conversion into natural language text; (5) *Categorizability*: Documents can be manually sorted based on certain attributes, such as case type, research theme, or company category, allowing for organized archival; (6) *Authoritativeness*: All documents are collected from scratch from official websites (e.g. China Judge Online, U.S. SEC, cniinf, Arxiv, Semantic Scholar), ensuring the quality and authority of the documents. The detailed URL can be seen in Appendix H.

Specifically, regarding financial reports, we primarily collect the latest quarterly and annual reports for the year 2024, totaling 574 documents. For legal documents, our collection consists exclusively of cases adjudicated by the higher and intermediate courts in 2024, amounting to 629 documents. As for academic papers, our focus is on procuring the latest articles from arXiv in 2024, with a total of 764 papers. Additionally, to meet the requirements of the chain of reasoning task, we gather a small portion of financial reports and academic papers from before 2024. Upon the collection of documents, we first parse these documents, converting them uniformly into TXT format. Subsequently, we carry out further data cleansing, removing any portions that contain personal information.

Model	Spotlight Locating		Comparison		Clustering		Chain of Reasoning		Overall	
GPT-4o (128K)	73.95	0.62	50.50	0.28	44.29	0.09	57.95	0.28	53.47	0.26
Gemini-1.5-pro (1000K)	75.02	0.56	49.94	0.27	44.10	0.09	64.97	0.37	55.37	0.27
Claude3.5-Sonnet (200K)	58.45	0.49	54.21	0.35	45.77	0.07	43.92	0.25	48.85	0.23
Qwen2-72B-Instruct (128K)	54.17	0.36	42.38	0.20	36.71	0.04	47.76	0.18	43.29	0.15
Claude3-Haiku (200K)	68.68	0.59	42.10	0.21	35.04	0.02	47.59	0.17	44.88	0.19
Kimi-Chat (200k)	60.98	0.50	34.74	0.13	28.76	0.04	38.52	0.15	37.49	0.16
GLM4-9B-Chat (1000K)	57.35	0.47	40.38	0.20	28.52	0.02	39.94	0.16	38.31	0.16

Table 3: Overall results on four evaluation tasks. For each task, the indicator on the left represents the *Avg Scores (0~100)*, while the right one represents the *Perfect Rate (0~1)*.

3.3.2 Annotation Process

Compared to annotating short texts, annotating long texts is more challenging. To address this issue, we designed innovative annotation workflows to reduce the cost of annotation while ensuring quality.

For *financial reports*, we compress the information contained within the long context, breaking down the annotation process into numerous simple tasks. We initially manually identify hundreds of key attributes that cover the important information in the long context. Subsequently, we employ GPT-4o to execute the relatively simple task of information extraction, pulling the values corresponding to these key attributes. After obtaining the key attributes and their corresponding values, we can proceed to annotate only the compressed information, eliminating the need to refer back to the original lengthy texts. For *legal cases*, we follow the classification provided by China Judge Online, manually downloading judgment documents sorted by different causes of action and case types. Additionally, we use a rule-based method to segment each judgment document into its factual statement and verdict sections. For *academic papers*, we leverage the Semantic Scholar website’s API to access the target paper’s citations and references. Moreover, by utilizing the bbl files of each arXiv paper, we write scripts to recursively collect articles that meet the requirements of the linear citation chain task.

During the question-and-answer annotation phase, we adopt two approaches: (1) *Template-based*: We design question types and templates, and based on pre-classified documents, we construct Q&A pairs using rules. (2) *Free annotation*: Referring to the compressed information of multiple documents, we design prompts with four different task descriptions. We employ GPT-4o to generate Q&A pairs for each task.

3.3.3 Quality Control

Throughout the annotation process, we employ several methods to ensure accuracy: (1) *Evidence Recall*: By designing prompts that not only prompt GPT-4o to generate labels but also to recall evidence supporting the labels from the text, significantly enhancing the accuracy in practical applications. (2) *Self-Check*: GPT-4o reviews the original text to re-evaluate and correct any mistakes in the generated labels. (3) *Manual Check*: We manually review and confirm the quality of annotations, eliminating any unreasonable or low-quality questions. Additionally, we also take into account the distribution and number of different length sets, sub-tasks, and language. From a pool of 2,814 entries, we conduct a secondary selection process, ultimately choosing 1,600 entries for our final benchmark.

4 Experiments

4.1 Experimental Setup

Models We evaluate six advanced long-context LLMs, with their context window sizes ranging from 128K to 1000K, including API-based LLMs: GPT-4o-128K (OpenAI, 2023), Gemini-1.5-pro-1000K (Reid et al., 2024), Claude3.5-Sonnet-200K (Anthropic, 2024b), Claude3-Haiku-200K (Anthropic, 2024a), Kimi-Chat-200K³ and Open-sourced LLMs: Qwen2-72B-Instruct-128K (Bai et al., 2023), GLM4-9B-Chat-1000K (Du et al., 2022).

Evaluation Metric In the long-context question-answering scenarios, traditional evaluation metrics F1 and Rouge-L may lead to inaccurate responses. Recent research (Zhang et al., 2023; Liu et al., 2024; Wang et al., 2024) indicates that the GPT-4 (OpenAI, 2023) evaluator demonstrates high consistency with human evaluations, making it a reasonably reliable annotator. Building on these con-

³<https://kimi.moonshot.cn/>

Model	Spotlight Locating		Comparison		Clustering		Chain of Reasoning		Overall	
Set1 (10K-50K)										
GPT-4o (128K)	85.67	0.81	64.27	0.33	57.01	0.24	81.58	0.55	70.40	0.44
Gemini-1.5-pro (1000K)	75.00	0.60	54.88	0.28	56.15	0.23	70.64	0.37	63.36	0.34
Claude3.5-Sonnet (200K)	60.85	0.55	69.07	0.47	58.63	0.13	68.57	0.50	63.69	0.37
Qwen2-72B-Instruct (128K)	68.49	0.55	60.60	0.37	47.08	0.08	70.39	0.36	60.11	0.29
Claude3-Haiku (200K)	60.94	0.55	59.97	0.40	45.53	0.04	66.85	0.34	57.14	0.28
Kimi-Chat (200k)	81.11	0.74	46.70	0.20	47.84	0.07	53.77	0.17	55.02	0.24
GLM4-9B-Chat (1000K)	63.11	0.53	54.10	0.27	39.50	0.08	56.32	0.28	51.43	0.25
Set2 (50K-100K)										
GPT-4o (128K)	86.76	0.72	59.81	0.40	47.83	0.11	62.09	0.34	58.38	0.29
Gemini-1.5-pro (1000K)	76.50	0.57	54.51	0.34	44.58	0.09	64.87	0.34	55.56	0.26
Claude3.5-Sonnet (200K)	63.83	0.53	58.90	0.39	50.96	0.10	46.09	0.26	52.73	0.24
Qwen2-72B-Instruct (128K)	64.53	0.43	42.60	0.21	38.52	0.05	51.18	0.20	45.71	0.17
Claude3-Haiku (200K)	73.71	0.66	41.90	0.22	36.18	0.02	50.20	0.15	45.45	0.17
Kimi-Chat (200k)	72.82	0.52	46.77	0.21	33.46	0.06	40.51	0.15	42.40	0.16
GLM4-9B-Chat (1000K)	65.04	0.54	41.80	0.23	30.72	0.02	42.34	0.17	40.19	0.17
Set3 (100K-200K)										
GPT-4o (128K)	74.84	0.65	42.40	0.21	38.70	0.04	45.06	0.09	46.95	0.19
Gemini-1.5-pro (1000K)	81.25	0.56	44.66	0.20	39.90	0.05	58.38	0.36	52.05	0.24
Claude3.5-Sonnet (200K)	65.36	0.56	50.32	0.34	37.79	0.03	25.95	0.11	42.06	0.19
Qwen2-72B-Instruct (128K)	46.99	0.27	37.06	0.13	31.50	0.02	35.01	0.07	35.94	0.09
Claude3-Haiku (200K)	77.81	0.67	37.07	0.17	30.94	0.01	36.87	0.12	41.41	0.18
Kimi-Chat (200k)	62.13	0.54	24.20	0.05	21.98	0.01	31.02	0.14	31.37	0.14
GLM4-9B-Chat (1000K)	69.19	0.56	37.99	0.18	26.63	0.01	32.30	0.09	37.36	0.16
Set4 (200K-250K)										
GPT-4o (128K)	36.79	0.19	23.97	0.08	30.40	0.00	32.89	0.07	31.11	0.07
Gemini-1.5-pro (1000K)	62.23	0.49	43.08	0.20	36.48	0.00	68.51	0.49	50.70	0.25
Claude3.5-Sonnet (200K)	36.91	0.24	28.82	0.05	28.68	0.00	28.77	0.08	30.51	0.08
Qwen2-72B-Instruct (128K)	33.18	0.16	26.59	0.08	29.84	0.01	25.81	0.04	28.92	0.06
Claude3-Haiku (200K)	53.26	0.40	27.00	0.03	25.36	0.00	28.11	0.05	32.15	0.10
Kimi-Chat (200k)	20.17	0.12	9.17	0.00	5.65	0.00	22.61	0.11	13.50	0.05
GLM4-9B-Chat (1000K)	15.67	0.12	21.33	0.05	12.35	0.00	21.04	0.05	16.84	0.05

Table 4: The performance of LLMs on four evaluation tasks with different length sets. For each task, the indicator on the left represents the *Avg Scores* (0~100), while the right one represents the *Perfect Rate* (0~1).

siderations, we prompt GPT-4 as a judge to evaluate the model’s output based on the golden answer and the question’s requirements from three aspects: *Accuracy*, *Hallucinations*, and *Completeness*, scoring from 0 to 100. For a detailed prompt, please refer to the Appendix A. We also design two indicators: (1) *Avg Scores*: the average value of scores given by GPT-4 for all questions; (2) *Perfect Rate*: the proportion of cases scoring 100 out of the total cases. The latter is a more stringent evaluation metric compared to the former.

Prompt Templates For different sub-tasks, we require the model to follow the given instructions and output the answer according to the specific prompts shown in Appendix B.

Input Truncation Due to input length limits, we assess whether adding a document would exceed the model’s processing length when concatenating multiple documents. If appending the document

would surpass the model’s capacity, we discard it from the concatenation process. The evaluation and selection process continues until we have reviewed all documents that need concatenation.

Implement Details We set ‘temperature = 0’ to eliminate randomness and keep other hyperparameters default. For API-Based LLMs, we directly utilize the official API for testing. Since the Kimi-Chat-200k currently does not provide an interface, we manually input content on the web. As for open-source models, we conduct experiments on a server with 8×A100 80GB.

4.2 Main Results

We assess seven advanced LLMs on the Loong benchmark. The main results are shown in Table 3 and Table 4. We can see that Gemini-1.5-pro shows the best overall performance, especially excelling in the processing of ultra-long context within Set3


and Set4. Its comprehensive score reached 55.37 with the perfect rate of 27%, followed by GPT-4o. Besides, the long-context modeling capacity of open-source models still falls short when compared to that of the most powerful closed-source models in the Loong. Additionally, larger-parameter models outperform their smaller counterparts within the same window size, indicating the advantages of scaling up model sizes for improved long-context modeling. The overall assessment results highlight that even the most advanced long-context LLMs currently fail to achieve passing marks, particularly in terms of the perfect rate. This suggests that there exists significant room for improvement in the long-context modeling capabilities of LLMs. We also compare the current results with other current Long-Context benchmarks, which can be seen in Appendix F.

4.3 Task Analysis

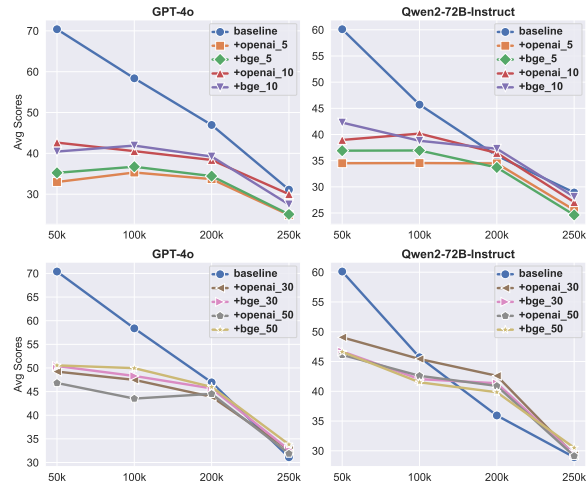
Analyzing performance across different tasks, models exhibit their best performance in the spotlight locating task. This can be attributed to the task’s relative simplicity, which tests the foundational capabilities of long-context modeling. Moreover, the evidence is only distributed within a single document, making it easier to locate and less prone to confusion. In contrast, due to the requirements of multi-source information inference, the comparison and cluster tasks present greater challenges, leading to model underperformance. These tasks necessitate not only the collection of evidence across documents but also involve complex reasoning processes such as matching, contrasting, and classification. Thus, they more rigorously test the higher-order capabilities of long-context modeling, revealing significant gaps in the current models’ abilities. Regarding the chain of reasoning task, models perform well within Set1. However, as the context length increases, their performance drastically declines. This suggests that within the scope of long-context modeling capabilities, LLMs possess adequate skills in temporal analysis, logical sequencing, and linking multiple concepts. Nevertheless, an overflow in context length leads to the loss of key evidence, severely impacting the accuracy of chain reasoning tasks.

4.4 Scaling Law of Context Window

It’s observed that the general performance of all models deteriorates with the increase in context size. As observed from Table 4, it is apparent



(a) The overall results on all length sets.



(b) The detailed results on different length sets. The baseline means the setting without RAG.

Figure 3: The results on all tasks after adding RAG module. We only represents the *Avg Scores (0-100)*.

that for the same task, models perform well within small length sets but exhibit a notable performance decline as the length increases. This indicates that the models possess a certain capability to process the task, yet their performance is constrained by the context window. Moreover, despite being trained on 128K data, the GPT-4o and Qwen2-72B-Instruct begin to show performance degradation within the 50-100K interval, revealing that their actual capability boundary is significantly lower than the claimed window size. This suggests the presence of an ineffective zone within the claimed window. There exists a Scaling Law for model window sizes: to truly equip an LLM with the ability to handle 128K long texts, it should be trained on data exceeding 128K, meaning the training length should be greater than the actual processable length. Among numerous models, only the Gemini is less affected by changes in context length, which was training on the ultra-long context of 1000K. To ensure your model genuinely possesses the desired context window size, train it on longer data!

4.5 RAG or Not

We have also incorporated the Embedding RAG module into the GPT-4o and Qwen2-72B-Instruct to explore whether RAG can enhance the model’s performance on Loong. For the Embedding choice, we employ two distinct models: the OpenAI Embedding model⁴ and the BGE Embedding model⁵. Besides, we set the top-k value of 5, 10, 30, and 50 for each model respectively, and the chunk size is 1024. The result is shown in Figure 3 and the details can be seen in Appendix D.

Benchmark Analysis It is evident that the inclusion of RAG does not enhance the model’s overall performance on the Loong, and there is a noticeable decline in assessment. This is because the evidence in the Loong is distributed relatively evenly across multiple documents, requiring a comprehensive understanding of long texts by the model. RAG, being more limited, only shows some effectiveness in the task with sparse evidence, such as spotlight locating. However, RAG’s negative impact is significant for tasks requiring a high level of comprehensiveness. Integrating RAG does not enhance the performance, indicating that Loong focuses on evaluating the model’s complex reasoning and comprehensive analysis capabilities for long contexts, thereby effectively assessing the LLM’s long-context modeling ability.

Model Analysis Comparing the performance between GPT-4o and Qwen2-72B-Instruct, it is evident that the powerful long-context LLM significantly outperforms RAG. On the other hand, the performance of RAG is closer to the original performance when used with a weak context model. This is because a strong long-context LLM can fully exploit the complete information flow of long contexts, capturing complex dependencies and semantic information. However, RAG causes context fragmentation and information loss, impairing the model’s understanding and reasoning capabilities, thereby preventing the full utilization of its inherent modeling advantages. Consequently, a strong long-text modeling capability is not suitable for enhancement through RAG. Conversely, a weak model with poor long-context modeling capability cannot effectively capture information, and RAG cannot compensate for this deficiency.

Length Analysis Within the context window size

⁴<https://huggingface.co/Xenova/text-embedding-ada-002>

⁵<https://huggingface.co/BAAI/bge-m3>

that the model can handle, RAG does not offer an advantage. However, for ultra-long context sets, a high top-k setting of RAG can produce certain effects. This is because, in short context sets, the model’s inherent modeling capability can effectively handle the entire text length without losing information. The introduction of RAG, conversely, may result in the loss of certain evidence, leading to information gaps. In ultra-long context collections, RAG can effectively compress information, recalling evidence that the LLM could not access due to length truncation, thereby enhancing the model’s performance on Loong.

We also provide analytical experiment to show that RAG does not cover all evidence, which can be seen in Appendix G. Relying on RAG cannot resolve all the problems associated with long-text modeling. To genuinely improve the long-context modeling capability, stronger training methods and effective training on longer texts are required, rather than merely integrating the RAG module.

5 Conclusion

In this study, we propose Loong, a question-answering format benchmark designed to evaluate long-context comprehension in real-world multi-document scenarios. We analyze six advanced language models (LLMs), considering variations in their parameter sizes and context windows, including GPT-4o and Gemini-Pro1.5. Notably, even the most powerful long-context LLMs fail to achieve satisfactory performance. Furthermore, we conduct in-depth analyses to enhance long-context modeling capabilities by comparing the RAG approach and the scaling laws related to context size.

Limitations

Here we list some of the limitations that are not considered when designing Loong: (1) Limited Domains. The purpose of Loong is to evaluate the long-context understanding capabilities in real-world multi-document scenarios. However, a sea of multi-document domains exists in the real world. Considering annotation costs and model evaluation efficiency, we only cover the most representative parts of them: financial, legal, and academic. (2) High Annotation Cost. To enhance the reliability of Loong in assessing the LLM’s long-context understanding capabilities, we recruited a group of experts for each of the three domains to proofread the data, and they are proficient in both English and

Chinese. They need to understand the question and search for relevant evidence in multiple documents with an average length of up to 110k to judge the consistency between the question and the answer, which requires a significant amount of time and effort.

Acknowledgments

Min Yang was supported by National Key Research and Development Program of China (2022YFF0902100), National Natural Science Foundation of China (Grant No. 62376262), the Natural Science Foundation of Guangdong Province of China (2024A1515030166), Shenzhen Science and Technology Innovation Program (KQTD20190929172835662), Shenzhen Basic Research Foundation (JCYJ20210324115614039).

References

- Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2024. [L-eval: Instituting standardized evaluation for long context language models](#). In *Proceedings of ACL*, pages 14388–14411.
- AI Anthropic. 2024a. [The claude 3 model family: Opus, sonnet, haiku](#). *Claude-3 Model Card*.
- AI Anthropic. 2024b. [Claude 3.5 sonnet model card addendum](#). *Claude-3.5 Model Card*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. [Qwen technical report](#). *arXiv preprint arXiv:2309.16609*.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [LongBench: A bilingual, multi-task benchmark for long context understanding](#). In *Proceedings of ACL*, pages 3119–3137.
- bloc97. 2023. [Ntk-aware scaled rope](https://www.reddit.com/r/LocalLLaMA/comments/141z7j5/ntkaware_scaled_rope_allows_llama_models_to_have/). https://www.reddit.com/r/LocalLLaMA/comments/141z7j5/ntkaware_scaled_rope_allows_llama_models_to_have/.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. [Improving language models by retrieving from trillions of tokens](#). In *Proceedings of ICML*, pages 2206–2240.
- Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024a. [Benchmarking large language models in retrieval-augmented generation](#). In *Proceedings of AAAI*, pages 17754–17762.
- Longze Chen, Ziqiang Liu, Wanwei He, Yinhe Zheng, Hao Sun, Yunshui Li, Run Luo, and Min Yang. 2024b. [Long context is not long at all: A prospector of long-dependency data for large language models](#). In *Proceedings of ACL*, pages 8222–8234.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. [Extending context window of large language models via positional interpolation](#). *arXiv preprint arXiv:2306.15595*.
- Zican Dong, Tianyi Tang, Junyi Li, Wayne Xin Zhao, and Ji-Rong Wen. 2024. [BAMBOO: A comprehensive benchmark for evaluating long text modeling capacities of large language models](#). In *Proceedings of LREC-COLING*, pages 2086–2099.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. [Glm: General language model pretraining with autoregressive blank infilling](#). In *Proceedings of ACL*, pages 320–335.
- Shahriar Golchin and Mihai Surdeanu. 2023. [Time travel in llms: Tracing data contamination in large language models](#). *arXiv preprint arXiv:2308.08493*.
- Chi Han, Qifan Wang, Hao Peng, Wenhan Xiong, Yu Chen, Heng Ji, and Sinong Wang. 2024. [LM-infinite: Zero-shot extreme length generalization for large language models](#). In *Proceedings of NAACL*, pages 3991–4008.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Krizan, Shantanu Acharya, Dima Rekesch, Fei Jia, and Boris Ginsburg. 2024. [RULER: What’s the real context size of your long-context language models?](#) In *Proceedings of COLM*.
- Greg Kamradt. 2023. [Needle in a haystack - pressure testing llms](#). https://github.com/gkamradt/LLMTest_NeedleInAHaystack.
- Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2024. [One thousand and one pairs: A "novel" challenge for long-context language models](#). *arXiv preprint arXiv: 2406.16264*.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024. [LooGLE: Can long-context language models understand long contexts?](#) In *Proceedings of ACL*, pages 16304–16333.
- Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun, and Qi Zhang. 2024. [Calibrating LLM-based evaluator](#). In *Proceedings of LREC-COLING*, pages 2638–2656.
- OpenAI. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Arka Pal, Deep Karkhanis, Manley Roberts, Samuel Dooley, Arvind Sundararajan, and Siddhartha Naidu. 2023. [Giraffe: Adventures in expanding context lengths in llms](#). *arXiv preprint arXiv:2308.10882*.

- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2024. [YaRN: Efficient context window extension of large language models](#). In *Proceedings of ICLR*.
- Ofir Press, Noah Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *Proceedings of ICLR*.
- Zexuan Qiu, Jingjing Li, Shijue Huang, Wanjun Zhong, and Irwin King. 2024. [Clongeval: A chinese benchmark for evaluating long-context large language models](#). *arXiv preprint arXiv:2403.03514*.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [Parallel context windows for large language models](#). In *Proceedings of ACL*, pages 6383–6402.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *arXiv preprint arXiv:2403.05530*.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. [Code llama: Open foundation models for code](#). *arXiv preprint arXiv:2308.12950*.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [REPLUG: Retrieval-augmented black-box language models](#). In *Proceedings of NAACL*, pages 8371–8384.
- Mingyang Song, Mao Zheng, and Xuan Luo. 2024. [Counting-stars: A simple, efficient, and reasonable strategy for evaluating long-context large language models](#). *arXiv preprint arXiv:2403.11802*.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*, 568:127063.
- Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shao-han Huang, Alon Benhaim, Vishrav Chaudhary, Xia Song, and Furu Wei. 2023. [A length-extrapolatable transformer](#). In *Proceedings of ACL*, pages 14590–14604.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of NeurIPS*, page 30.
- Cunxiang Wang, Sirui Cheng, Qipeng Guo, Yuanhao Yue, Bowen Ding, Zhikun Xu, Yidong Wang, Xi-angkun Hu, Zheng Zhang, and Yue Zhang. 2024. [Evaluating open-qa evaluation](#). In *Proceedings of NeurIPS*, page 36.
- Kevin Wu, Eric Wu, and James Zou. 2024. [How faithful are rag models? quantifying the tug-of-war between rag and llms’ internal prior](#). *arXiv preprint arXiv:2404.10198*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. [Efficient streaming language models with attention sinks](#). *arXiv preprint arXiv:2309.17453*.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajjwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, Madian Khabisa, Han Fang, Yashar Mehdad, Sharan Narang, Kshitiz Malik, Angela Fan, Shruti Bhosale, Sergey Edunov, Mike Lewis, Sinong Wang, and Hao Ma. 2024. [Effective long-context scaling of foundation models](#). In *Proceedings of NAACL*, pages 4643–4663.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2024. [Retrieval meets long context large language models](#). In *Proceedings of ICLR*.
- Lei Zhang, Yunshui Li, Ziqiang Liu, Jiayi Yang, Junhao Liu, Longze Chen, Run Luo, and Min Yang. 2024a. [Marathon: A race through the realm of long context with large language models](#). In *Proceedings of ACL*, pages 5201–5217.
- Xinghua Zhang, Bowen Yu, Haiyang Yu, Yangyu Lv, Tingwen Liu, Fei Huang, Hongbo Xu, and Yongbin Li. 2023. [Wider and deeper llm networks are fairer llm evaluators](#). *arXiv preprint arXiv:2308.01862*.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024b. [∞Bench: Extending long context evaluation beyond 100K tokens](#). In *Proceedings of ACL*, pages 15262–15277.
- Dawei Zhu, Nan Yang, Liang Wang, Yifan Song, Wenhao Wu, Furu Wei, and Sujian Li. 2024. [PoSE: Efficient context window extension of LLMs via positional skip-wise training](#). In *Proceedings of ICLR*.

A GPT4-as-the-Judge Prompt

In Loong, GPT4 is used as a Judger to evaluate the correctness of the model-generated content, with the prompt used shown in the following. With this evaluation method, we expect the Judger model to output a percentage score along with its corresponding explanation.

[Gold Answer]

<answer>

[The Start of Assistant's Predicted Answer]

<LLM's response>

[The End of Assistant's Predicted Answer]

[System]

We would like to request your feedback on the performance of the AI assistant in response to the user question displayed above according to the gold answer. Please use the following listed aspects and their descriptions as evaluation criteria:

- Accuracy and Hallucinations: The assistant's answer is semantically consistent with the gold answer; The numerical value and order need to be accurate, and there should be no hallucinations.
 - Completeness: Referring to the reference answers, the assistant's answer should contain all the key points needed to answer the user's question; further elaboration on these key points can be omitted.
- Please rate whether this answer is suitable for the question. Please note that the gold answer can be considered as a correct answer to the question.

The assistant receives an overall score on a scale of 1 to 100, where a higher score indicates better overall performance. Please note that if the assistant's answer and the gold answer fully meet the above criteria, its overall rating should be the full marks (100). Please first provide a comprehensive explanation of your evaluation, avoiding any potential bias. Then, output a line indicating the score of the Assistant.

PLEASE OUTPUT WITH THE FOLLOWING FORMAT, WHERE THE SCORE IS A SCALE OF 1 TO 100 BY STRICTLY FOLLOWING THIS FORMAT: "[[score]]", FOR EXAMPLE "Rating: [[100]]":

<Start Output>

Evaluation evidence: your evaluation explanation here, no more than 100 words Rating: [[score]]

<End Output>

Now, start your evaluation:

B Test Case

To facilitate understanding of Loong's data examples, we present examples of 11 sub-tasks in the following, showing the format we input to the model as well as the prompts we used.

B.1 Spotlight Locating

<Multi_Documents>

Prompt: Please answer the following questions based only on the judgment documents you have seen above. You only need to give the titles of the judgment documents that meet the requirements.

Question: Among the above judgment documents, which one is the case of 'crime of endangering public security'?

Answer: Judgment Document 5

B.2 Sequential Enumeration

<Multi_Documents>

Prompt: We kindly ask you to review the financial statements of the companies provided above and answer the following questions based solely on the information you have seen. If the question involves content not found in the financial statements, you may ignore this part and only answer the other parts.

Question: Please list the Cash and Cash Equivalents of each of the aforementioned companies in ascending order?

Answer: \$ 1,273 thousand, \$ 1,360 thousand, \$ 9,364 thousand

B.3 Extremum Acquisition

<Multi_Documents>

Prompt: We kindly ask you to review the financial statements of the companies provided above and answer the following questions based solely on the information you have seen. If the question involves content not found in the financial statements, you may ignore this part and only answer the other parts.

Question: Which company has the highest 'Total Non-current Assets'?

Answer: BLUE DOLPHIN ENERGY CO with \$56,787,000

B.4 Range Awareness

<Multi_Documents>

Prompt: We kindly ask you to review the financial statements of the companies provided above and answer the following questions

based solely on the information you have seen. If the question involves content not found in the financial statements, you may ignore this part and only answer the other parts.

Question: How many companies have 'Total Shares Outstanding' exceeding 10,000,000 shares?

Answer: 4 companies

B.5 Report Integration

<Multi_Documents>

Prompt: We kindly ask you to review the financial statements of the companies provided above and answer the following questions based solely on the information you have seen. If the question involves content not found in the financial statements, you may ignore this part and only answer the other parts.

Question: Please categorize the companies listed above by 'Total Shares Outstanding' into the following groups: below 10,000,000 shares and 10,000,000 shares or more. Place companies into the same collection for the same category and into different collections for different categories.

Answer: {"below 10,000,000 shares": ["GSE SYSTEMS INC", "CROSS TIMBERS ROYALTY TRUST"], "10,000,000 shares or more": ["HUGOTON ROYALTY TRUST"]}

B.6 Citation&Reference

<Multi_Documents>

Prompt: We hope you will carefully study the provided papers and determine the citation relationships between them. Please follow the instructions below strictly to complete the task:

#Specific Requirements:

1. Reference: When a given paper mentions other provided papers, those other papers are considered as "references" for the given paper. To summarize in this specific context, references are about what the given paper is using.
2. Citation: Conversely, when other provided papers mention the given paper in their works, the given paper is being "cited" by those other papers. To summarize in this specific context, citations are about who is using the given pa-

per.

3. Given a paper, you need to determine the citation or reference relationship between this paper and the other papers. Do not consider papers that are not provided.

3. Please present the paper titles in a json format as follows: {"Reference":["Reference Title 1", "Reference Title 2", ..., "Reference Title n"], "Citation":["Citation Title 1", "Citation Title 2", ..., "Citation Title n"]}

4. If a paper does not have any references or citations, please leave the corresponding list empty, e.g. {"Refernce":[]}, {"Citation":[]}

Question: The paper you need to analyze:Self-Discover: Large Language Models Self-Compose Reasoning Structures

Answer: {'Reference': ['# Plan, Verify and Switch: Integrated Reasoning with Diverse X-of-Thoughts ', '# StrategyLLM: Large Language Models as Strategy Generators, Executors, Optimizers, and Evaluators for Problem Solving '], 'Citation': ['# Can LLMs Solve Longer Math Word Problems Better? ']}

B.7 Case Classification

<Multi_Documents>

Prompt: Please answer the following questions based only on the judgment documents you have seen above. You only need to give the titles of the judgment documents that meet the requirements.

Question: After reading the above judgments, please classify all the judgments according to the following three types of cases: 'Civil Cases', 'Enforcement Cases', and 'Administrative Cases'.

Answer: {"Civil Cases": ["Judgment Document 2"], "Enforcement Cases": ["Judgment Document 4"], "Administrative Cases": ["Judgment Document 1", "Judgment Document 3"]}

B.8 Temporal Analysis

<Multi_Documents>

Prompt: We kindly ask you to review the financial statements of the companies provided above and answer the following questions based solely on the information you have seen.

If the question involves content not found in the financial statements, you may ignore this part and only answer the other parts.

Question: What is the trend in ARVANA INC's share capital from 2021 to 2024?

Answer: ARVANA INC's share capital has consistently increased from \$4,611 in 2021 to \$34,149 in 2022, \$35,949 in 2023, and \$107,847 in 2024.

B.9 Citation Chain

<Multi_Documents>

Prompt: We kindly ask you to thoroughly review the provided papers and construct a citation chain from them. Please adhere to the following instructions strictly while completing the task:

#Task Instructions:

Given several papers, you are required to identify and list the longest citation chain, which demonstrates the citation relationship among the provided papers.

#Specific Requirements:

1. Please present the titles of the papers in the form of a list, as follows: ["Title of Paper 1", "Title of Paper 2", ..., "Title of Paper n"].

2. Ensure that the citation chain in the list is linear and continuous, meaning that the first paper title in the list (Paper 1) should not cite any other works. Instead, it should be cited by the next paper in the list (Paper 2); subsequently, each paper should then be cited by the next one in the list, continuing up to the last paper (Paper n).

3. Consider only the citation relationships within the supplied collection of papers, and ensure that the citation chain accurately reflects the sequential citation order among these documents.

4. Do not take into account any articles not provided, and disregard other non-linear citation relationships.

Answer: ["# Very Deep Transformers for Neural Machine Translation", "# Understanding the Difficulty of Training Transformers", "# MonaCoBERT: Monotonic attention based ConvBERT for Knowledge Tracing"]

B.10 Link the Links

<Multi_Documents>

Prompt: Answer the following questions based solely on the judgment document you have seen above.

Question: After reading the above judgment document, I will give you several judgment results: <a list of judgment result: judgment1~judgment6> You need to determine the most likely judgment result for each of the above judgment documents.

Answer: {"Judgment Document 1": "Judgment Result 1", "Judgment Document 2": "Judgment Result 6", "Judgment Document 3": "Judgment Result 2", "Judgment Document 4": "Judgment Result 5"}

B.11 Solitaire

<Multi_Documents>

Prompt: Answer the following questions based solely on the judgment document you have seen above.

Question: Reading the above judgments, I will provide several case types arranged in a left-to-right sequence: ['CaseType1', 'CaseType2', 'CaseType3', 'CaseType4']. You need to sort all the judgment documents according to the above sequence of case types. The judgment documents only need to include the titles. Please provide the answer:

Answer: {"CaseType1": "Judgment Document 3", "CaseType2": "Judgment Document 1", "CaseType3": "Judgment Document 4", "CaseType4": "Judgment Document 6"}

C Length Distribution

As shown in Figure 4 and Table 7, we present the distribution of data lengths in Loong. It can be observed that the data is primarily distributed around 30-150k. Moreover, we have sufficient data in both shorter and longer ranges, allowing us to assess the model's capabilities across each length interval.

D RAG Detailed Results

We conducted experiments on GPT-4o and Qwen2-72B-Instruct with the addition of a RAG module. As shown in Table 8, Table 9, and Table 10, we have published detailed experimental results. It can




Figure 4: Test Case Length Distribution in Loong.

be seen that RAG achieved subpar results on our Loong, indicating that Loong requires the model to have genuine long-context understanding capabilities.

E Comparison of Evidence Distribution

In the same multi-document question-answering task, we compared the distribution of evidence related to the answers in the context for Loong (Ours) and LongBench (Bai et al., 2024). As shown in Table 11, we present data examples from Loong and LongBench. It can be observed that although LongBench contains a large number of passages, the evidence is only distributed within Passage 1. In contrast, in our Loong, the evidence is distributed across every document, requiring the model to understand each document in order to provide the correct answer.

F Comparison of Results with Other Benchmarks

Due to the rapid iteration of models and the high cost of closed-source model APIs, previous studies have rarely evaluated models that support context windows of 128k or more. Evaluating the latest models with ultra-long context windows and expensive closed-source models is our advantage compared to previous methods. Users can evaluate the gap between their models and the latest open-source and closed-source models at a low cost. To discover more interesting conclusions, we compared Loong with RULER (Hsieh et al., 2024),

Models	Loong	RULER	NOCHA
GPT-4o	53.47	-	55.80
GPT-4-Turbo	-	91.60	40.20
Gemini-1.5-pro	55.37	95.80	48.10
Claude3.5-Sonnet	48.85	-	41.00
Qwen2-72B-Instruct	43.29	85.90	43.35
GLM4-9B-Chat	38.31	89.90	27.07
Llama3.1-8B	36.31	88.30	16.53
Phi-3-mini-3.8B	14.54	68.80	9.30

Table 5: Compare Loong with RULER and NOCHA.

and the most recent work, NOCHA (Karpinska et al., 2024). The experimental results are shown in the Table 5.

From the result, we discovered that: (1) Although RULER focuses on synthetic tasks and NOCHA focuses on the domain of novels, the performance of the models on the three benchmarks is essentially consistent. (2) Gemini-1.5-pro achieved the best results in Loong and RULER, but GPT-4o had a significant lead on NOCHA. We believe that since NOCHA’s data sources only include novels, GPT-4o might have stronger capabilities in the domain of novel literature. (3) The large-parameter Qwen2-72B-Instruct underperformed in RULER compared to the smaller-parameter models GLM4-9B-Chat and Llama3.1-8B, which is inconsistent with the results of Loong. We believe that RULER emphasizes synthetic tasks, where the performance of such tasks is as sensitive as Needle In A Haystack. Specifically, simply adding some synthetic task data to the training corpus can yield good results. In real-world tasks like Loong, such anomalous phenomena generally do not occur.

G Results of Recall Rate by RAG

To show that RAG does not cover all evidence, we sampled all the questions from the financial test case in the Loong and analyzed RAG’s retrieval results. Since the answers to our questions are distributed across all documents on average, we evaluated whether the retrieved top-k passages cover all documents to reflect the RAG’s recall of evidence. *This metric will be higher than the actual evidence recall rate because even if all required documents are retrieved, the specific passages may not necessarily contain the evidence. If the model fails to retrieve all documents, it will certainly be unable to retrieve all evidence.* The evaluation metric used is $Recall@n$, where n represents the number

of retrieved passages. If the n passages contain all documents, $Recall@n$ is set to 1; otherwise, it is 0. The results can be seen in Table 6. Evidently, in the Loong dataset, RAG struggles to retrieve all relevant documents. Even when the topK reaches 50, the highest recall is only 0.64. In reality, the recall rate for all evidence will be even lower.

H Detailed URL of Document Source

China Judge Online:

<https://wenshu.court.gov.cn/>;

U.S. SEC:

<https://www.sec.gov/>;

cninf:

<http://www.cninfo.com.cn/>;

Arxiv:

<https://arxiv.org/>;

Semantic Scholar:

<https://www.semanticscholar.org/>

Model	Recall@10	Recall@20	Recall@30	Recall@40	Recall@50
Openai Embedding, chunk size=1024	0.20	0.34	0.43	0.51	0.58
Openai Embedding, chunk size=2048	0.20	0.36	0.48	0.57	0.62
BGE Embedding, chunk size=1024	0.19	0.34	0.44	0.53	0.59
BGE Embedding, chunk size=2048	0.21	0.38	0.49	0.58	0.64

Table 6: The Recall Rate of all documents by using RAG.

Dataset	#data in 10-50k	#data in 50-100K	#data in 100K-200K	#data in 200-250K
<i>Spotlight Locating</i>	53	70	80	47
<i>Comparison</i>	60	105	95	40
Sequential Enumeration	24	29	20	14
Extremum Acquisition	16	55	59	13
Range Awareness	20	21	16	13
<i>Clustering</i>	113	246	194	88
Report Integration	40	90	90	30
Citation&Reference	37	120	79	34
Case Classification	36	36	25	24
<i>Chain of Reasoning</i>	97	143	112	57
Temporal Analysis	10	40	35	15
Citation Chain	33	50	41	6
Link the Links	35	25	28	25
Solitaire	28	19	8	11
<i>Overall</i>	323	564	281	232
Chinese	240	284	251	130
English	83	280	230	102

Table 7: Data length distributions in Loong benchmark.

Model	Spotlight Locating		Comparison		Clustering		Chain of Reasoning		Overall	
GPT4o (128K)	73.95	0.62	50.50	0.28	44.29	0.09	57.95	0.28	53.47	0.26
w/ Openai Embedding, Top k=5	56.97	0.36	31.28	0.14	27.71	0.03	26.65	0.04	32.85	0.11
w/ BGE Embedding, Top k=5	63.32	0.44	34.63	0.17	26.74	0.03	26.21	0.04	34.01	0.13
w/ Openai Embedding, Top k=10	65.20	0.46	36.80	0.19	33.06	0.04	33.26	0.08	38.80	0.14
w/ BGE Embedding, Top k=10	68.27	0.50	39.51	0.22	31.91	0.04	30.71	0.07	38.71	0.15
w/ Openai Embedding, Top k=30	64.32	0.43	42.15	0.26	41.02	0.08	40.14	0.14	44.62	0.18
w/ BGE Embedding, Top k=30	64.76	0.45	47.56	0.32	40.43	0.08	40.82	0.17	45.67	0.21
w/ Openai Embedding, Top k=50	65.59	0.45	41.69	0.28	34.49	0.04	39.74	0.14	42.70	0.18
w/ BGE Embedding, Top k=50	63.28	0.42	47.05	0.32	42.64	0.10	41.97	0.18	46.52	0.21
Qwen2-72B-Instruct (128K)	54.17	0.36	42.38	0.20	36.71	0.04	47.76	0.18	43.29	0.15
w/ Openai Embedding, Top k=5	57.57	0.40	30.98	0.14	27.91	0.02	28.30	0.04	33.22	0.10
w/ BGE Embedding, Top k=5	62.02	0.44	32.90	0.16	27.05	0.02	29.26	0.06	34.18	0.12
w/ Openai Embedding, Top k=10	62.52	0.44	35.79	0.18	30.16	0.03	32.67	0.08	36.92	0.13
w/ BGE Embedding, Top k=10	69.24	0.51	36.78	0.18	29.07	0.02	31.90	0.07	37.50	0.14
w/ Openai Embedding, Top k=30	64.11	0.43	41.91	0.26	35.61	0.04	42.61	0.19	42.98	0.18
w/ BGE Embedding, Top k=30	58.59	0.38	42.66	0.23	33.77	0.05	40.39	0.17	40.94	0.16
w/ Openai Embedding, Top k=50	57.87	0.37	42.10	0.25	32.78	0.03	42.39	0.19	40.86	0.17
w/ BGE Embedding, Top k=50	56.93	0.37	39.51	0.20	32.67	0.04	43.44	0.21	40.46	0.16

Table 8: Overall results (%) of adding RAG module on GPT4o and Qwen2-72B-Instruct.

Model	Spotlight Locating		Comparison		Clustering		Chain of Reasoning		Overall	
Set1 (10K-50K)										
GPT-4o (128K)	85.67	0.81	64.27	0.33	57.01	0.24	81.58	0.55	70.40	0.44
w/ Openai Embedding, Top k=5	47.60	0.31	29.75	0.10	29.10	0.06	31.46	0.08	32.98	0.11
w/ BGE Embedding, Top k=5	57.17	0.43	34.15	0.12	30.71	0.07	28.77	0.08	35.23	0.14
w/ Openai Embedding, Top k=10	61.25	0.44	38.33	0.17	37.00	0.08	41.67	0.16	42.63	0.18
w/ BGE Embedding, Top k=10	61.00	0.44	39.74	0.19	36.14	0.08	34.90	0.11	40.44	0.17
w/ Openai Embedding, Top k=30	55.15	0.37	46.60	0.28	45.54	0.13	51.98	0.27	49.23	0.24
w/ BGE Embedding, Top k=30	57.40	0.38	52.25	0.32	46.54	0.18	50.02	0.25	50.41	0.26
w/ Openai Embedding, Top k=50	55.47	0.40	49.62	0.33	39.61	0.10	46.08	0.20	46.82	0.24
w/ BGE Embedding, Top k=50	52.08	0.38	53.42	0.37	49.83	0.21	48.88	0.24	50.55	0.27
Set2 (50K-100K)										
GPT-4o (128K)	86.76	0.72	59.81	0.40	47.83	0.11	62.09	0.34	58.38	0.29
w/ Openai Embedding, Top k=5	56.01	0.35	39.56	0.22	31.84	0.04	27.01	0.03	35.31	0.11
w/ BGE Embedding, Top k=5	67.33	0.43	43.90	0.28	29.37	0.04	27.84	0.04	36.72	0.14
w/ Openai Embedding, Top k=10	64.77	0.45	45.44	0.31	36.07	0.05	32.29	0.05	40.54	0.15
w/ BGE Embedding, Top k=10	72.07	0.52	50.15	0.32	34.35	0.05	33.49	0.07	41.90	0.17
w/ Openai Embedding, Top k=30	65.87	0.42	50.05	0.34	44.08	0.11	42.60	0.15	47.48	0.20
w/ BGE Embedding, Top k=30	65.26	0.49	55.21	0.43	43.82	0.10	42.80	0.18	48.31	0.23
w/ Openai Embedding, Top k=50	67.21	0.46	51.38	0.38	31.65	0.03	40.69	0.12	43.52	0.19
w/ BGE Embedding, Top k=50	67.43	0.46	53.98	0.39	45.04	0.12	46.94	0.21	49.96	0.24
Set3 (100K-200K)										
GPT-4o (128K)	74.84	0.65	42.40	0.21	38.70	0.04	45.06	0.09	46.95	0.19
w/ Openai Embedding, Top k=5	67.45	0.49	29.00	0.13	25.09	0.01	27.22	0.02	33.69	0.12
w/ BGE Embedding, Top k=5	71.12	0.56	31.36	0.14	25.32	0.00	25.78	0.04	34.43	0.13
w/ Openai Embedding, Top k=10	72.37	0.55	31.41	0.13	30.59	0.01	33.14	0.08	38.38	0.14
w/ BGE Embedding, Top k=10	79.04	0.67	34.29	0.18	30.59	0.02	29.69	0.06	39.22	0.17
w/ Openai Embedding, Top k=30	74.03	0.57	37.00	0.22	39.53	0.04	36.07	0.09	43.91	0.18
w/ BGE Embedding, Top k=30	75.45	0.59	45.96	0.31	36.91	0.04	39.16	0.14	45.68	0.21
w/ Openai Embedding, Top k=50	77.24	0.59	25.02	0.12	37.55	0.05	40.93	0.15	44.52	0.19
w/ BGE Embedding, Top k=50	74.19	0.55	46.15	0.31	39.71	0.04	36.56	0.15	45.99	0.20
Set4 (200K-250K)										
GPT-4o (128K)	36.79	0.19	23.97	0.08	30.40	0.00	32.89	0.07	31.11	0.07
w/ Openai Embedding, Top k=5	50.76	0.22	17.25	0.00	19.53	0.00	16.61	0.00	24.91	0.05
w/ BGE Embedding, Top k=5	51.02	0.26	18.75	0.03	17.83	0.00	18.77	0.02	25.07	0.06
w/ Openai Embedding, Top k=10	57.98	0.31	23.00	0.03	25.08	0.00	21.29	0.02	30.00	0.07
w/ BGE Embedding, Top k=10	51.48	0.25	23.36	0.05	22.55	0.00	18.95	0.02	27.48	0.06
w/ Openai Embedding, Top k=30	55.85	0.26	26.38	0.08	29.94	0.00	21.81	0.02	32.66	0.07
w/ BGE Embedding, Top k=30	53.94	0.21	24.82	0.07	30.77	0.00	23.61	0.05	32.68	0.07
w/ Openai Embedding, Top k=50	55.00	0.28	12.79	0.04	29.64	0.00	24.67	0.07	31.88	0.08
w/ BGE Embedding, Top k=50	51.17	0.21	23.36	0.10	33.08	0.02	28.39	0.09	33.82	0.09

Table 9: The result of adding RAG module on GPT-4o with different length sets.

Model	Spotlight Locating		Comparison		Clustering		Chain of Reasoning		Overall	
Set1 (10K-50K)										
Qwen2-72B-Instruct (128K)	68.49	0.55	60.60	0.37	47.08	0.08	70.39	0.36	60.11	0.29
w/ Openai Embedding, Top k=5	54.62	0.45	26.17	0.08	29.60	0.03	34.41	0.08	34.51	0.12
w/ BGE Embedding, Top k=5	62.92	0.53	30.92	0.08	31.28	0.03	32.95	0.11	36.91	0.15
w/ Openai Embedding, Top k=10	59.81	0.43	34.93	0.15	29.33	0.02	41.27	0.15	38.96	0.15
w/ BGE Embedding, Top k=10	72.13	0.62	32.42	0.12	31.90	0.05	44.12	0.20	42.27	0.20
w/ Openai Embedding, Top k=30	57.26	0.40	45.43	0.28	40.04	0.06	57.32	0.35	49.06	0.24
w/ BGE Embedding, Top k=30	56.37	0.33	46.27	0.30	38.35	0.10	51.49	0.29	46.69	0.23
w/ Openai Embedding, Top k=50	51.08	0.35	44.53	0.27	37.96	0.05	53.95	0.35	46.11	0.23
w/ BGE Embedding, Top k=50	53.47	0.37	47.31	0.29	36.42	0.06	54.65	0.35	46.62	0.24
Set2 (50K-100K)										
Qwen2-72B-Instruct (128K)	64.53	0.43	42.60	0.21	38.52	0.05	51.18	0.20	45.71	0.17
w/ Openai Embedding, Top k=5	56.64	0.40	36.68	0.19	30.91	0.03	28.38	0.01	34.54	0.10
w/ BGE Embedding, Top k=5	67.29	0.47	43.39	0.28	28.31	0.03	32.22	0.07	36.95	0.14
w/ Openai Embedding, Top k=10	67.07	0.53	44.30	0.27	34.31	0.05	34.03	0.06	40.17	0.15
w/ BGE Embedding, Top k=10	71.74	0.54	47.68	0.30	30.55	0.03	30.57	0.03	38.80	0.14
w/ Openai Embedding, Top k=30	66.27	0.46	46.28	0.31	38.95	0.05	46.15	0.22	45.42	0.19
w/ BGE Embedding, Top k=30	57.35	0.41	46.92	0.29	35.30	0.05	42.82	0.20	42.04	0.18
w/ Openai Embedding, Top k=50	55.94	0.32	47.94	0.31	34.32	0.03	46.64	0.21	42.60	0.16
w/ BGE Embedding, Top k=50	59.41	0.39	38.52	0.21	35.40	0.06	45.47	0.24	41.51	0.17
Set3 (100K-200K)										
Qwen2-72B-Instruct (128K)	46.99	0.27	37.06	0.13	31.50	0.02	35.01	0.07	35.94	0.09
w/ Openai Embedding, Top k=5	63.91	0.44	33.56	0.17	25.98	0.01	28.98	0.04	34.48	0.12
w/ BGE Embedding, Top k=5	64.81	0.47	30.27	0.14	25.88	0.01	27.86	0.05	33.70	0.12
w/ Openai Embedding, Top k=10	67.50	0.46	33.44	0.16	27.94	0.02	31.62	0.06	36.47	0.13
w/ BGE Embedding, Top k=10	75.88	0.56	33.76	0.15	27.20	0.01	30.17	0.04	37.28	0.14
w/ Openai Embedding, Top k=30	73.69	0.55	42.20	0.27	32.78	0.02	37.65	0.13	42.60	0.18
w/ BGE Embedding, Top k=30	67.50	0.47	42.42	0.18	32.34	0.03	37.85	0.12	41.35	0.15
w/ Openai Embedding, Top k=50	67.44	0.50	41.82	0.24	31.59	0.04	37.29	0.12	40.90	0.18
w/ BGE Embedding, Top k=50	62.56	0.42	40.41	0.18	29.82	0.02	40.31	0.14	39.84	0.15
Set4 (200K-250K)										
Qwen2-72B-Instruct (128K)	33.18	0.16	26.59	0.08	29.84	0.01	25.81	0.04	28.92	0.06
w/ Openai Embedding, Top k=5	51.49	0.26	17.12	0.03	21.59	0.00	16.37	0.00	25.59	0.06
w/ BGE Embedding, Top k=5	48.40	0.26	14.55	0.00	20.69	0.00	18.07	0.00	24.63	0.05
w/ Openai Embedding, Top k=10	50.32	0.28	20.30	0.03	24.56	0.00	16.38	0.00	27.08	0.06
w/ BGE Embedding, Top k=10	51.02	0.28	21.88	0.03	25.45	0.00	17.29	0.00	28.10	0.06
w/ Openai Embedding, Top k=30	52.17	0.24	24.60	0.10	26.78	0.00	17.79	0.00	29.29	0.07
w/ BGE Embedding, Top k=30	47.98	0.21	26.82	0.10	26.70	0.00	20.02	0.00	29.44	0.06
w/ Openai Embedding, Top k=50	51.63	0.26	23.62	0.08	24.49	0.00	21.84	0.02	29.14	0.07
w/ BGE Embedding, Top k=50	47.23	0.28	27.78	0.08	26.48	0.00	24.44	0.02	30.52	0.08

Table 10: The result of adding RAG module on Qwen2-72B-Instruct with different length sets.

LongBench

Passage 1: The Real Glory. The Real Glory is a 1939 Samuel Goldwyn Productions adventure film starring Gary Cooper, David Niven, Andrea Leeds and Broderick Crawford released by United Artists in the weeks immediately following Nazi Germany's invasion of Poland. Based on a 1937 novel of the same name by Charles L. Clifford and directed by Henry Hathaway, the film is set against the backdrop of the Moro Rebellion during the American occupation of the Philippines at the beginning of the 20th century . . .

Passage 2: Jay Sheffield. Jay Howard Sheffield (September 25, 1934 – June 25, 1998) was an American actor, who appeared on the stage, in films, and on television. He married Barbara Babcock on June 9, 1962, in San Mateo, California. They later divorced . . .

Passage 3: David Niven. James David Graham Niven (; 1 March 1910 – 29 July 1983) was a British actor, soldier, memoirist, and novelist. Niven was known as a handsome and debonair leading man in Classic Hollywood films. He received an Academy Award and a Golden Globe Award . . .

Passage 4: Phileas Fogg snacks. Phileas Fogg Ltd is a company that produces snack products in the United Kingdom that was created in 1982 by Derwent Valley Foods. The brand is named for Phileas Fogg, the protagonist of Jules Verne's Around the World in Eighty Days . . .

Passage 5: Jules Verne Trophy. The Jules Verne Trophy is a prize for the fastest circumnavigation of the world by any type of yacht with no restrictions on the size of the crew provided the vessel has registered with the organization and paid an entry fee. A vessel holding the Jules Verne trophy will not necessarily hold the absolute round the world record . . .

Question: The actor that plays Phileas Fogg in "Around the World in 80 Days", co-starred with Gary Cooper in a 1939 Goldwyn Productions film based on a novel by what author?

Answer: Charles L. Clifford

Loong (Ours)

Document 1: THE ARENA GROUP HOLDINGS, INC. . Proceeds from Simplify loan: \$7,748. Unearned revenue: \$(11,665). Amortization of debt discounts: \$536. Cash and cash equivalents: \$4,003 . Noncash and accrued interest: \$2,839. Loss on impairment of assets: \$40,589. Accounts receivable, net: \$12,029. Subscription refund liability: \$18. Accounts payable: \$(102). Subscription acquisition costs: \$6,131. Change in fair value of contingent consideration: \$313. (\$ in thousands, except share data) . . .

Document 2: General Enterprise Ventures, Inc. . For purposes of balance sheet presentation and reporting of cash flows, the Company considers all unrestricted demand deposits, money market funds and highly liquid debt instruments with an original maturity of less than 90 days to be cash and cash equivalents. The Company did not have any cash equivalents at March 31, 2024. The Company had cash of \$549,755 at March 31, 2024 . . .

Document 3: BROAD STREET REALTY, INC. . The carrying amounts of cash and cash equivalents, restricted cash, receivables and payables are reasonable estimates of their fair value as of March 31, 2024 due to the short-term nature of these instruments. Reconciliation of cash and cash equivalents and restricted cash: Cash and cash equivalents: \$14,631 (in thousands) . . .

Question: Please list the 'Cash and Cash Equivalents' of the aforementioned companies in ascending order.

Answer: 1. General Enterprise Ventures, Inc.: \$549,755. 2. Arena Group Holdings, Inc.: \$4,003 in thousands. 3. Broad Street Realty, Inc.: \$14,631 in thousands.

Table 11: Comparison of Evidence Distribution in Examples from LongBench and Loong (Ours). Evidence related to the answers is highlighted with Orange Background .