

# Training with “Paraphrasing the Original Text” Teaches LLM to Better Retrieve in Long-Context Tasks

Yijiong Yu<sup>1</sup>, Yongfeng Huang<sup>1 2 3</sup>, Zhixiao Qi<sup>1</sup>, Zhe Zhou<sup>1</sup>

<sup>1</sup>Tsinghua University

<sup>2</sup>Zhongguancun Laboratory

<sup>3</sup>Institute for Precision Medicine, Tsinghua University

{yuyj22, qzx21, zhouzhe20}@mails.tsinghua.edu.cn, yfhuang@tsinghua.edu.cn

## Abstract

As Large Language Models (LLMs) continue to evolve, more are being designed to handle long-context inputs. Despite this advancement, most of them still face challenges in accurately handling long-context tasks, often showing the “lost in the middle” issue. We identify that insufficient retrieval capability is one of the important reasons for this issue. To tackle this challenge, we propose a novel approach to design training data for long-context tasks, aiming at augmenting LLMs’ proficiency in extracting key information from long context. Specially, we incorporate an additional part named “paraphrasing the original text” when constructing the answer of training samples and then fine-tuning the model. Experimenting on LongBench and NaturalQuestions Multi-document-QA dataset with models of Llama and Qwen series, our method achieves an improvement of up to 8.48% and 4.48% in average scores, respectively, showing effectiveness in improving the model’s performance on long-context tasks.

## 1 Introduction

Large Language Models (LLMs) have recently emerged as top performers in a wide range of natural language processing tasks. Nevertheless, they are usually trained on segments of text of a fixed length, which means they have a predetermined context window size. Typically, their performance tends to decline markedly when the input text exceeds the context window size.

Some classic works like Yarn (Peng et al. 2023), LongChat (Li et al. 2023a) and LongAlpaca (Chen et al. 2023c) have explored ways to make short-context LLMs better adaptable to long-context tasks. Based on these, recently, an increasing number of powerful LLMs have been equipped with the capability to handle long contexts, such as Mistral (Jiang et al. 2023a) and Qwen2 (Yang et al. 2024), which can handling the text length of 32k or longer.

However, while LLMs have made remarkable progress in handling long-context tasks, as shown in many evaluations (Li et al. 2023b; An et al. 2023; Wang et al. 2024), they still lack satisfactory accuracy, and often suffer from a severe “lost in the middle” problem (Liu et al. 2024) when the




Figure 1: Our method adds “paraphrasing the original text” to the training samples.

context is getting longer or more complex. “Lost in the middle” refers to the phenomenon where the model’s utilization of the context significantly weakens when the key information is located in the middle of the context. This limitation hinders the further application of LLMs in long-context scenarios.

Long-context tasks can be divided into two categories: short-dependency and long-dependency (Li et al. 2023b). The former means only a small part of the long context is truly needed for the task, while the latter means most parts are needed. Because short-dependency tasks are usually more common and easier to study, for simplicity, in this paper, we start with the short-dependency tasks and mainly focus on multi-document-QA, which is one of the most representative long-context tasks and also easy to be constructed.

In short-dependency long-context tasks, the useful information in the context is sparse, thus the task can be regarded as a composite task, which can be split to “first retrieval, then another follow-up task” process, while the retrieval step is abstract and implicit but necessary. However, we find LLMs usually perform much worse in such composite tasks, even if they are good at each sub-task individually.

In this case, it is natural to think that using CoT (Wei et al. 2022) prompt could help. Yet, in our experiments, despite using CoT-like prompt, we find many LLMs still perform poorly, which generate wrong answers just in the first step, retrieval, even though some of them can perfectly pass “Needle in a Haystack” test, which represents qualified retrieval ability. While in the other hand, if they retrieve correctly, then the follow-up task becomes easier and the final answer

is usually right.

We posit that this is either because these models inherently have weak retrieval capabilities, or because their retrieval abilities cannot be fully activated by just designing prompts, which hinders them from accurately locating key information in the long context. This prompts us to seek a more effective method to enhance the accuracy of models in retrieval-based composite tasks, through teaching them to better retrieve.

Thus, we propose a method based on fine-tuning with specially designed samples to enhance and activate LLMs' retrieval capability over long contexts. Specifically, for the goal of explicitly separating and highlighting the "retrieval" step, when designing answers of a long-context QA sample, different from normal ways only designing a brief answer, we incorporate an additional part named "original text paraphrasing", which paraphrases or repeats the original text of the sentences in the context containing relevant information required for answering, as shown in Figure 1. This part corresponds to a direct "retrieval" operation, while maintaining the answers' quality and coherence, which aims to not only enhance the model's retrieval capability but also teach the model to use its inherent retrieval capability more actively.

With the help of GPT-4 (OpenAI 2023), we automatically construct a dataset consisting of thousands of training samples. Through evaluate models fine-tuned on the training samples designed by us, we prove our method can generally benefit the model's performance across various long-context tasks (not just retrieval or QA tasks). Besides, our method only need a light-weight dataset and a cost-effective fine-tuning stage, which is very applicable.

Our main contributions are summarized as follows:

1. We find that LLMs do not guarantee the full utilization of their retrieval capability in composite long-context tasks, even if their individual capabilities in simple tasks are strong.
2. We propose the "original text paraphrasing" approach, which explicitly isolates the "retrieval" step, to construct training samples, aiming at teaching LLMs to better use their retrieval capability.
3. Through fine-tuning and then evaluating, we prove our method can improve LLMs' retrieval capability as well as overall long-context performance.

## 2 Related Work

There have been many studies that aim to improve the long-context performance of LLMs and address "lost in the middle" issue, which can be summarized into the following four aspects:

### Input Context and Prompt

It is well known that using an appropriate prompt such as CoT (Wei et al. 2022) can significantly improve the model's performance in a zero-shot way. A common method is to prompt the model to first give relevant evidence and then answer the question. This way can guide the model to first retrieve relevant information from the long text and then give

answers based on this information. For example, the Claude-2.1 team proposed that adding the sentence "Here is the most relevant sentence in the context" (Anthropic 2023) can improve the long-context QA accuracy from 27% to 98%.

In addition, reorganizing the long input context is also effective. For example, LongLLMLingua (Jiang et al. 2023b) compresses long contexts, and Attention Sorting (Peysakhovich and Lerer 2023) improves the model's utilization of contexts through reordering the input documents.

### Training Data

Some works attempt to fine-tune models with long-context datasets to improve their long-context performance. Ziya-Reader (He et al. 2023) proposes an attention strengthening method, designing the answer with 3 parts: "question repetition", "index prediction" and "answer summarization". FILM (An et al. 2024) proposes "Information-Intensive Training" to teach the model that any position in a long context can contain crucial information.

### Position Embedding

The remote attenuation characteristic of ROPE (Su et al. 2021) is often considered a significant factor contributing to the "lost in the middle" phenomenon. Thus Chen et al. propose "Attention Buckets", which sets up three different frequencies of ROPE and integrates them together, filling in the troughs of ROPE's remote attenuation curve. MSPOE (Zhang et al. 2024) defines "position-aware" scores for each attention head and then assigns different ROPE interpolation (Chen et al. 2023a) factors to each attention head, which significantly improves the response accuracy of Llama2 (Touvron et al. 2023) in long-context scenarios.

### Attention Weights

Hsieh et al. claim that the biased attention distribution of the model is the direct cause of its poor performance on long context, thus they decompose attention into two components, determined by position and semantics, respectively, and then calibrate the attention by eliminating the position-related component. Gao et al. analyze the attention matrix to directly eliminate the less important parts of attention and compensate the important parts, thereby making the attention distribution more rational.

## 3 Method

### Models Often Fail to Fully Utilize Retrieval Ability

We take Qwen1.5-4b-Chat (Yang et al. 2024) for example, which is a model with 32k context window. As shown in Figure 4, it can nearly perfectly pass "Needle in a Haystack" (gkamradt 2023) test within 32k length, which is a task requiring LLM to retrieve a sentence containing relevant information (i.e. the "needle") from a long context made up of a large amount of irrelevant information, when the "needle" is inserted at various positions of the context. Passing this test represents a long-context LLM has the ability to directly retrieve information at anywhere of the context.

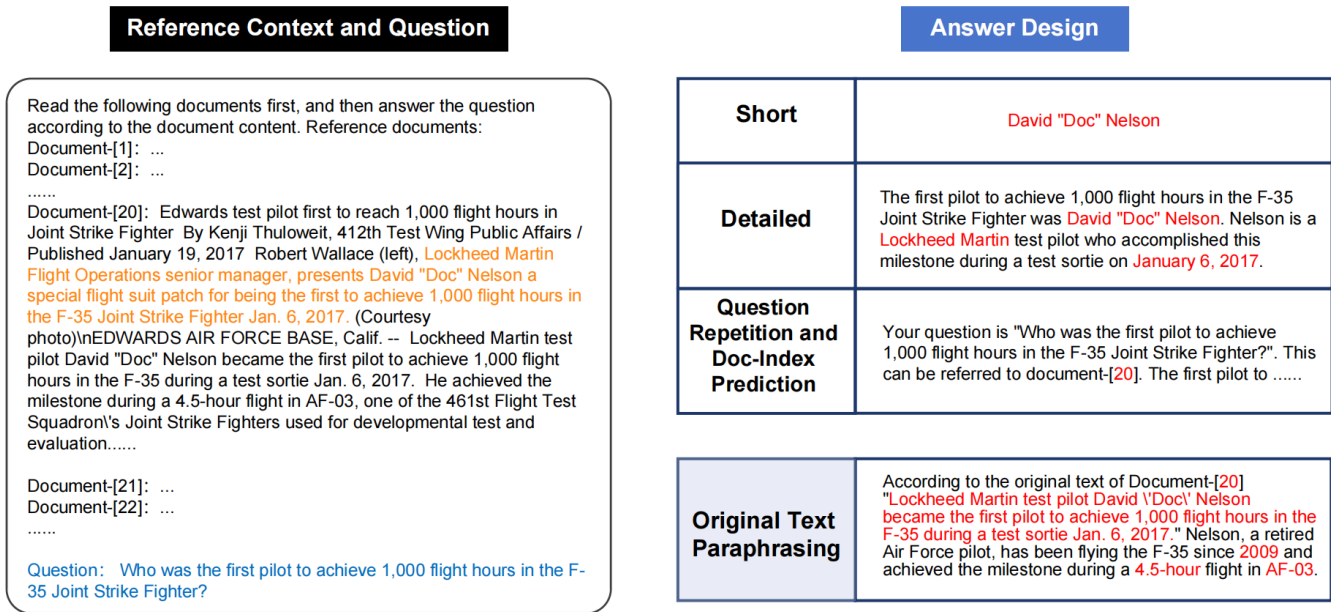


Figure 2: An example of different answer design methods for a multi-doc-QA sample. In the context and answers, key information for answering the question are highlighted.

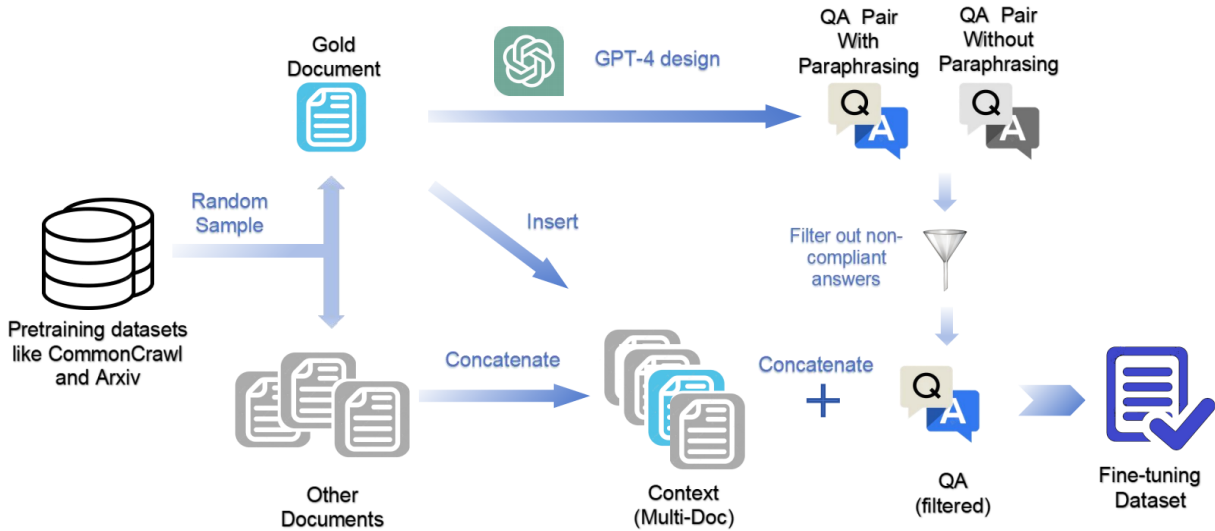


Figure 3: The pipeline of constructing our datasets with multi-doc-QA samples.

However, as for composite tasks such as Natural Questions Multi-Document-QA (Liu et al. 2024) and passage retrieval tasks in Longbench (Bai et al. 2023b), it performs much worse. Multi-Document-QA requires model to answer the question based on the context consisting of 20 documents with only 1 document containing truly useful information (i.e. the gold document), and passage retrieval task requires models to answer the corresponding document number given the abstract of one of 30 documents. They can both be considered as composite tasks that require both document understanding and retrieval. As shown in Figure 5, it achieves an accuracy of 82% in multi-doc-QA when the

context only consists of the gold document, which proves its ability of understanding the document and question is fine. But when the context consists of more documents with the gold document placed in the middle, the accuracy decreases rapidly.

We have tried to awaken the retrieval capability of the model through CoT-like (Wei et al. 2022) prompt (see Appendix A for detailed prompts). However, as shown in Table 1, even though CoT-like (Wei et al. 2022) prompt raises the accuracy in passage retrieval to some extent, the accuracy still cannot be considered satisfactory. As for Multi-Document-QA, it even cannot help.




Figure 4: Qwen1.5-4b-Chat can nearly perfectly pass “Needle in a Haystack” test. The x-axis represents the length of the context, and the y-axis represents the position of the “needle” in the context.




Figure 5: The accuracy of Qwen1.5-4b-Chat in multi-doc-QA task drops rapidly, as the number of documents in the context grows, with the gold document always placed in the middle position of the context.

We conclude that LLMs may not fully utilize their retrieval ability in some composite tasks, even though they perform well in simple retrieval tasks. And this issue cannot be effectively addressed by just designing more refined prompts. Therefore, to make LLMs utilize their retrieval ability more effectively in long-context tasks, we decided to fine-tuning them.

### Original Text Paraphrasing

We have identified that an accurate “retrieval” step is crucial for completing the whole task. To explicitly isolate and emphasize the “retrieval” process, while also ensuring that the retrieved content aids in the better completion of subsequent tasks, when constructing long-context training samples’ answers, we design a “paraphrasing the original text” task at the beginning part of the answer, and the remaining part of the answer is crafted based on the question and the paraphrased original text to provide a final response. In “paraphrasing the original text” task, we extract and then repeat the original text of the most crucial parts (usually several sentences directly related to answering the question) of the context.

Figure 2 illustrates how we design an multi-doc-QA sam-

Task	Prompt	Acc
multi-doc-QA w/ 20 docs	default	47.20
multi-doc-QA w/ 20 docs	CoT	46.48
passage retrieval	default	13.00
passage retrieval	CoT	26.00

Table 1: Performance of Qwen1.5-4b-Chat on Natural Questions Multi-Document-QA and passage\_retrieval\_en of LongBench with different types of prompts.

ple with “original text paraphrasing”, which is compared with other methods: “short” and “detailed” are traditional ways used by Multi-passage-QA-from-Natural-Questions (Together 2023), “question repeating and doc-index prediction” is proposed by Ziya-Reader (He et al. 2023). The most crucial part in the reference documents for answering the question has been highlighted, and correspondingly, the words or sentences in the answer that pertain to this part have also been highlighted. It can be clearly seen that our design is capable of encompassing the maximum amount of key information, not only emphasizing the retrieval process to the greatest extent but also enabling the subsequent generation to directly and effectively utilize the retrieved information.

### Dataset Construction

We leverage the power of GPT-4 (gpt-4-0613) (OpenAI 2023) to generate high-quality in-context question-answer pairs according to our requirements. Figure 3 shows how we build multi-doc-QA samples with “original text paraphrasing”. For English dataset, we randomly select about 100k samples from the CommonCrawl dataset (Patel 2020) as reference documents for constructing the context, and for Chinese, Wudao 200GB open-source Chinese corpus (BAAI 2023) is used. Then, we let every 100 reference documents form a group, and take the first document from each group as the gold document, based on which, GPT-4 (OpenAI 2023) designs a question-answer pair under the prompt that the relevant original text must be provided in the answer, as follows:

Document:  
 {document content}  
 Please design an question answer pair base on the document.  
 The answer of the QA pair must start with “According to the original text ‘.....’”, first give the relevant original text in the reference content, and then answer the question in detail.

If the designed answer does not start with “According to the original text ‘.....’”, we let GPT-4 generate again until it meets our requirements.

The remaining documents are distractors. To control the context length not exceeding 32k, we randomly discard a certain number of distractor documents. Then, all the documents in the group are randomly shuffled and concatenated

to form a single context containing these documents, with each document preceded by its corresponding serial number. The question designed by GPT-4 is appended at the end of the context, with the answer designed by GPT-4 serving as the ground-truth answer. In addition, to make the answer more complete, we add the relevant document’s number before the “original text paraphrasing” part, like “According to the original text of document [5] ‘.....’”, so that the answer contains both “document index predicting” and “original text paraphrasing” tasks.

For constructing single-doc-QA samples, we use papers from Arxiv and CNKI as the context, and for each paper, we randomly select one paragraph based on which GPT-4 designs a QA-pair. Other operations are the sample as multi-doc-QA.

In fact, we originally used gpt-3.5-turbo, but when transferred to gpt-4, we found the overall performance was not improved much. Therefore, although intuitively using a more advanced LLM should be better, our method does not necessitate a very strong LLM or a specific version, as long as it can follow the format requirement in most cases.

Following LongAlpaca dataset(Chen et al. 2023c), to maintaining models’ generalization, we also add 2,000 short-form instruction samples, which are sampled from Alpaca (Taori et al. 2023) dataset and Llama2-chinese dataset (FlagAlpha 2023), into our dataset. Finally, our dataset comprises 2,000 short-form samples and 3,000 long-context samples with “original text paraphrasing” that we synthesized, with an equal number of Chinese and English samples. The length of the long-context samples varies from 8k to 32k and 80% of them are multi-document-QA samples while the rest is single-document-QA.

## 4 Experiments

### Implementation Details

We use models in Qwen series (Bai et al. 2023a) and the Chinese version of Llama series (Touvron et al. 2023), which support both English and Chinese, as our base models for fine-tuning on our dataset.

Using the Qlora (Detmeters et al. 2023) training method, we quantized the base model to int4 and fine-tuned all the linear layers of the models, with the following parameters: lora rank is 16, lora alpha is 64, global batch size is 16 and learning rate starts from 1e-5 with linearly decreasing. During fine-tuning, only the answer part was set as labels and involved in the loss calculation. Each model was trained for 1 epoch on a single A100 GPU, taking about 10 hours.

As an ablation experiment, we removed the “original text paraphrasing” part in the answers of our dataset while keeping other parts unchanged, and again fine-tuned the model as a contrast (SFT w/o ours). We also construct a dataset with the same size using Ziya-Reader’s methods (He et al. 2023) as a comparable baseline.

In evaluation, we use greedy decoding strategy for all the models and all the datasets to ensure reproducibility, and the model precision is set to bfloat16.

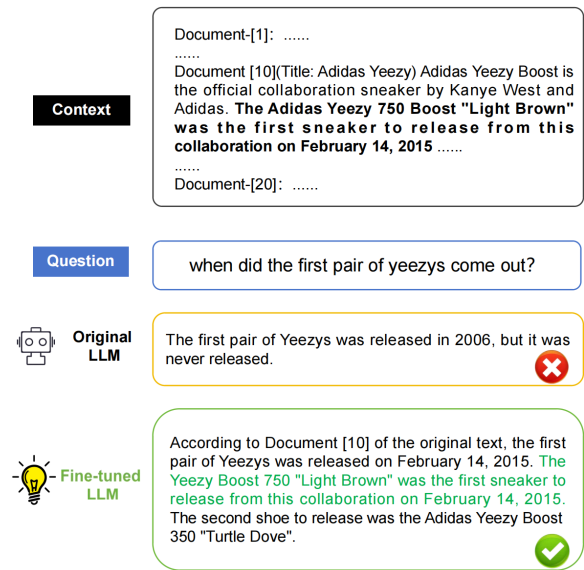


Figure 6: A case showing how the fine-tuned model (Qwen1.5-4b-Chat) gives a better answer in multi-doc-QA task. The key sentence is highlighted.

### Evaluation

**Evaluation on Long-context Tasks** We evaluated our model on LongBench (Bai et al. 2023b), a commonly used bilingual and multi-task dataset for long-context evaluation. It contains 14 English tasks and 5 Chinese tasks, across 5 task types: multi-doc-QA, single-doc-QA, synthetic tasks, summarization and few-shot learning. For each task type, to ensure a balanced sample size for both languages, we choose a English sub-dataset and a Chinese sub-dataset (specific dataset selection is in appendix A). Every sub-dataset contains 200 samples.

We evaluate our models on these 5 types of long-context tasks. For synthetic (i.e. passage retrieval) and few-shot tasks, we use exact-match to calculate accuracy and for summarization tasks, we use rouge-L. For multi-doc-qa and single-doc-qa tasks, we also use accuracy (see appendix A for how to calculate accuracy for QA). To avoid OOM, the maximum input length is set to 16k for all models.

The evaluation results are shown in Table 2. Our method achieves the best average performance, and is particularly outstanding on the retrieval tasks (Synthetic), even if our dataset does not include an specialized retrieval task but just dissociates the implicit retrieval process in in-context QA tasks into an explicit process, which proves that adding “original text paraphrasing” to the training samples does significantly improve the model’s retrieval capability.

However, the improvement on single-document-QA tasks is unstable, which may be due to the relatively short sample length of such tasks in Longbench (about 5k in average length), making it difficult to consistently reflect the model’s long-context capability.

**Evaluation on “Lost in the middle” Issue** The “lost in the middle” phenomenon (Liu et al. 2024) represents

Models	MultiDoc	SingleDoc	Synthetic	Summurization	FewShot	AVG
Qwen1.5-4b-Chat	65.06	<b>77.70</b>	14.75	19.27	<b>57.00</b>	46.76
Qwen1.5-4b-Chat SFT w/ ours	<b>65.39</b>	72.22	<b>62.00</b>	19.57	<b>57.00</b>	<b>55.24</b>
Qwen1.5-4b-Chat SFT w/o ours	65.28	71.86	12.75	<b>19.84</b>	56.50	45.25
Qwen1.5-4b-Chat SFT w/ Ziya	65.32	76.28	25.75	19.50	56.88	48.75
Qwen2-7b-Instruct	60.89	80.32	63.00	19.62	60.00	56.77
Qwen2-7b-Instruct SFT w/ ours	65.80	82.17	<b>65.00</b>	<b>20.34</b>	61.50	<b>58.96</b>
Qwen2-7b-Instruct SFT w/o ours	<b>66.49</b>	82.12	56.75	16.98	61.75	56.82
Qwen2-7b-Instruct SFT w/ Ziya	65.42	<b>82.40</b>	61.50	19.88	<b>62.50</b>	58.34
Llama3-8b-chinese-Chat	64.04	79.20	80.75	18.56	50.62	58.63
Llama3-8b-chinese-Chat SFT w/ ours	65.22	<b>80.92</b>	<b>97.75</b>	<b>21.17</b>	<b>57.38</b>	<b>64.39</b>
Llama3-8b-chinese-Chat SFT w/o ours	<b>66.21</b>	80.62	97.00	20.64	57.12	64.32
Llama3-8b-chinese-Chat SFT w/ Ziya	63.80	80.90	68.75	19.44	51.75	56.93

Table 2: Performance of different training methods with different models on LongBench (Bai et al. 2023b).

Models	1st	5th	10th	15th	20th	AVG
Qwen1.5-4b-Chat	<b>57.60</b>	46.60	45.20	43.80	42.80	47.20
Qwen1.5-4b-Chat SFT w/ ours	55.40	<b>49.40</b>	<b>60.00</b>	<b>54.80</b>	38.80	<b>51.68</b>
Qwen1.5-4b-Chat SFT w/o ours	57.00	47.00	48.00	43.60	44.80	48.08
Qwen1.5-4b-Chat SFT w/ Ziya	52.20	46.60	44.60	43.00	<b>51.00</b>	47.47
Qwen2-7b-Instruct	72.20	56.00	<b>55.00</b>	55.40	<b>57.40</b>	59.20
Qwen2-7b-Instruct SFT w/ ours	<b>78.40</b>	<b>59.40</b>	<b>55.00</b>	52.20	56.60	<b>60.32</b>
Qwen2-7b-Instruct SFT w/o ours	70.20	57.80	53.40	53.40	51.40	57.24
Qwen2-7b-Instruct SFT w/ Ziya	75.40	58.80	54.20	<b>56.00</b>	54.60	59.80
Llama3-8b-chinese-Chat	69.00	61.80	61.20	58.40	64.00	62.88
Llama3-8b-chinese-Chat SFT w/ ours	<b>71.20</b>	<b>63.60</b>	<b>64.60</b>	60.80	64.60	<b>64.96</b>
Llama3-8b-chinese-Chat SFT w/o ours	69.00	61.40	63.40	<b>62.00</b>	<b>67.80</b>	64.72
Llama3-8b-chinese-Chat SFT w/ Ziya	69.80	60.60	61.40	58.80	66.40	63.40

Table 3: Performance of different methods with different models on NaturalQuestions Multi-Doc-QA (Liu et al. 2024) when the gold document is placed at 5 different positions (at 1st, 5th, 10th, 15th, 20th with 20 documents in total).

a prevalent issue for LLMs when processing long-context tasks, signifying a low efficiency in utilizing information in the middle of the context. We use the multi-doc-QA datasets provided by “lost in the middle” research (Liu et al. 2024), which contains 2,655 samples. For each test sample, its context comprises a total of 20 documents, with only one document containing the information corresponding to the question attached to the context. The key document is placed at positions 1st, 5th, 10th, 15th, and 20th, respectively, to evaluate the model’s performance when the key information is located at different positions within the context. We also use accuracy (see appendix A) as the metric and the prompt is the same as the original paper.

The results in Table 3 shows, our method achieves the highest average score on this dataset, and in most cases, it demonstrates great improvement when the key document is located in the middle. This substantiates that our method to fortify the model’s retrieval capabilities provides an effective way to mitigate the ‘lost in the middle’ issue and make LLMs more effectively utilize the information in the middle of the long context.

**Case Study** An example of how our fine-tuned model generates a more accurate answer in multi-doc-QA is shown in Figure 6. The original model seems to have completely failed to retrieve the correct paragraph, and gives a wrong answer, even though it has the context window of 32k length, which is much larger than the prompt length, 4k. In contrast, the fine-tuned model can not only correctly give the document number, but also retell the relevant original text as a reference (although there is a slight discrepancy in the word order and format of the model’s responses compared to our training samples), which proves it has learned to better utilize its retrieval skills to give a more accurate answer.

**Influence of Prompt Style** By incorporating “original text paraphrasing” into our training data, it may change the format of the model’s responses to a CoT-like (Wei et al. 2022) form after fine-tuning. We want to find out, if during inference, no special response format is used and we ask the model to provide answers straightforwardly, will our method still show improvement? Thus we use 2 opposite prompts to control model’s response format: 1. CoT-like: we require the model to provide reference text before answering (adding

Prompt	Method	1st	5th	10th	15th	20th	Avg
CoT	baseline	<b>59.80</b>	42.80	44.60	42.60	<b>42.60</b>	46.48
	SFT w/ ours	54.20	<b>51.00</b>	<b>59.40</b>	<b>51.80</b>	36.20	<b>50.52</b>
not CoT	baseline	<b>56.80</b>	42.40	42.80	41.00	<b>41.00</b>	44.79
	SFT w/ ours	52.60	<b>46.60</b>	<b>52.60</b>	<b>44.00</b>	36.20	<b>46.39</b>

Table 4: Performance of Qwen1.5-4b-Chat on NaturalQuestions Multi-Doc-QA (Liu et al. 2024), when using different prompts. The gold document is placed at 5 different positions respectively (at 1st, 5th, 10th, 15th, 20th with 20 documents in total).

Model	0 shot	5 shot
Qwen1.5-4B-Chat	25.77	52.83
Qwen1.5-4B-Chat w/ ours	52.57	51.16
Qwen2-7B-Instruct	70.82	70.96
Qwen2-7B-Instruct w/ ours	70.66	70.75
Llama3-8B-Chinese-Chat	64.94	67.22
Llama3-8B-Chinese-Chat w/ ours	63.83	65.85

Table 5: Performance of the original models and the models fine-tuned by our method, in MMLU with 0-shot and 5-shot settings.

“you must first paraphrase the original text of the relevant paragraphs”); 2. not CoT-like: we ask the model to provide the answer straightforwardly and concisely (adding “you mustn’t repeat paragraphs or sentences in the provided results, but give the answer directly and concisely”).

As shown in Table 4, the model trained with our method achieves higher average scores in both prompts. Moreover, though CoT-like prompt can generally improve accuracy, with our fine-tuning, the model can benefit more from CoT-like prompt (the improvement of CoT is 4.13% for the fine-tuned model and 1.69% for the base model). This demonstrates that our method does not merely alter the format of the model’s responses, but rather enhances the intrinsic capabilities of the model.

**Evaluation on General Ability** Most training samples of our datasets are long-context, so we want to test if such training will impair models’ performance in general tasks, which are usually short-context (degrading in short-context tasks is a common trade-off for long-context training). We evaluate the fine-tuned models on MMLU benchmark (Hendrycks et al. 2021), as shown in Table 5. Although there is a slight decrease in the score in most cases, the maximum decline does not exceed 2%, which remains within an acceptable range. Hence it can be regarded that our method substantially does not impair the model’s generalization ability.

## 5 Conclusion

In this paper, we propose a novel and practical method to enhance model’s long-context capability. We first discover models often fail to correctly retrieval in composite tasks. Then, we use “paraphrasing the original text” to construct a training dataset which can effectively enhancing model’s

retrieval capability. Finally, through low-cost fine-tuning, we obtain models with superior performance on long-context tasks. Our research provides a new way to help LLMs handle long context better, and mitigate the problem of “lost in the middle”.

## Appendix A Additional Details

### Datasets in Longbench

The datasets we selected is shown in 6. For each task type, we choose one English dataset and one Chinese dataset, and the score of this task is averaged across these two.

Task	dataset en	dataset zh
Multi-doc-QA	hotpotqa	dureader
Single-doc-QA	multifieldqa	multifieldqa
Retrieval	passage_retrieval	passage_retrieval
Summarization	qmsum	vcsum
Few-shot	trec	lsht

Table 6: Datasets names for different tasks in Longbench.

### Metric for QA Tasks

We use the same metric as which has been used in “lost in the middle” issue (Liu et al. 2024). This metric evaluate whether any of the correct answers appear in the predicted output. The calculation is very simple. If the ground truth (usually a few words) is in the model’s prediction, it is considered a correct sample, otherwise it is incorrect, and the accuracy is calculated accordingly. Before calculating, punctuation and articles will be removed from the text, and all letters will be lowercase.

### Time Consumption

Although our method needs to generate more text, in practice, our method does not increase the time cost too much in long-context scenarios. As shown in Table 7, we compare the inference time cost of our method and the baseline, which indicates there is only a minor increase on time consumption when the context is long.

## Acknowledgments

This work was supported by National Natural Science Foundation of China (U2336208,82090053,62262002) and Tsinghua University Initiative Scientific Research Program of Precision Medicine (2022ZLA007).

prompt length	baseline	ours
3k	1.4	2.5
24k	12.4	15.1

Table 7: The average time consumption (seconds) per QA sample, in different input lengths, in the Multi-Doc-QA task.

We are grateful to Huiqiang Jiang from Microsoft Research Asia for his guidance in writing and revising this paper.

## References

- An, C.; Gong, S.; Zhong, M.; Zhao, X.; Li, M.; Zhang, J.; Kong, L.; and Qiu, X. 2023. L-Eval: Instituting Standardized Evaluation for Long Context Language Models.
- An, S.; Ma, Z.; Lin, Z.; Zheng, N.; and Lou, J.-G. 2024. Make Your LLM Fully Utilize the Context.
- Anthropic. 2023. Long context prompting for Claude 2.1.
- BAAI. 2023. Wudao Corpus.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; Hui, B.; Ji, L.; Li, M.; Lin, J.; Lin, R.; Liu, D.; Liu, G.; Lu, C.; Lu, K.; Ma, J.; Men, R.; Ren, X.; Ren, X.; Tan, C.; Tan, S.; Tu, J.; Wang, P.; Wang, S.; Wang, W.; Wu, S.; Xu, B.; Xu, J.; Yang, A.; Yang, H.; Yang, J.; Yang, S.; Yao, Y.; Yu, B.; Yuan, H.; Yuan, Z.; Zhang, J.; Zhang, X.; Zhang, Y.; Zhang, Z.; Zhou, C.; Zhou, J.; Zhou, X.; and Zhu, T. 2023a. Qwen Technical Report.
- Bai, Y.; Lv, X.; Zhang, J.; Lyu, H.; Tang, J.; Huang, Z.; Du, Z.; Liu, X.; Zeng, A.; Hou, L.; Dong, Y.; Tang, J.; and Li, J. 2023b. LongBench: A Bilingual, Multitask Benchmark for Long Context Understanding.
- Chen, S.; Wong, S.; Chen, L.; and Tian, Y. 2023a. Extending Context Window of Large Language Models via Positional Interpolation.
- Chen, Y.; Lv, A.; Lin, T.-E.; Chen, C.; Wu, Y.; Huang, F.; Li, Y.; and Yan, R. 2023b. Fortify the Shortest Stave in Attention: Enhancing Context Awareness of Large Language Models for Effective Tool Use.
- Chen, Y.; Qian, S.; Tang, H.; Lai, X.; Liu, Z.; Han, S.; and Jia, J. 2023c. LongLoRA: Efficient Fine-tuning of Long-Context Large Language Models.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- FlagAlpha. 2023. Llama2-Chinese.
- Gao, Y.; Wang, L.; Fang, J.; Hu, L.; and Cheng, J. 2023. Empower Your Model with Longer and Better Context Comprehension.
- gkamradt. 2023. LLMTest\_NeedleInAHaystack: Doing simple retrieval from LLM models at various context lengths to measure accuracy.
- He, J.; Pan, K.; Dong, X.; Song, Z.; Liu, Y.; Liang, Y.; Wang, H.; Sun, Q.; Zhang, S.; Xie, Z.; and Zhang, J. 2023. Never Lost in the Middle: Improving Large Language Models via Attention Strengthening Question Answering.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Hsieh, C.-Y.; Chuang, Y.-S.; Li, C.-L.; Wang, Z.; Le, L. T.; Kumar, A.; Glass, J.; Ratner, A.; Lee, C.-Y.; Krishna, R.; and Pfister, T. 2024. Found in the Middle: Calibrating Positional Attention Bias Improves Long Context Utilization.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. I.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.-A.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023a. Mistral 7B.
- Jiang, H.; Wu, Q.; Luo, X.; Li, D.; Lin, C.-Y.; Yang, Y.; and Qiu, L. 2023b. LongLLMLingua: Accelerating and Enhancing LLMs in Long Context Scenarios via Prompt Compression.
- Li, D.; Shao, R.; Xie, A.; Sheng, Y.; Zheng, L.; Gonzalez, J. E.; Stoica, I.; Ma, X.; ; and Zhang, H. 2023a. How Long Can Open-Source LLMs Truly Promise on Context Length?
- Li, J.; Wang, M.; Zheng, Z.; and Zhang, M. 2023b. LooGLE: Can Long-Context Language Models Understand Long Contexts?
- Liu, N. F.; Lin, K.; Hewitt, J.; Paranjape, A.; Bevilacqua, M.; Petroni, F.; and Liang, P. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics*, 12: 157–173.
- OpenAI. 2023. GPT-4 Technical Report. Technical report.
- Patel, J. M. 2020. Introduction to Common Crawl Datasets. In Patel, J. M., ed., *Getting Structured Data from the Internet: Running Web Crawlers/Scrapers on a Big Data Production Scale*, 277–324. Berkeley, CA: Apress. ISBN 978-1-4842-6576-5.
- Peng, B.; Quesnelle, J.; Fan, H.; and Shippole, E. 2023. YaRN: Efficient Context Window Extension of Large Language Models.
- Peysakhovich, A.; and Lerer, A. 2023. Attention Sorting Combats Recency Bias In Long Context Language Models.
- Su, J.; Lu, Y.; Pan, S.; Murtadha, A.; Wen, B.; and Liu, Y. 2021. RoFormer: Enhanced Transformer with Rotary Position Embedding.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Together. 2023. Long-Data-Collections.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Ferrer, C. C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Hosseini,

S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P. S.; Lachaux, M.-A.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E. M.; Subramanian, R.; Tan, X. E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J. X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; and Scialom, T. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models.

Wang, M.; Chen, L.; Fu, C.; Liao, S.; Zhang, X.; Wu, B.; Yu, H.; Xu, N.; Zhang, L.; Luo, R.; Li, Y.; Yang, M.; Huang, F.; and Li, Y. 2024. Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E. H.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024. Qwen2 Technical Report.

Zhang, Z.; Chen, R.; Liu, S.; Yao, Z.; Ruwase, O.; Chen, B.; Wu, X.; and Wang, Z. 2024. Found in the Middle: How Language Models Use Long Contexts Better via Plug-and-Play Positional Encoding.