

BAMBOO: A Comprehensive Benchmark for Evaluating Long Text Modeling Capacities of Large Language Models

Zican Dong¹, Tianyi Tang¹, Junyi Li^{1,2}, Wayne Xin Zhao^{1*}, Ji-Rong Wen^{1,3}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²DIRO, Université de Montréal

³School of Information, Renmin University of China

{dongzican, steven_tang, lijunyi, jrwen}@ruc.edu.cn, batmanfly@gmail.com

Abstract

Large language models (LLMs) have achieved dramatic proficiency over NLP tasks with normal length. Recently, multiple studies have committed to extending the context length and enhancing the long text modeling capabilities of LLMs. To comprehensively evaluate the long context ability of LLMs, we propose **BAMBOO**, a multi-task long context benchmark. BAMBOO has been designed with four principles: comprehensive capacity evaluation, avoidance of data contamination, accurate automatic evaluation, and different length levels. It consists of 10 datasets from 5 different long text understanding tasks, *i.e.*, question answering, hallucination detection, text sorting, language modeling, and code completion, to cover various domains and core capacities of LLMs. We conduct experiments with five widely-used long-context models and further discuss five key questions for long text research. In the end, we discuss problems of current long-context models and point out future directions for enhancing long text modeling capacities. We release our data, prompts, and code at <https://github.com/RUCAIBox/BAMBOO>.

Keywords: BAMBOO, Long Text, Large Language Models, Benchmark, Evaluation

1. Introduction

Recently, large language models (LLMs) have exhibited huge success across various NLP tasks (Ouyang et al., 2022; Touvron et al., 2023a; Zhao et al., 2023). Despite their remarkable capacities, existing LLMs still face constraints due to their limited context length, suffering from performance decline when the input text exceeds their length limit. To enhance the capability of LLMs to model long texts, several studies extend the context window through modifying position embeddings (kaio-kendev, 2023; Chen et al., 2023), and continuing to fine-tune the models on multi-turn conversations or long text datasets (Zheng et al., 2023). Moreover, numerous research also explores transforming the original long text into short segments within the context length based on context compression techniques (Zhou et al., 2023; Jiang et al., 2023).

To push forward the research of long context modeling, it is crucial to establish an automatic, reliable, and comprehensive evaluation benchmark for assessing the performance of LLMs in tasks involving lengthy text. In existing literature, several benchmarks have been proposed to evaluate the long context modeling ability of LLMs, such as Zero-SCROLLS (Shaham et al., 2023), L-Eval (An et al., 2023), and LongBench (Bai et al., 2023). However, a crucial aspect overlooked by these studies is the potential issue of data contamination (Golchin and Surdeanu, 2023), *i.e.*, there might exist overlaps between the pre-training data and evaluation data.

In the meanwhile, these benchmarks consist of some text generation tasks evaluated by automatic metrics such as BLEU (Wang et al., 2023a). However, these metrics are blamed for their inaccurate evaluation of model performance.

To this end, we propose **BAMBOO**, a comprehensive benchmark to analyze the long text modeling capacities of LLMs. In BAMBOO benchmark, we manually construct ten datasets from five tasks, including question answering, hallucination detection, text sorting, language modeling, and code completion. It comprehensively evaluates language generation, knowledge utilization, reasoning, and tool manipulation abilities over long texts. We aim to ensure that the salient information is spread throughout the complete long texts. Consequently, the LLM needs to model long-range dependencies to identify both the coarse-grained information related to the global comprehension of the complete text and fine-grained information about details across several sentences. Moreover, with datasets from multiple domains, BAMBOO can evaluate the performance of LLMs across various sources.

BAMBOO is divided into two subsets at *different length levels*, *i.e.*, BAMBOO-4k and BAMBOO-16k, according to the number of tokens in the prompt¹. Each subset has 1502 samples with an average length of 3152 and 7500 tokens, respectively. To alleviate the *data contamination* with pre-training corpora, BAMBOO is constructed based on the data

¹All references to the number of tokens in our paper are computed using the gpt-3.5-turbo tokenizer.

*Corresponding author.

Benchmark	ZeroSCROLLS	L-Eval	LongBench	BAMBOO
Comprehensive Capacity Evaluation	✓	✓	✓	✓
Accurate Automatic Evaluation	×	×	×	✓
Avoidance of Data Contamination	×	×	×	✓
Different Length Levels	×	×	✓	✓
#Tasks	3	4	6	5

Table 1: Comparison with other long text evaluation benchmarks. #Tasks denotes the number of tasks in the benchmark.

sources released in 2023. Furthermore, we select diverse tasks with *accurate automatic metrics* for reliable evaluation. Table 1 presents a comparison of BAMBOO with existing benchmarks.

We extensively evaluate five long-context LLMs on the BAMBOO benchmark. We observe that ChatGPT-16k consistently demonstrates optimal performance across most datasets, whereas other models usually struggle, especially on uncommon tasks. We conduct on a series of analytical experiments to shed light on the key questions in long text modeling. Our findings are listed as follows:

- Extending the context window of LLMs is a double-edged sword, which is beneficial for medium-length texts but harmful for short texts.
- Using evidence and properly compressing long texts can generally enhance the performances of LLMs. However, smaller LLMs still exhibit limited performance even given more concise input.
- LLMs’ performances on different datasets are sensitive to positions of instructions. Additionally, locating evidence on both ends is beneficial due to larger attention scores.

Ultimately, we undertake a discussion of problems associated with instruction forgetting, format errors, reasoning capabilities, and uncommon tasks on long text modeling. We conclude by identifying future research directions for advancing the long text modeling capabilities of LLMs.

2. Related Work

Long-Context Architectures Owing to the quadratic memory and computational complexity of self-attention, original Transformer models encountered significant challenges in scaling to extended contexts. To enhance the efficiency of Transformers, researchers explored multiple innovations in the attention mechanism. These include recurrent Transformers (Dai et al., 2019; Rae et al., 2020), local attention (Child et al., 2019; Beltagy et al., 2020; Zaheer et al., 2020), content-based sparse attention (Kitaev et al., 2020; Tay et al., 2020; Roy et al., 2021), and attention approximation (Wang et al., 2020; Ma

et al., 2021; Katharopoulos et al., 2020; Peng et al., 2021). In addition to the above Transformer variants, several models were introduced, replacing attention with alternative modules. S4 (Gu et al., 2022), DSS (Gupta et al., 2022), and GSS (Mehta et al., 2023) employed state space models, while RetNet (Sun et al., 2023a) made use of retention mechanism, achieving parallel training and recurrent inference for long sequences.

Adapting LLMs for Long Contexts Recently, with the booming of LLMs, there has been an increase in research efforts for overcoming the length limitations of off-the-shelf LLMs. Through modifying RoPE (Su et al., 2021) with position interpolation (kaiokendev, 2023; Chen et al., 2023), truncated basis (Pal et al., 2023), or modifying bases (Xiong et al., 2023; Peng et al., 2023b; Liu et al., 2023b; Rozière et al., 2023), the context windows of LLMs could be extended to 32k or more tokens. Equipped with fusion-in-decoder (Izacard and Grave, 2021; Ratner et al., 2023) or external memory techniques (Bertsch et al., 2023; Tworkowski et al., 2023; Wang et al., 2023b), LLMs could access tokens from a distant past during the generation stage. In addition, context compression techniques, such as retrieval-augmentation and recurrence, could significantly expand their capacity to model longer contexts, where LLMs only need to process a short segment at once (Zhou et al., 2023; Liang et al., 2023; Packer et al., 2023; Xu et al., 2023; Jiang et al., 2023).

Evaluation for Long Text Modeling For evaluating long text modeling capabilities, diverse datasets were proposed, mainly classified into three task categories: language modeling, question answering, and summarization. The objective of language modeling was to sequentially predict the next tokens based on previous context (Rae et al., 2020). However, popular benchmarks did not adopt it since most tokens could be easily predicted within a limited context (Sun et al., 2021). Long text summarization datasets encompassed normal summarization (Cohan et al., 2018; Huang et al., 2021; Chen et al., 2022) and query-based summarization (Zhong et al., 2021), involving the condensation of either the entire or a portion of lengthy texts into concise summaries. As for question answering, given a question and a single document (Pang et al., 2022; Dasigi et al., 2021) or multiple documents (Yang et al., 2018), the model needed to provide the answer of the question.

To facilitate comprehensive evaluations of long text modeling abilities, several long text benchmarks were introduced. Scrolls (Shaham et al., 2022) collected and standardized several long text reasoning datasets spanning various tasks and do-

Dataset	#Input	#Example	Metric	Domain	Source
AltQA	3243/13084	200/200	accuracy	Wikipedia	https://github.com/abacusai/long-context
PaperQA	3101/6838	100/100	accuracy	Paper	https://aclanthology.org/
MeetingQA	2738/9838	100/100	accuracy	Meeting	https://record.assembly.wales/
SenHallu	3170/6357	200/200	P/R/F1	Paper	https://aclanthology.org/
AbsHallu	3314/6445	200/200	P/R/F1	Paper	https://aclanthology.org/
ShowsSort	2992/6411	200/200	CI	TV Shows	https://tvmeg.com/
ReportSumSort	3753/8309	150/150	CI	Reports	https://www.gao.gov/
ShowsPred	2389/4860	100/100	accuracy	TV Shows	https://tvmeg.com/
MeetingPred	3689/11578	100/100	accuracy	Meeting	https://record.assembly.wales/
PrivateEval	3149/6230	152/152	pass@1	Code	https://github.com/microsoft/PyCodeGPT

Table 2: Overview of our BAMBOO benchmark. #Input and #Example represent the average number of tokens in the input text and the number of examples of each dataset, respectively. For #Input and #Example, “number1/number2” is the number of the middle and long levels, respectively. P/R/F1 denotes precision/recall/F1, and CI stands for concordance index.

mains into a unified format. In the wake of the heightened interest in LLMs, ZeroScrolls (Shaham et al., 2023), L-Eval (An et al., 2023), and Long-Bench (Bai et al., 2023) were released to assess the zero-shot long text modeling ability of LLMs. Nevertheless, they all encountered potential data contamination problems and suffered from inaccurate automatic metrics.

3. BAMBOO Benchmark

Our BAMBOO benchmark commits to equitably and thoroughly evaluating LLMs’ long text modeling proficiency. It contains ten elaborately constructed datasets across a broad range of domains and evaluations. The overview of the BAMBOO benchmark is shown in Table 2.

3.1. Design Principles

Different from existing long text benchmarks in Table 1, our BAMBOO benchmark is meticulously designed to comprehensively evaluate the long text modeling capacity of LLMs. Our benchmark can potentially circumvent the issue of data contamination in previous benchmarks, pursue more accurate automatic evaluation, and cater to LLMs at different length levels.

Comprehensive Capacity Evaluation The primary objective of BAMBOO is comprehensively evaluating the abilities of LLMs in modeling long texts via different tasks and domains. To meet the real needs in long text scenarios, all the tasks assess the LLMs’ ability to capture long-range dependencies. Inspired by Zhao et al. (2023), BAMBOO covers diverse tasks evaluating language generation, knowledge utilization, reasoning, and tool manipulation capacities over long texts. BAMBOO can also evaluate the coarse-grained comprehension of entire texts and the fine-grained reasoning

concerning specific details. Moreover, we collect data from various sources, assessing and comparing abilities across diverse domains.

Avoidance of Data Contamination In BAMBOO, we make a deliberate effort to minimize the overlap between test data and the training corpora of LLMs as much as possible. One of the key challenges when evaluating the long text capability of LLMs is the model has “seen” the input text during the training stages, *i.e.*, data contamination (Golchin and Surdeanu, 2023). To alleviate this issue, we take a cautious approach by retaining data only released in 2023, since training data cutoff for gpt-3.5-turbo and Claude2 are up to September 2021² and the early 2023³, respectively. Yet, since the specifics of their training data remain undisclosed for some models, we cannot guarantee the absence of any pre-training data used in our benchmark. In addition, we also utilize datasets with old data sources following Pal et al. (2023), but any appearances of the answers have been modified. Thus, the model must respond based solely on the input, reducing the potential effect of exposure to similar data.

Accurate Automatic Evaluation Each task in BAMBOO is designed for precise automatic evaluation. Previous benchmarks include text generation tasks (*e.g.*, text summarization and question answering). The generations of LLMs may exhibit variations in expression compared to the references, making it challenging to accurately evaluate their outputs with automatic metrics (Tang et al., 2023), even using LLMs as evaluator (Wang et al., 2023a). Hence, we transform some generation tasks into multi-choice tasks and select tasks that can be precisely evaluated. All the tasks can be measured by

²<https://platform.openai.com/docs/models/gpt-3-5>

³<https://www.anthropic.com/index/claude-2>

metrics such as accuracy, pass@1, etc.

Different Length Levels In BAMBOO, each task encompasses two distinct length levels. The most popular LLMs, e.g., Llama2 (Touvron et al., 2023b) and gpt-3.5-turbo², only have a context window of 4k tokens while some of them extend their context windows to 16k or more tokens, including gpt-3.5-turbo-16k² and ChatGLM2-6b-32k (Zeng et al., 2023). However, previous benchmarks did not thoroughly consider different length distributions of different datasets, making it difficult to analyze the impact of length changes on LLMs in the same task. Instead, we divide each task of BAMBOO into two subsets based on the input length, with the maximum length of 4k and 16k tokens, respectively. Thus, we can effectively measure the impact of lengths on the proficiency of LLMs.

3.2. Data Collection

To ensure domain diversity, we collect long texts from four distinct sources, namely NLP research papers, government reports, TV show transcripts, and committee meetings transcripts. All these data sources are publicly accessible and released in 2023 to avoid pre-training data contamination. We crawl and parse the source websites, extracting only the plain text while excluding images, tables, and footnotes. Texts containing fewer than 1000 tokens are excluded, and overly lengthy texts are truncated. In addition to these latest corpora, we also consider the pseudo long-context datasets where potential data contamination has attempted to be avoided.

In pursuit of the objective of encompassing diverse length levels, we divide samples within the BAMBOO benchmark into two subsets based on the length of the complete input text. One subset contains prompts up to 4k tokens (BAMBOO-4k), while the other comprises prompts with tokens from 4k to 16k (BAMBOO-16k).

3.3. Task Construction

We aim to assess long text modeling ability in BAMBOO thoroughly. To attain the goal, we devise five distinct tasks and comprise ten corresponding datasets based on the fundamental principle.

Question Answering In the question answering task, we primarily focus on evaluating knowledge utilization and reasoning ability. With the assistance of human labelers, we manually construct two multi-choice question answering datasets, **PaperQA** and **MeetingQA**, to evaluate the LLMs' comprehension of a long document or dialogue. We ask annotators to rephrase the expressions of questions and options instead of directly copying the

original expressions in given texts. We also ensure that most of the answers need to be reasoned from evidence from multiple paragraphs. In addition, we utilize **AitQA** (Pal et al., 2023), comprising multiple Wikipedia documents. Each question requires a numerical answer, and every occurrence of the answer within the documents is modified to another number to avoid data contamination.

Hallucination Detection Hallucination is a prevalent phenomenon in LLMs, i.e., the generated content may conflict with or be beyond the source. In this task, models must utilize the knowledge of given contexts to detect hallucinations. We create two novel datasets, i.e., **SenHallu** and **AbsHallu**, using ChatGPT to generate hallucinations based on the correct hypothesis (e.g., a few words modification or a make-up sentence insertion) (Li et al., 2023b). Based on the concept of natural language inference tasks, we present LLMs with a paper and a hypothesis. Then, we request LLMs to determine whether the hypothesis entails or contradicts the paper as a binary classification task.

Text Sorting To evaluate model's capacity of reasoning over the logical order of texts, we devise two text sorting datasets that require LLMs to reorder the shuffled texts according to contextual information of the whole texts. **ShowsSort** is a sorting task involving shuffled segments of a TV show transcript. The original order is reconstructed based on contextual continuity and the inherent meaning of each segment. **ReportSumSort** is another sorting task with the input of a complete government report and a list of shuffled summaries. The objective is to reorder these summaries according to their positions within the segments represented in the report.

Language Modeling Language modeling task is a typical language generation task where the next token is sequentially predicted based on the past context. In order to evaluate the capacity to handle long-range dependencies, we design the task intending to predict the last utterance's speaker within a long dialogue. Specifically, we adopt multi-turn conversation as the source and format each utterance as "{utterance} said by {speaker}". Then, we remove the speaker of the last turn of conversation as the target. Successfully predicting the speaker necessitates the model's grasp of the characteristics of existing speakers and the contextual conversations. We construct **ShowsPred** and **Meeting-Pred** with transcripts of TV shows and meetings as sources, respectively.

Code Completion To assess the capability of LLMs to effectively utilize external tools through

Models	ShowsSort		ReportSumSort		ShowsPred		MeetingPred		SenHallu	
	16k	4k	16k	4k	16k	4k	16k	4k	16k	4k
gpt-3.5-turbo-16k	59.0	58.2	69.6	77.6	55.0	53.0	57.0	73.0	62.7/99.0/76.7	68.1/96.0/79.7
Claude2-100k	53.1	54.0	59.7	66.7	8.0	26.0	0.0	21.0	67.3/99.0/80.2	69.3/97.0/80.8
ChatGLM2-32k	32.0	42.7	31.6	43.5	13.0	8.0	14.0	25.0	52.9/ 100.0/69.2	52.7/96.0/68.1
Vicuna-v1.5-16k	53.2	53.6	50.1	51.4	5.0	11.0	5.2	64.0	56.4/92.9/70.2	67.7/99.0/69.3
Longchat-v1.5-16k	53.2	53.8	44.5	47.9	1.0	1.0	29.9	36.0	50.3/ 100.0/66.9	51.2/ 100.0/67.8
Random Baseline	50.0	50.0	50.0	50.0	7.6	10.3	8.0	13.9	50.0/50.0/50.0	50.0/50.0/50.0

Models	PrivateEval		AltQA		PaperQA		MeetingQA		AbsHallu	
	16k	4k	16k	4k	16k	4k	16k	4k	16k	4k
gpt-3.5-turbo-16k	21.7	21.1	72.0	76.5	75.0	78.0	72.0	75.0	55.6/ 100.0/71.4	57.2/99.0/72.5
Claude2-100k	0.7	7.2	4.5	27.0	51.0	69.0	47.0	63.0	56.8/100.0/72.5	58.2/99.0/73.3
ChatGLM2-32k	3.5	0.7	67.0	64.0	65.0	62.0	65.0	52.0	50.5/ 100.0/66.7	50.0/99.0/66.4
Vicuna-7b-v1.5-16k	2.0	3.3	25.0	30.5	37.0	29.0	27.0	31.0	51.0/99.0/67.3	53.3/88.0/66.4
Longchat-7b-v1.5-16k	0.7	3.9	41.7	64.0	44.0	44.0	36.0	42.0	50.3/ 100.0/66.9	50.0/ 100.0/66.7
Random Baseline	0.0	0.0	0.0	0.0	25.0	25.0	25.0	25.0	50.0/50.0/50.0	50.0/50.0/50.0

Table 3: Results (%) of selected long-context LLMs on our benchmark. We report both the performance on 4k and 16k versions of all datasets. For SenHallu and AbsHallu, we sequentially exhibit precision, recall, and F1.

API calls in tackling complex tasks, we construct the **PrivateEval** dataset on the basis of the benchmark from Zan et al. (2022). In the dataset, LLMs are presented with API documents derived from private libraries and a partial code snippet. Then, the model needs to identify key API documents to complete the code snippet. To build the dataset, we manually alter the keywords within API documents of the Torchdata library following previous methods (Zan et al., 2022) on Monkey and BeatNum libraries. Furthermore, we adjust the number of the provided documents for length control.

3.4. Metrics

As shown in Table 2, we employ four metrics to evaluate different tasks in our benchmark, *i.e.*, accuracy for question answering and language modeling, concordance index (Shaham et al., 2023) for sorting, pass@1 (Chen et al., 2021) for code completion, and precision, recall, and F1 for hallucination detection. Notably, all the answers are generated freely, and any answer with a wrong format is treated as incorrect.

3.5. Benchmark Usage

The BAMBOO benchmark offers researchers a comprehensive platform to thoroughly evaluate the long text capabilities of LLMs. First, with multiple tasks covering the evaluation of different capacities, BAMBOO provides a direct comparison of different long-context LLMs’ comprehensive capacities. Meanwhile, the BAMBOO benchmark enables a detailed analysis of LLMs’ weaknesses in different aspects. Thus, more powerful long-context LLMs

can be built by making up for these deficiencies. Finally, the BAMBOO benchmark introduces varying length levels, enabling assessing the significance of the lengths of prompts and context windows. In total, BAMBOO can promote better handling of the multifaceted challenges and opportunities in the realm of long text modeling.

Users can utilize the BAMBOO benchmark with instructions provided in our project repository or instructions for your design. Then, users can evaluate the generations with our code.

4. Experiment

With the BAMBOO benchmark, we conduct experiments to evaluate the zero-shot long text modeling capability of long-context LLMs. Additionally, we further delve into key questions of long text modeling and discuss problems of long-context models.

4.1. Baselines

We select five instruction-tuned models with a context length of over 16k tokens as our baselines. For closed models available via API, we select gpt-3.5-turbo-16k² from OpenAI and Claude2-100K³ from Anthropic. For open-sourced models, we select Vicuna-7b-v1.5-16k (Zheng et al., 2023), Longchat-7b-v1.5-32k (Li et al., 2023a), and ChatGLM2-6b-32k (Du et al., 2022; Zeng et al., 2023)⁴.

⁴We abbreviate gpt-3.5-turbo as ChatGPT, Vicuna-7b-v1.5 as Vicuna, Claude2-100k as Claude2, Longchat-7b-v1.5 as Longchat, and ChatGLM2-6b as ChatGLM2 in the following.

We add a random baseline for each task. For text sorting and language modeling tasks, we randomly return the permuted sequences and pre-existing speakers in the dialogues, respectively. For hallucination detection and multi-choice QA datasets, we randomly select one choice from (true, false) or (A, B, C, D). For PrivateEval and AltQA, we directly set the random baseline to 0 due to the infinite search space for answers.

4.2. Overall Results

The overall results for the model performance on the BAMBOO benchmark are displayed in Table 3. Across different datasets and length levels, ChatGPT-16k consistently exhibits excellent performance, surpassing other LLMs in nearly all tasks, except a minor performance gap observed in the hallucination detection task. Conversely, the other models’ performances are often poor, even falling below performances of random baselines.

Moreover, models usually struggle with datasets with uncommon objectives and complex requirements, even the most powerful ChatGPT. In hallucination detection tasks where distinctions between the hallucinated and correct hypotheses are subtle, most LLMs can hardly capture those hallucinations, resulting in low precision. In addition, LLMs also underperform in ShowsSort and PrivateEval, highlighting their inability to aggregate, manipulate tools, and synthesize code. We infer that limited diversity of training data types plays a role in the subpar performance of LLMs.

Comparing performances across different length levels, a notable trend is that there is a discernible decrease in performance for most datasets and models as extending input length. Due to the incorrect generating format, the trend is more prominent for tasks with diverse outputs.

4.3. Key Questions of Long Text Modeling

To delve into the mechanism and influential factors affecting the long text modeling capacities of LLMs, our investigation centers on five key research questions. In this section, we conduct analytical experiments to address these questions⁵.

RQ1: Do LLMs with Long Context Windows Pay for Extension Tax?

We first discuss whether extending the LLMs’ context windows and fine-tuning LLMs on long texts harm performance over tasks with shorter

⁵We only report F1 score for hallucination detection tasks followingly.

texts, *i.e.*, “extension tax”. We compare ChatGPT-16k and Vicuna-16k to their short-context variants on selected tasks from BANBOO-4k and MMLU (Hendrycks et al., 2021).

As shown in Figure 1, ChatGPT and Vicuna exhibit disparate trends. Performances between ChatGPT and ChatGPT-16k over different tasks are almost identical. However, an “extension tax” is evident when replacing Vicuna with Vicuna-16k on MMLU. The longer training data and position interpolation detrimentally affect the ability of short-context tasks. Interestingly, Vicuna-16k demonstrates improved performance in BAMBOO-4k, suggesting that longer training data can enhance performance in medium-length tasks.




Figure 1: Results of long and short-context LLMs on BAMBOO-4k and MMLU.

RQ2: To What Extent Do LLMs Struggle Due to The Long Input?

In the long text scene, a critical question emerges: is the challenge associated with the length of the text or the inherent capacity to solve the task? In MeetingQA, PaperQA, and PrivateEval datasets, only concise evidence in the long text is sufficient to support the completion of tasks. Within these datasets, we make a comparison of the performance between inputting only evidence and the complete text.

Models	Input	MeetingQA		PaperQA		PrivateEval	
		16k	4k	16k	4k	16k	4k
ChatGPT-16k	evidence	74.0	78.0	76.0	79.0	31.6	31.6
	complete	72.0	75.0	75.0	78.0	21.7	21.1
Vicuna-16k	evidence	40.0	36.0	35.0	36.0	5.3	5.3
	complete	27.0	31.0	37.0	29.0	2.0	3.3

Table 4: Comparison of LLMs’ performances with the input of only evidence or the complete text.

Table 4 illustrates a general trend of improved performance across various datasets. Mislocating evidence within long text contributes to some errors. Nevertheless, the improvements are relatively minor, and Vicuna’s performance even declines in

Model	Position	ShowsSort		SenHallu		AbsHallu		MeetingPred		ShowsPred		MeetingQA		PaperQA	
		16k	4k	16k	4k	16k	4k	16k	4k	16k	4k	16k	4k	16k	4k
ChatGPT-16k	Pre-Ins	54.5	56.4	76.7	79.7	69.3	70.3	20.5	45.0	35.0	33.0	78.0	73.0	70.0	71.0
	Post-Ins	55.4	54.6	75.8	77.3	69.7	71.9	9.5	16.0	14.0	20.0	75.0	80.0	74.0	77.0
	Both-Ins	54.6	60.2	75.8	75.5	68.7	70.7	16.0	29.0	26.0	31.0	81.0	79.0	73.0	80.0
Vicuna-16K	Pre-Ins	42.8	53.9	66.7	70.5	66.7	66.4	5.0	24.0	6.0	16.0	6.0	32.0	33.0	24.0
	Post-Ins	51.4	51.4	67.8	73.7	66.9	66.7	13.0	2.0	0.0	2.0	21.0	27.0	24.0	38.0
	Both-Ins	52.4	53.4	67.9	66.7	66.7	69.0	12.0	8.0	0.0	4.0	24.0	29.0	32.0	29.0
LongChat-32k	Pre-Ins	51.5	53.8	66.7	67.6	66.2	66.7	12.0	37.0	4.0	21.0	14.0	37.0	28.0	29.0
	Post-Ins	49.8	48.9	66.7	66.9	66.7	66.7	19.0	19.0	1.0	8.0	34.0	33.0	30.0	39.0
	Both-Ins	50.3	53.8	66.7	68.7	67.1	67.6	18.0	18.0	4.0	8.0	32.0	34.0	24.0	32.0
ChatGLM2-32k	Pre-Ins	42.6	45.4	69.2	69.5	66.4	69.4	13.0	18.0	9.0	15.0	34.0	46.0	47.0	54.0
	Post-Ins	29.2	45.7	69.7	71.2	67.3	70.2	8.0	11.0	7.0	15.0	61.0	63.0	65.0	66.0
	Both-Ins	37.8	47.5	69.4	67.8	66.9	69.0	12.0	15.0	9.0	11.0	54.0	56.0	61.0	67.0

Table 5: Results of LLMs with three different instruction positions relative to content.

PaperQA. Additionally, the model consistently commits the same errors in both settings. Therefore, we infer that the awful performances of LLMs are predominantly attributed to subpar reasoning and coding ability rather than the mislocalization of evidence in the long text.

RQ3: How Do Instruction Positions Affect Long Text Modeling?

A recent study observed that positioning the instruction at the end of the input can enhance generation performance (Liu et al., 2023c). Therefore, we explore the influence of instruction placement in long text tasks. Specifically, we segment the complete prompt into two distinct components: "context" encompasses the long text, and "instruction" contains task descriptions, optional questions, hypotheses, etc. Following Liu et al. (2023c), We respectively position the instruction at the beginning, end, and both ends of the prompt, denoted as "Pre-Ins," "Post-Ins," and "Both-Ins." We present a comparison of their effects in Table 5.

To start with, when the instruction is exclusively placed at the beginning, a discernible decline in performance is commonly encountered as extending input lengths. Notably, the "Pre-Ins" configuration within BAMBOO-16k often exhibits problems related to forgetting instructions, which is potentially attributed to the long-range decay of attention scores of RoPE (Xiong et al., 2023).

Moreover, the optimal instruction positions vary depending on the datasets and models. "Pre-Ins" typically result in lower performance for tasks where the instructions contain content-relevant information, such as question answering and hallucination detection. Conversely, "Post-Ins" exhibit reversed trends, consistent with findings from previous research (Liu et al., 2023c). Furthermore, "Both-Ins" consistently yield sub-optimal results across most datasets, with only minor performance gaps compared to the best-performing strategies.

RQ4: Can LLMs Model Longer Text with Context Compression Methods?

Prior research has effectively leveraged context compression techniques such as retrieval and truncation to address long text tasks with pretrained language models (Gong et al., 2020; Zhang et al., 2022; Zhao et al., 2022). Thus, it holds promise to integrate these techniques with LLMs.

To verify whether context compression techniques can enhance long text modeling abilities of short-context LLMs, we employ three different methods for hallucination detection and question answering tasks. (1) *Truncation*: We truncate the initial 3000 tokens. (2) *Retrieval*: we partition the long text into segments with 256 tokens, representing the question and each segment into an embedding with openai's text-embedding-ada-002⁶. Then, we retrieve top-10 relevant segments to the question. (3) *Summarization*: We partition the input into chunks with 1024 tokens, sequentially summarize each chunk, and concatenate these summaries to form the final input.

Method	AQA	PQA	MQA	SHallu	AHallu
ChatGPT-16K	72.0	75.0	72.0	76.7	71.4
ChatGPT+Retrieval	69.5	79.0	78.0	78.3	71.2
ChatGPT+Truncation	71.5	64.0	68.0	74.1	70.0
ChatGPT+Summarization	51.5	63.0	67.0	75.0	70.0
Vicuna-16K	25.0	37.0	36.0	66.9	66.7
Vicuna+Retrieval	54.5	33.0	39.0	67.6	66.7
Vicuna+Truncation	55.8	33.0	32.0	66.7	66.7
Vicuna+Summarization	43.5	28.0	25.0	70.9	66.9

Table 6: Comparison of short-context LLMs with context compression techniques and long-context LLMs on selected tasks in BAMBOO-16k. We abbreviate AltQA, PaperQA, MeetingQA, Senhallu, and AbsHallu as AQA, PQA, MQA, SHallu, and AHallu, respectively.

The effect of different compression techniques is shown in Table 6. Evidently, retrieval-augmented

⁶<https://platform.openai.com/docs/guides/embeddings>




Figure 2: Logarithmic value of attention scores in three settings of evidence positions.

LLMs can achieve comparable or superior performance to long-context LLMs, aligning with the findings in Xu et al. (2023). However, truncation and summarization methods often perform poorly due to the omission of substantial relevant information.

RQ5: Why “Lost in The Middle” Happens?

Previous study (Liu et al., 2023a) has observed that optimal performance is often attained when evidence information is positioned at the beginning or end of the input. To validate this phenomenon, we conducted experiments by placing evidence at the beginning or end of the long document (denoted as “Pre-Evi” and “Post-Evi”, respectively). As illustrated in Table 7, despite disrupting the paper’s original discourse structure, we consistently observed similar or improved performance when the evidence was situated at either end of the paper. Moreover, the performance gains were more substantial for longer input sequences.

To further investigate the rationale behind these remarkable properties, we analyze the attention maps. We compute the average attention scores for predictions across all heads and layers, and then partition the sequence into chunks of 32 tokens. Subsequently, we employ curve fitting to model the average scores of each chunk, as shown in Figure 2. Regardless of the location of the evidence within the input, we observe a U-shaped curve with higher attention allocated to tokens positioned at the two ends of the content. Consequently, it becomes evident that LLMs tend to utilize information located at the beginning or end of the input more effectively.

Model	PaperQA-16k			PaperQA-4k		
	Pre-Evi	Original	Post-Evi	Pre-Evi	Original	Post-Evi
ChatGPT-16k	78.0	75.0	77.0	77.0	78.0	79.0
Vicuna-16k	44.0	37.0	39.0	29.0	29.0	30.0

Table 7: Impact of evidence position.

4.4. Discussions

Finally, we give discussions about the problems of LLMs over long texts.

- **A severe problem is the cataclysmic forgetting of instructions.** Frequently, the generating responses fail to align with the task demands. During the generation process, the instruction may be forgotten or even not understood at all, especially when the instructions are put at the beginning of long inputs. Smaller LLMs are more prone to adhering to the instructions, possibly owing to limitations in memorization. Thus, more robust long instruction datasets are required to enhance the instruction-following ability of long-context models

- **LLMs are prone to format errors in long text tasks.** In certain cases, the predictions of LLMs may be presented in an informal format, though they convey the same meaning as the answers. The models may offer additional explanations and repetitions or adopt different formats. Thus, they can not be identified by the evaluation code. To alleviate the issue, we suggest reducing the format problems during the fine-tuning or RLHF stages and employing post-process techniques to rewrite the responses.

- **LLMs exhibit poor performance beyond the long texts.** Provided evidence or condensed texts, LLM may still generate false responses and make the same fault as inputting the complete long texts. This error can be attributed to insufficient ability, across various lengths and tasks. Thus, it is imperative to augment the overall capabilities of LLMs, rather than concentrating solely on the context of long text domains.

- **LLMs perform poorly on uncommon tasks.** Notably, LLMs generally exhibit poor performance on tasks that are uncommon, *e.g.*, text sorting and code completion, and even underperform by random baselines. Most open-sourced LLMs were fine-tuned on multi-turns of conversation, question answering, and summarization tasks. The limited diversity in fine-tuning data hinders LLMs’ generalization to uncommon long text tasks. We believe that broadening the variety of tasks and domains on fine-tuning data can better comprehensively enhance LLMs’ long text modeling capacities.

5. Conclusion

We propose BAMBOO, a benchmark for comprehensively evaluating the long text modeling capabilities of LLMs. BAMBOO consists of five tasks with two length levels, enabling the evaluation of LLMs' main capacities across various dimensions and domains. Based on the evaluation of several long-context models on BAMBOO, we give an overall analysis of the performances of different models and tasks. Additionally, we analyze key questions of long text modeling, discuss the problems of long-context LLMs, and suggest directions for improving long text comprehension abilities. We believe BAMBOO has the potential to serve as a valuable tool for analyzing the comprehensive capacities and promoting the enhancement of long text modeling abilities of LLMs in the future.

6. Acknowledgments

This work was partially supported by National Natural Science Foundation of China under Grant No. 62222215, Beijing Natural Science Foundation under Grant No. 4222027 and L233008. Xin Zhao is the corresponding author.

7. Bibliographical References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *CoRR*, abs/2004.05150.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew R. Gormley. 2023. Unlimiformer: Long-range transformers with unlimited length input. *CoRR*, abs/2305.01625.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating large language models trained on code. *CoRR*, abs/2107.03374.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. Extending context window of large language models via positional interpolation. *CoRR*, abs/2306.15595.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *CoRR*, abs/1904.10509.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2978–2988. Association for Computational Linguistics.
- Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. In *NeurIPS*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Zican Dong, Tianyi Tang, Junyi Li, and Wayne Xin Zhao. 2023. A survey on long text modeling with transformers. *CoRR*, abs/2302.14502.
- 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 the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335.
- Shahriar Golchin and Mihai Surdeanu. 2023. Time travel in llms: Tracing data contamination in large language models. *CoRR*, abs/2308.08493.
- Hongyu Gong, Yelong Shen, Dian Yu, Jianshu Chen, and Dong Yu. 2020. Recurrent chunking mechanisms for long-text machine reading comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10,*

- 2020, pages 6751–6761. Association for Computational Linguistics.
- Albert Gu, Karan Goel, and Christopher Ré. 2022. Efficiently modeling long sequences with structured state spaces. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Ankit Gupta, Albert Gu, and Jonathan Berant. 2022. Diagonal state spaces are as effective as structured state spaces. In *NeurIPS*.
- Gautier Izacard and Edouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 874–880. Association for Computational Linguistics.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. Longllm1ngua: Accelerating and enhancing llms in long context scenarios via prompt compression. *CoRR*, abs/2310.06839.
- kaiokendev. 2023. Things I’m learning while training superhot.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6769–6781. Association for Computational Linguistics.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rns: Fast autoregressive transformers with linear attention. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5156–5165. PMLR.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Dacheng Li, Rulin Shao, Anze Xie, Ying Sheng, Lianmin Zheng, Joseph E. Gonzalez, Ion Stoica, Xuezhe Ma, and Hao Zhang. 2023a. How long can open-source llms truly promise on context length?
- Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023b. Halueval: A large-scale hallucination evaluation benchmark for large language models. *CoRR*, abs/2305.11747.
- Xinnian Liang, Bing Wang, Hui Huang, Shuangzhi Wu, Peihao Wu, Lu Lu, Zejun Ma, and Zhoujun Li. 2023. Unleashing infinite-length input capacity for large-scale language models with self-controlled memory system. *CoRR*, abs/2304.13343.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023a. Lost in the middle: How language models use long contexts. *CoRR*, abs/2307.03172.
- Xiaoran Liu, Hang Yan, Shuo Zhang, Chenxin An, Xipeng Qiu, and Dahua Lin. 2023b. Scaling laws of rope-based extrapolation. *CoRR*, abs/2310.05209.
- Yijin Liu, Xianfeng Zeng, Fandong Meng, and Jie Zhou. 2023c. Instruction position matters in sequence generation with large language models. *CoRR*, abs/2308.12097.
- Xuezhe Ma, Xiang Kong, Sinong Wang, Chunting Zhou, Jonathan May, Hao Ma, and Luke Zettlemoyer. 2021. Luna: Linear unified nested attention. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 2441–2453.
- Harsh Mehta, Ankit Gupta, Ashok Cutkosky, and Behnam Neyshabur. 2023. Long range language modeling via gated state spaces. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell,

- Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- Charles Packer, Vivian Fang, Shishir G Patil, Kevin Lin, Sarah Wooders, and Joseph E Gonzalez. 2023. Memgpt: Towards llms as operating systems. *CoRR*, abs/2310.08560.
- Arka Pal, Deep Karkhanis, Manley Roberts, Samuel Dooley, Arvind Sundararajan, and Siddhartha Naidu. 2023. Giraffe: Adventures in expanding context lengths in llms. *CoRR*, abs/2308.10882.
- Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, Kranthi Kiran G. V., Xuzheng He, Haowen Hou, Przemyslaw Kazienko, Jan Kocon, Jiaming Kong, Bartlomiej Kopyta, Hayden Lau, Krishna Sri Ipsit Mantri, Ferdinand Mom, Atsushi Saito, Xiangru Tang, Bolun Wang, Johan S. Wind, Stanislaw Wozniak, Ruichong Zhang, Zhenyuan Zhang, Qihang Zhao, Peng Zhou, Jian Zhu, and Ruijie Zhu. 2023a. RWKV: reinventing rnns for the transformer era. *CoRR*, abs/2305.13048.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2023b. Yarn: Efficient context window extension of large language models. *CoRR*, abs/2309.00071.
- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A. Smith, and Lingpeng Kong. 2021. Random feature attention. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Ofir Press, Noah A. Smith, and Mike Lewis. 2022. Train short, test long: Attention with linear biases enables input length extrapolation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. 2020. Compressive transformers for long-range sequence modelling. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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 the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Toronto, Canada. Association for Computational Linguistics.
- Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. 2021. Efficient content-based sparse attention with routing transformers. *Trans. Assoc. Comput. Linguistics*, 9:53–68.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. Code llama: Open foundation models for code. *CoRR*, abs/2308.12950.
- Tevan Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilic, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muenighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, and et al. 2022. BLOOM: A 176b-parameter open-access multilingual language model. *CoRR*, abs/2211.05100.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *CoRR*, abs/2104.09864.
- Simeng Sun, Kalpesh Krishna, Andrew Mattarella-Micke, and Mohit Iyyer. 2021. Do long-range language models actually use long-range context? In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 807–822. Association for Computational Linguistics.
- Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. 2023a. Retentive network: A successor to

- transformer for large language models. *CoRR*, abs/2307.08621.
- Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shaohan Huang, Alon Benham, Vishrav Chaudhary, Xia Song, and Furu Wei. 2023b. A length-extrapolatable transformer. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 14590–14604. Association for Computational Linguistics.
- Tianyi Tang, Hongyuan Lu, Yuchen Eleanor Jiang, Haoyang Huang, Dongdong Zhang, Wayne Xin Zhao, and Furu Wei. 2023. Not all metrics are guilty: Improving NLG evaluation with LLM paraphrasing. *CoRR*, abs/2305.15067.
- Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. 2020. Sparse sinkhorn attention. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9438–9447. PMLR.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288.
- Szymon Tworkowski, Konrad Staniszewski, Mikolaj Pacek, Yuhuai Wu, Henryk Michalewski, and Piotr Milos. 2023. Focused transformer: Contrastive training for context scaling. *CoRR*, abs/2307.03170.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. Is chatgpt a good NLG evaluator? A preliminary study. *CoRR*, abs/2303.04048.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *CoRR*, abs/2006.04768.
- Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. 2023b. Augmenting language models with long-term memory. *CoRR*, abs/2306.07174.
- Yuhuai Wu, Markus Norman Rabe, DeLesley Hutchins, and Christian Szegedy. 2022. Memorizing transformers. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, Madian Khabsa, Han Fang, Yashar Mehdad, Sharan Narang, Kshitiz Malik, Angela Fan, Shruti Bhosale, Sergey Edunov, Mike Lewis, Sinong Wang, and Hao Ma. 2023. Effective long-context scaling of foundation models. *CoRR*, abs/2309.16039.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2023. Retrieval meets long context large language models. *arXiv preprint arXiv:2310.03025*.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big bird: Transformers for longer sequences. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Zhiyuan Liu, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. GLM-130B: an open bilingual

pre-trained model. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Hassan Awadallah, Dragomir R. Radev, and Rui Zhang. 2022. Summⁿ: A multi-stage summarization framework for long input dialogues and documents. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1592–1604. Association for Computational Linguistics.

Wayne Xin Zhao, Jing Liu, Ruiyang Ren, and Ji-Rong Wen. 2022. Dense text retrieval based on pretrained language models: A survey. *CoRR*, abs/2211.14876.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A survey of large language models. *CoRR*, abs/2303.18223.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *CoRR*, abs/2306.05685.

Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan Cotterell, and Mrinmaya Sachan. 2023. Recurrentgpt: Interactive generation of (arbitrarily) long text. *CoRR*, abs/2305.13304.

8. Language Resource References

Chenxin An, Shansan Gong, Ming Zhong, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2023. L-eval: Instituting standardized evaluation for long context language models. *CoRR*, abs/2307.11088.

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. 2023. Longbench: A

bilingual, multitask benchmark for long context understanding. *CoRR*, abs/2308.14508.

Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2022. Summscreen: A dataset for abstractive screenplay summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 8602–8615. Association for Computational Linguistics.

Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 615–621. Association for Computational Linguistics.

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 4599–4610. Association for Computational Linguistics.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Luyang Huang, Shuyang Cao, Nikolaus Nova Parulian, Heng Ji, and Lu Wang. 2021. Efficient attentions for long document summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1419–1436. Association for Computational Linguistics.

Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel R. Bowman. 2022. Quality: Question answering with long input texts, yes! In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language*

Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022, pages 5336–5358. Association for Computational Linguistics.

Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. 2023. Zeroscrolls: A zero-shot benchmark for long text understanding. *CoRR*, abs/2305.14196.

Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, and Omer Levy. 2022. SCROLLS: standardized comparison over long language sequences. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 12007–12021. Association for Computational Linguistics.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380. Association for Computational Linguistics.

Daoguang Zan, Bei Chen, Zeqi Lin, Bei Guan, Yongji Wang, and Jian-Guang Lou. 2022. When language model meets private library. In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 277–288. Association for Computational Linguistics.

Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir R. Radev. 2021. Qmsum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 5905–5921. Association for Computational Linguistics.