

# Unleashing Multi-Hop Reasoning Potential in Large Language Models through Repetition of Misordered Context

Sangwon Yu<sup>1</sup>    Ik-hwan Kim<sup>1</sup>    Jongyoon Song<sup>1</sup>  
Saehyung Lee<sup>1</sup>    Junsung Park<sup>1</sup>    Sungroh Yoon<sup>1,2\*</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Seoul National University

<sup>2</sup>AIS, ASRI, INMC, ISRC, and IPAI, Seoul National University

{dbtkddnjs96, 12kimih, coms1580, halo8218, jerryray, sryoon}@snu.ac.kr

## Abstract

Multi-hop reasoning, which requires multi-step reasoning based on the supporting documents within a given context, remains challenging for large language models (LLMs). LLMs often struggle to filter out irrelevant documents within the context, and their performance is sensitive to the absolute position of supporting documents within that context. In this paper, we identify an additional challenge: LLMs’ performance is also sensitive to the order, relative position, in which the supporting documents are presented. We refer to this as the **misordered context** problem. To address this issue, based on the theoretical approach, we propose a simple yet effective method called **context repetition (CoRe)**, which involves prompting the model by repeatedly presenting the context. This ensures that certain contiguous reasoning segments within supporting documents are presented in the optimal order, effectively guiding the model’s reasoning in the appropriate direction. Applying CoRe, we improve the F1 score by up to 30%p on multi-hop QA tasks and increase accuracy by up to 70%p on a synthetic task. Additionally, CoRe helps mitigate the well-known “lost-in-the-middle” problem in LLMs and can be effectively combined with retrieval-based approaches utilizing Chain-of-Thought (CoT) reasoning.

## 1 Introduction

Large language models (LLMs) (Brown et al., 2020; Ouyang et al., 2022; Touvron et al., 2023) are capable of performing various complex natural language processing tasks due to their emergent abilities in in-context learning, instruction following, and step-by-step reasoning (Wei et al., 2022; Kojima et al., 2022; Zhao et al., 2023). However, multi-hop reasoning still remains a challenging task

\* Corresponding author

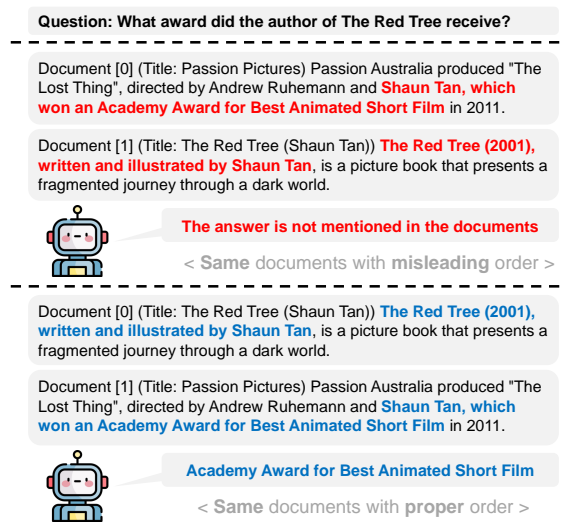


Figure 1: An example of the misordered context problem in multi-hop reasoning of large language models, sampled from the MuSiQue dataset. Model performance is sensitive to the order of given documents.

for LLMs. In this task, multiple external source documents related to the given query are provided as context, and the model searches for useful supporting documents within this context to find the answer. The model derives an answer through multiple steps of reasoning based on these supporting documents (Mavi et al., 2024).

During this process, the model faces several challenges in terms of grounding to the context (Gao et al., 2023). First, the model struggles to filter out noisy documents that are irrelevant to the correct answer (Shi et al., 2023). In practice, noisy documents can significantly degrade the reasoning performance of LLMs (Cuconasu et al., 2024; Wu et al., 2024). Additionally, LLMs are highly sensitive to the absolute position of supporting documents within the context. Specifically, when a document is positioned in the middle of the con-

text, the model may fail to recognize it properly, resulting in a “lost-in-the-middle” problem (Liu et al., 2023). These issues become more critical as the maximum length of the context grows and the number of retrieved documents increases (Lee et al., 2024; Li et al., 2024).

In this study, we address additional context grounding problems related to multi-hop reasoning that have not been deeply discussed so far. In a multi-hop reasoning task, there are multiple supporting documents, and the model derives the answer by considering each document through step-by-step reasoning. Due to the nature of the decoder-only structure, which performs causal attention, LLMs can only perceive documents in a left-to-right order. If this order is presented unfavorably for the model’s reasoning, the model’s performance may be significantly degraded. Following this rationale, we observe that, even in ideal situations without noisy documents, LLMs exhibit up to a 26%p difference in F1 score depending on the order, relative position, of supporting documents. Since there is no guarantee that supporting documents will be presented to the model in the appropriate order for the model reasoning during actual tasks, it is inevitable that this problem arises. We call it the **misordered context** problem. Figure 1 shows an example of the misordered context problem.

To address the problem, we approach it from the perspective of context augmentation. We aim to construct the augmented context that presents the supporting documents in a more appropriate order than the original context. Through theoretical analysis, we demonstrate that by repeating the context containing  $k$  supporting documents  $k$  times, all possible orders of presenting the documents can be covered. That is, *the  $k$ -repeated context can always present the supporting documents to the model in the optimal order*. In real-world scenarios, however, the value of  $k$  for a given context is unknown. Furthermore, repeating the context leads to an increase in inference time costs. Based on this theoretical analysis and its identified limitations, we propose the **context repetition (CoRe)** method: *the given context is repeated  $\hat{k}$  times in which  $\hat{k}$  is the predetermined hyperparameter*. This practical approach ensures that all  $\hat{k}$ -contiguous reasoning chains within supporting documents are presented in the optimal order, guiding LLM’s reasoning in a more coherent and structured manner.

To evaluate CoRe, we conduct experiments on 3 multi-hop QA benchmarks—HotpotQA (Yang

et al., 2018), 2WikiMultihopQA (Ho et al., 2020), and MuSiQue (Trivedi et al., 2022)—as well as an additional synthetic task. In the synthetic task, the model is required to search for elements in lists composed of non-negative integer elements, based on context where information is presented in a reverse order. In the evaluation, CoRe significantly improves multi-hop QA task performance across various LLMs. For example, in the 2WikiMultihopQA task using the Llama-3.1-8B-Instruct model, we observe up to 30%p improvement in F1 score, and in the synthetic task, we observe up to a 70%p improvement in accuracy. Through in-depth empirical analysis, we demonstrate that CoRe provides performance benefits to the model by improving order-related reasoning. Additionally, we show that CoRe alleviates the previously known “lost-in-the-middle” problem and can be effectively combined with Chain-of-Thought (CoT) methods in retrieval-augmented tasks.

In summary, our work (1) introduces the misordered context problem, a critical challenge for LLMs in multi-hop reasoning, (2) theoretically proposes context augmentation as a solution to this issue, which forms the foundation of our practical and effective CoRe method, and (3) demonstrates that CoRe significantly enhances model performance across various tasks and effectively mitigates the misordered context problem.

## 2 Related Work

The process of modeling the multi-step reasoning required in a multi-hop QA task typically involves methods that decompose the query step-by-step to predict the answer (Min et al., 2019; Yadav et al., 2019; Malon and Bai, 2020). From the perspective of utilizing LLMs, methods have been proposed that combine with CoT techniques, allowing the model to self-generate subsequent questions and predict intermediate answers while performing multi-step reasoning (Yao et al., 2023; Khot et al., 2022; Lazaridou et al., 2022; Press et al., 2023; Jiang et al., 2023b; Xu et al., 2024). However, these methods often involve complex processes where errors propagate and accumulate. To reduce this complexity, we aim to enhance LLMs’ innate multi-hop reasoning capabilities.

In retrieval-augmented generation (Lewis et al., 2020; Izacard et al., 2023; Shi et al., 2024a), the typical approach for multi-hop QA is to retrieve documents related to the query, and have the LLM




Figure 2: Evaluation results of F1 score for each query type in MuSiQue with permuted clean contexts.

reason based on this context in a "retrieve-and-reason" fashion. A key issue here is handling the noise introduced by irrelevant documents within the context. To address this, methods have been proposed to help the model correctly locate the evidence within the context (Shi et al., 2024b; Kim et al., 2024; Yu et al., 2023; Yan et al., 2024). With the recent increase in the LLMs' context length, Ziyang Jiang (2024) have also explored including more documents and leveraging the LLM's reasoning capabilities. We aim to propose a method focusing on document order to enable the LLM to reason correctly within the context.

From the perspective of repetition in LLMs, Springer et al. (2024) demonstrate that by repeating sentences during the sentence embedding extraction process, it is possible to compute embeddings that encode bidirectional context even in unidirectional LMs. Xu et al. (2023) show that in arithmetic reasoning tasks, repeating the query can improve performance by leveraging bidirectional understanding. While these methods are similar to ours in that they employ repetition, we differ fundamentally in that we focus on the issue of document order in multi-hop reasoning rather than on simple bidirectional understanding.

From the perspective of context sensitivity in LLMs, recent works focus on addressing the issue of absolute positional bias in tasks where LLMs are required to choose between multiple options (Pezeshkpour and Hruschka, 2024; Zheng et al., 2024). Also, Liu et al. (2023) also address the issue of absolute positional bias, specifically the "lost in the middle" problem. This problem pertains to the model's difficulty in effectively recognizing information that appears in the middle positions of the context provided to the LLM. This absolute positional bias issue differs from the ordering problem we address in our work, which involves ensuring that LLMs interpret supporting documents

within the context in the correct reasoning chain order, relative position, in the multi-hop reasoning task. Chen et al. (2024) analyze the sensitivity of reasoning performance in LLMs to the order in which rules or premises are presented during logical and mathematical reasoning. Our work investigates longer-context scenarios and addresses a broader spectrum of multi-hop reasoning tasks, including open-domain question answering and synthetic tasks, where document-level ordering issues are more pronounced. Furthermore, we propose a concrete methodology to mitigate order-related challenges, enhancing the novelty of our approach in addressing this problem.

### 3 CoRe: Context Augmentation for Misordered Supporting Documents

In this section, we empirically demonstrate that the performance of LLMs in multi-hop reasoning is sensitive to the order of supporting documents within the context. We refer to this phenomenon as the **misordered context problem**. We theoretically formalize this problem from the perspective of context structure. Building on the formalization and considering its practical limitation, we propose **context repetition (CoRe)**. This method enables LLMs to interpret specific contiguous reasoning segments in an optimal order, facilitating more effective multi-hop reasoning within a given context.

#### 3.1 Problem of misordered Context in Multi-Hop Reasoning

We begin by addressing the following research question: *Is the performance of LLMs on multi-hop reasoning tasks sensitive to the order of supporting documents within the context?* To explore this question, we conduct an empirical study in the context of the multi-hop QA task using the MuSiQue dataset. The dataset consists of samples categorized into 3 types: 2-hop, 3-hop, and 4-hop. For a given

$k$ -hop sample, there are  $k$  supporting documents and several noisy documents that distract from the correct answer. To isolate the effect of supporting document order, we remove the noisy documents and construct clean contexts solely from the supporting documents for our experiments.

For each  $k$ -hop query, we permute the order of the  $k$  supporting documents in every possible way, resulting in a total of  $k!$  different contexts. We then evaluate the model’s response to each of these contexts. In other words, for each  $k$ -hop query, the model generates a total of  $k!$  answers, and we assess these answers using the F1 score. We categorize the evaluation results based on the types of order of the documents within contexts associated with each query. The type is determined by the conditional probability of the correct answer  $a$  for the model  $\theta$ ,  $p_\theta(a | Q, C, I)$  where  $Q$  is the query,  $C$  is the context in the specific order of the supporting documents, and  $I$  is the instruction for question answering. The higher the probability, the more optimally the supporting documents are ordered within the context for the answering of the model. We evaluate the model’s performance across the best-case (highest probability), the worst-case (lowest probability), and the mean performance across all possible contexts for each query.

Figure 2 presents the results of an empirical study on 2-hop, 3-hop, and 4-hop reasoning samples. For each  $k$ -hop type, we randomly sample 200 samples from the validation set for evaluation. We evaluate 4 different LLMs: Llama-3.1-8B-Instruct (Dubey et al., 2024), Mistral-7B-Instruct-v0.3 (Jiang et al., 2023a), Qwen2.5-7B-Instruct (Team, 2024), and Phi-3.5-mini-instruct (Abdin et al., 2024). Across all models, the performance gap between the best and worst cases is consistently pronounced, with at least 10 points in terms of F1 score. As the number of hops increases, indicating more complex reasoning types, the performance variation grows, reaching up to 26 points in terms of F1 score. Given that the context in this study consists solely of supporting documents, the performance disparity is expected to widen further in real-world tasks where noisy documents are included in the context.

Based on these findings, we demonstrate the answer to the research question: *the multi-hop reasoning performance of LLMs is indeed significantly sensitive to the order in which supporting documents are presented*. This also aligns with the rationale that decoder-only LLMs, operating with

causal attention, interpret the context in a fixed left-to-right order. In real-world scenarios, common prompting methods such as retrieval or search-based approaches do not guarantee that supporting documents will be presented to the model in an optimal order. Consequently, the model is often exposed to a misordered context that is disadvantageous for reasoning, leading to degraded performance. We term this issue the **misordered context** problem, and in the following sections, we describe our approach to addressing this problem.

### 3.2 Context Augmentation Approach

For the misordered context problem, we formalize a solution from the perspective of context manipulation. We consider a  $k$ -hop reasoning scenario, where the model must derive an answer based on a prompt consisting of a query  $Q$ , context  $C$ , and instruction  $I$ . The context  $C$  is composed of  $k$  supporting documents as well as multiple noisy documents, structured as follows:

$$C = (n_0, d_1, n_1, \dots, d_k, n_k), \quad (1)$$

where  $d_i$  is the  $i$ -th supporting document within the context, and  $n_i$  is an arbitrary number of noisy documents located between supporting documents, as well as at the beginning and the end of the context. At this point, we define an order  $\sigma$  for the supporting documents as follows:

**Definition 1.** Let  $\sigma(\cdot)$  be a permutation function for  $(1, 2, \dots, k)$ . Then, order  $\sigma$  is defined as

$$\sigma(d_1, d_2, \dots, d_k) \triangleq (d_{\sigma(1)}, d_{\sigma(2)}, \dots, d_{\sigma(k)}).$$

Based on the definition of order, we define the set of contexts that include supporting documents in a specific order  $\sigma$  as follows:

**Definition 2.** Let  $\sigma$  be an order of supporting documents. Then, we define  $\mathcal{C}_\sigma$  as the set of all contexts that include the supporting documents in order  $\sigma$ :

$$\mathcal{C}_\sigma \triangleq \{C \mid C = (n_0, d_{\sigma(1)}, n_1, \dots, d_{\sigma(k)}, n_k), \\ \forall n_0, n_1, \dots, n_k\}.$$

The order  $\sigma$  defined above can exist in a total of  $k!$  different configurations, and thus there are  $k!$  possible sets  $\mathcal{C}_\sigma$ . The initially provided context for answering the query belongs to some arbitrary  $\mathcal{C}_\sigma$ , and when  $\sigma$  is the optimal order, the model is expected to perform at its best. We define the optimal order  $\sigma_\theta^*$  for a specific model  $\theta$  as follows:



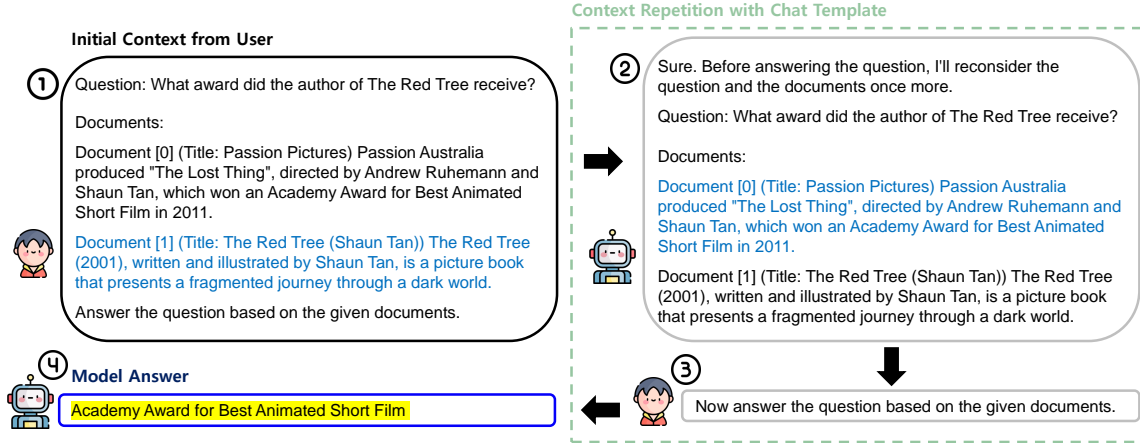


Figure 3: Illustration of CoRe where the model understands the context with *the optimal order of documents [1] and [0]*. Text with a white background is the prompt, and text with a yellow background is the model output.

**Definition 3.** Let  $\theta$  be a model that performs multi-hop reasoning. Then, we define  $\sigma_\theta^*$  as the optimal order of the given documents  $d$  for the model  $\theta$ :

$$\sigma_\theta^* \triangleq \operatorname{argmax}_\sigma \mathbb{E}_{C \in \mathcal{C}_\sigma} [p_\theta(a | Q, C, I)].$$

That is, the optimal order of supporting documents for a given model is defined as the order that maximizes the expected probability of generating the correct answer in the presence of arbitrary noisy documents. Based on the formalization so far, we aim to find a function that maps a given context  $C$  to another context  $C'$ , which has the supporting documents in the optimal order. The goal is to increase the probability of the model generating the correct answer. Therefore, we introduce an augmentation function  $f(\cdot)$  that satisfies the following conditions:

$$f(C) \in \mathcal{C}_{\sigma_\theta^*}. \quad (2)$$

### 3.3 Context Repetition

Now, we study the augmentation function discussed earlier. We define an augmentation function  $f_{\text{rep}}^{(k)}(\cdot)$  as follows:

$$f_{\text{rep}}^{(k)}(C) = \underbrace{C \oplus C \oplus \dots \oplus C}_{k \text{ times}}. \quad (3)$$

Here,  $C$  represents the given context, and  $\oplus$  denotes the concatenation operation. Hence, the function  $f_{\text{rep}}^{(k)}(C)$  creates an augmented context by just repeating the given context  $k$  times. Please note that  $k = 1$  indicates the absence of repetition. The following theorem provides an important property of  $f_{\text{rep}}^{(k)}(\cdot)$  in relation to  $\sigma$ :

**Theorem 1.** For a given context  $C$  with its supporting documents ordered according to some permutation  $\tau$ , the augmented context  $f_{\text{rep}}^{(k)}(C)$  belongs to the set  $\mathcal{C}_\sigma$  for any permutation  $\sigma$ . Formally,

$$f_{\text{rep}}^{(k)}(C) \in \mathcal{C}_\sigma, \quad \forall \sigma.$$

The proof of this theorem can be found in Appendix A. From Theorem 1, we can derive the following corollary:

**Corollary 1.1.**  $f_{\text{rep}}^{(k)}(C)$  always belongs to the set of all contexts that present the supporting documents in the optimal order. Formally,

$$f_{\text{rep}}^{(k)}(C) \in \mathcal{C}_{\sigma_\theta^*}.$$

Based on Corollary 1.1, we conclude that  $f_{\text{rep}}^{(k)}(\cdot)$  is an appropriate augmentation function that satisfies Eq. 2.

### 3.4 Practical Approach

From the previous section 3.3, we can conclude that  $f_{\text{rep}}^{(k)}(\cdot)$  can serve as a methodology for mitigating the misordered context problem. However, for direct application, the following two issues must be addressed:

1. It is necessary to know the number of supporting documents in the given context,  $k$ . (Issue 1)
2. The model should effectively handle the computational cost induced by the increased length when the context is repeated  $k$  times. (Issue 2)

In real-world multi-hop reasoning scenarios, it is not feasible to fully address issue 1. Additionally, issue 2 becomes intractable as the number of documents included in the context grows, and more

importantly, it incurs a significant additional cost in terms of memory.

To overcome these critical limitations, we replace  $k$  with a predetermined hyperparameter  $\hat{k}$ , resulting the function  $f_{\text{rep}}^{(\hat{k})}(\cdot)$  which can be considered an approximation of  $f_{\text{rep}}^{(k)}(\cdot)$ . This directly addresses the issue 1 by constraining the unknown value of  $k$  to predetermined  $\hat{k}$ . While this approach does not guarantee that all supporting documents will be recognized in the optimal order, it ensures that *at least the  $\hat{k}$ -contiguous reasoning chains within supporting documents are interpreted in the optimal order*. This allows the LLM to perform multi-hop reasoning from the optimal starting point, guiding the model to execute subsequent reasoning steps correctly. From the perspective of the issue 2, the introduction of  $\hat{k}$  significantly reduces the additional computational and memory costs, bringing them to a manageable level.

Consequently, we propose to use  $f_{\text{rep}}^{(\hat{k})}(\cdot)$  as a method for solving the misordered context problem. In other words, we introduce the **context repetition (CoRe)** methodology, which repeatedly presents the given context to guide the model recognize the supporting documents in the appropriate order. Figure 3 illustrates the detailed operation of our methodology when  $\hat{k} = 2$ . In practice, we implement this method by utilizing a chat template to ensure that the model naturally repeats the context given by the user.

## 4 Experiments

In this section, we evaluate CoRe on multi-hop QA tasks and a synthetic task that requires complex reasoning. In addition to the main experiments, we conduct an analysis of the mechanism through which CoRe achieves performance improvements in the model.

### 4.1 Experimental Settings

We conduct experiments on three multi-hop QA benchmarks: HotpotQA, 2WikiMultihopQA, and MuSiQue. For HotpotQA and 2WikiMultihopQA, reasoning-type annotations are provided for each query. For these two datasets, we create evaluation sets by sampling an equal number of queries from each reasoning type within the validation set. In the case of MuSiQue, as described in Section 3.1, query types are categorized by the number of hops. We use the entire validation set as the evaluation set for this dataset. To construct the context

of multiple documents, we combine the supporting documents with distracting (noisy) documents assigned to each sample in the datasets. Detailed statistics for each evaluation set are provided in Appendix B. In the multi-hop QA evaluation, we set the  $\hat{k}$  within the range of  $[1, 3]$  ( $\hat{k} = 1$  means the naive LLM baseline without repetition).

In addition to the multi-hop QA tasks, we evaluate model performance on an additional synthetic task. Specifically, we provide the model information about a list that contains non-negative integers as elements in a chained form consisting of two consecutive elements. Afterward, we ask the model for the first element of the list that contains a specific element. At this point, the specific element is the last element of a list. For example, when there are two lists with 3 elements each, the context and question are structured as follows:

```
All the 2 lists described below
    contain exactly 3 elements.
In the list 0, 381 is positioned
    immediately before 512.
In the list 1, 7123 is positioned
    immediately before 34.
In the list 0, 512 is positioned
    immediately before 1021.
In the list 1, 34 is positioned
    immediately before 6397.
Question: What is the first element
    of the list that contains 1021?
```

Since the elements of each list are presented in sequential order from the first to the last, the process of finding the list that contains the last element and then searching for the first element of that list must be performed in a right-to-left order. This poses a challenge for models that typically process context in a left-to-right order. Moreover, since this type of task does not involve specific factual knowledge, it minimizes the interference of the model’s pre-existing factual knowledge (Xiong et al., 2024). Thus, these characteristics of the synthetic task enable us to establish an environment that clearly exposes the model’s inherent misordered context problem. We create a synthetic dataset consisting of 1000 samples for the evaluation. Each sample is composed of 10 lists, each containing 3 integers between 0 and 9999 that do not overlap.

We conduct experiments using a total of five models, including GPT-4o-mini (OpenAI, 2024), in addition to the four LLMs used in section 3.1. For Multi-hop QA tasks, we evaluate by measuring the F1 score based on the short answers gener-

Models	HotpotQA		2WikiMultihopQA				MuSiQue		
	Comparison	Bridge	Compositional	Comparison	Inference	Bridge-Comparison	2-hop	3-hop	4-hop
Llama-3.1-8B	49.31	61.19	38.10	55.77	40.75	34.86	32.05	26.24	27.06
+ CoRe ( $\hat{k} = 2$ )	46.93	66.82	48.15	58.13	44.81	65.11	39.10	38.23	34.08
+ CoRe ( $\hat{k} = 3$ )	<b>50.46</b>	<b>70.54</b>	<b>50.74</b>	<b>62.58</b>	<b>47.44</b>	<b>68.18</b>	<b>43.64</b>	<b>41.02</b>	<b>36.08</b>
Mistral-7B	34.42	41.14	18.82	32.90	29.77	15.43	15.14	10.31	11.00
+ CoRe ( $\hat{k} = 2$ )	30.41	45.79	27.78	35.67	36.39	<b>29.09</b>	20.6	15.74	12.36
+ CoRe ( $\hat{k} = 3$ )	<b>42.50</b>	<b>54.51</b>	<b>38.93</b>	<b>42.91</b>	<b>38.67</b>	28.70	<b>23.96</b>	<b>20.17</b>	<b>14.78</b>
Qwen2.5-7B	61.96	61.66	41.36	65.47	37.26	39.45	33.56	34.87	37.49
+ CoRe ( $\hat{k} = 2$ )	<b>67.03</b>	70.33	<b>54.20</b>	<b>70.30</b>	42.40	64.54	44.76	<b>43.10</b>	45.08
+ CoRe ( $\hat{k} = 3$ )	63.37	<b>71.59</b>	54.13	69.64	<b>44.30</b>	<b>68.00</b>	<b>46.12</b>	42.84	<b>45.88</b>
Phi-3.5-mini	26.44	44.17	27.70	45.26	27.00	35.19	17.60	16.80	12.67
+ CoRe ( $\hat{k} = 2$ )	23.29	49.67	37.64	<b>50.92</b>	32.34	<b>36.71</b>	28.60	26.31	26.59
+ CoRe ( $\hat{k} = 3$ )	<b>41.41</b>	<b>59.07</b>	<b>40.62</b>	46.91	<b>33.76</b>	22.98	<b>31.49</b>	<b>26.83</b>	<b>32.40</b>
GPT-4o-mini	<b>66.92</b>	68.66	51.96	76.63	43.84	60.96	45.26	42.04	40.04
+ CoRe ( $\hat{k} = 2$ )	64.72	<b>72.67</b>	<b>57.43</b>	79.14	<b>52.03</b>	66.96	52.47	<b>49.47</b>	42.53
+ CoRe ( $\hat{k} = 3$ )	66.52	<b>74.06</b>	56.72	<b>79.84</b>	51.09	<b>68.84</b>	<b>52.66</b>	49.37	<b>42.89</b>

Table 1: Main results of F1 score in the multi-hop QA tasks. All models are instruction-tuned LLMs.




Figure 4: Main results in the synthetic task. Repetition step denotes the number of additional repetitions ( $\hat{k} - 1$ ).

ated by the models. We evaluate synthetic tasks by measuring accuracy based on whether the answers generated by the models match the correct answers. All prompts used in the experiments can be found in the Appendix C.

## 4.2 Main Results

Table 1 shows the F1 scores of baseline models and models with CoRe applied for multi-hop QA tasks. Overall, performance improvements were observed across all models and datasets. Notably, the performance appears to reach a plateau in the  $\hat{k} = 2$  setting, where the context is repeated only once. Furthermore, we find additional performance improvements when further repetitions are applied ( $\hat{k} = 3$ ). In the case of HotpotQA, a significant performance improvement was seen in the bridge type. This is because answering bridge-type questions requires

reasoning across documents by linking information between them, which is closely related to the order of documents. For 2WikiMultihopQA, a considerable performance improvement was observed in the bridge-comparison type, with Llama achieving a 30-point improvement and Qwen achieving a 25-point improvement in the F1 score. This is because the bridge-comparison type is a complex 4-hop reasoning task that combines bridge and comparison from HotpotQA. The MuSiQue dataset’s overall performance improved regardless of the number of hops. These results indicate that CoRe substantially enhances the model’s multi-hop reasoning performance, particularly in tasks where the order of supporting documents is crucial, effectively guiding the model along the correct reasoning paths.

We measure the accuracy of each model as we increase the number of context repetitions up to 10 ( $\hat{k} = 11$ ) in the synthetic task. As shown in Figure 4, we can see that the accuracy when no repetition is performed has a low value between 0% and 60%. This demonstrates that existing LLMs are unable to perform the synthetic task directly at once. However, when CoRe is applied, we observe a significant performance improvement, with accuracy increasing by at least 30% points and up to 70% points. Additionally, we can see that many models tend to improve performance as the number of repetitions increases. This suggests that the model leads to the more appropriate order of reasoning chains as repetition progresses multiple times. However, for some models, particularly Qwen, there is a tendency for performance to drop again as the repetition steps increase. This appears to be due to noisy




Figure 5: Performance of Llama-3.1-8B-Instruct with permuted contexts of MuSiQue during repetitions. The red line denotes the context in the worst order, and the purple line denotes the context in the best order.

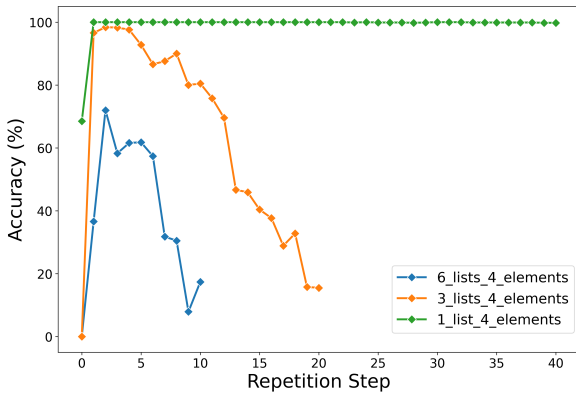


Figure 6: Performance of Qwen2.5-7B-Instruct for contexts with various noise powers in the synthetic task.

information being repeated in the context, negatively affecting its reasoning ability. We provide a detailed analysis of this in Section 4.3.

### 4.3 Analysis

**Does CoRe really guide the model along the better reasoning chain?** We conduct additional analysis to investigate whether the performance improvement from CoRe is indeed related to the order of the documents. Similar to the experiments described in Section 3.1, we permute the order of the documents in the context across all possible combinations and measure the model’s performance for each permuted context. Specifically, for 200  $k$ -hop samples from MuSiQue, we construct a context with five documents, including the  $k$  supporting documents and some random  $5 - k$  noisy documents assigned to each sample. As a result, each sample has 120 permuted contexts, and we sort the performance of each context from the worst order to the best order, based on  $p_{\theta}(a | Q, C, I)$ . Figure 5 shows the results for Llama. We plot the worst

order with a red line and the best with a purple line. For other cases, we plot their performance with the line of a continuous gradient of colors from red to purple, corresponding to the performance range from the worst to the best case. In other words, the closer the color is to red in the rainbow spectrum, the lower  $p_{\theta}(a | Q, C, I)$  is, and the closer to purple, the higher  $p_{\theta}(a | Q, C, I)$  is. Across all query types, we find compelling evidence that CoRe contributes to performance improvements in terms of document order: **The closer the context is to the worst order case, the greater the performance improvement from repetition.** Additionally, we observe that as the number of hops increases, performance improves progressively with iterative repetition, indicating that multiple repetitions are more effective when the required reasoning steps increase. Appendix D shows the results for other LLMs in this analysis.

**Impact of noisy documents** We further analyze the impact of noisy documents on CoRe. In the synthetic task, we measure accuracy by increasing the number of repetitions in cases where the number of lists is 6, 3, and 1, respectively. Here, the decrease in the number of lists means a reduction in the proportion of information unrelated to the search target. Figure 6 shows the results of this analysis. Considering that the length of the context decreases as the number of lists handled in the context decreases, we perform more repetitions in cases with fewer lists. As a result, we confirm that as the proportion of noisy information within the context decreases, the performance degradation caused by excessive repetition is reduced. In particular, when there is no noisy information, performance did not decrease even with a large number of repetitions. Since noisy information adversely affects model reasoning at



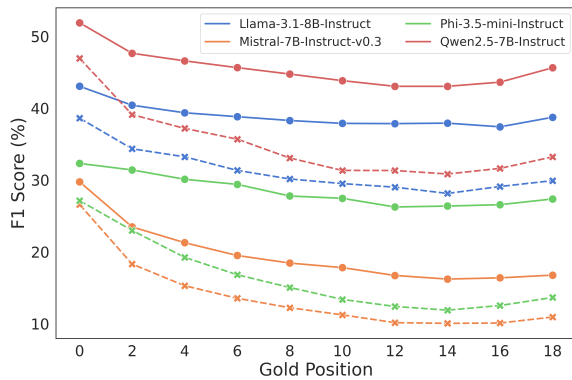


Figure 7: Performance of LLMs for various positions of the supporting document. The solid line represents CoRe, and the dotted line represents the baseline.

high repetition counts, exploring the valid number of repetitions required for model reasoning could be an interesting future work based on our findings.

## 5 Further Studies

### 5.1 Robustness to Positional Bias

We analyze the absolute positional bias in the context based on the position of supporting documents. Here, positional bias refers to the phenomenon where the QA performance of an LLM varies depending on the position of the supporting documents within the context, known as the "lost-in-the-middle" problem (Liu et al., 2023). We concatenate  $k$  supporting documents in MuSiQue and measure the F1 scores of the baseline without repetition ( $\hat{k} = 1$ ) and CoRe ( $\hat{k} = 2$ ) while shifting the position of the documents from 0 to 18 by increments of 2. As shown in 7, we observe that the performance gap between the baseline and CoRe increases as the position of the supporting documents moves toward the middle. Particularly for Llama (blue line) and Phi (green line), while the baseline suffers from the "lost-in-the-middle" problem, CoRe demonstrates robustness against positional bias, exhibiting minimal performance variation depending on the position of the supporting documents. Based on these results, we demonstrate that CoRe is a solution to mitigate the "lost-in-the-middle" problem.

### 5.2 Enhancement in Retrieve-and-Reason

We evaluate the effectiveness of CoRe ( $\hat{k} = 2$ ) in a retrieve-and-reason task, which closely mirrors real-world scenarios. We utilize Contriever (Izacard

Methods	2-hop	3-hop	4-hop
RAG	12.88	8.06	6.70
+ CoRe ( $\hat{k} = 2$ )	<b>17.04</b>	<b>12.44</b>	<b>9.07</b>
IRCoT	23.67	10.64	4.33
+ Dec.	29.03	11.78	6.62
+ CoRe ( $\hat{k} = 2$ )	<b>37.65</b>	<b>19.35</b>	8.90
+ CoRe & Dec.	36.62	17.87	<b>11.52</b>

Table 2: F1 results for retrieve-and-reason tasks in MuSiQue. Dec. refers to Decompose method.

et al., 2021) as the retriever model, constructing the context with up to 100 retrieved documents from MuSiQue samples. We use Llama for this experiment. We first assess CoRe’s performance with (1) the standard RAG, which directly retrieves the top 100 relevant documents to form the context. We also evaluate CoRe with two additional methods: (2) IRCoT (Trivedi et al., 2023), where retrieval is performed sequentially during the CoT reasoning process, and (3) Decompose, where the query is decomposed into sub-queries, and reasoning is conducted by deriving intermediate answers for each sub-query (Press et al., 2023). As shown in Table 2, CoRe significantly improves performance across all methods, demonstrating its effectiveness when integrated into various retrieval-based reasoning approaches. Appendix E presents the experimental details of this section.

Additionally, we report the performance of CoRe when applied in Chain-of-Thought (CoT) scenarios in Appendix D. We also conduct various ablation studies in Appendix F.

## 6 Conclusion

In this work, we introduce the misordered context problem, which critically impacts LLMs’ multi-hop reasoning capabilities. Formalizing this problem regarding context augmentation, we propose a simple yet highly effective solution: the CoRe method, which optimally arranges contiguous reasoning segments within the context to enhance performance. We demonstrate the contributions of our approach across various multi-hop reasoning benchmarks, including synthetic tasks, and elucidate its effectiveness through further studies. Our work highlights a key issue crucial for future discussions in multi-hop reasoning tasks with LLMs. Furthermore, we believe CoRe could provide meaningful insights for research on the test-time scaling of LLMs that perform complex reasoning.

## Limitations

The context repetition method proposed in this work increases the input context length by repeating the prompt’s contents, which introduces additional costs for memory and time during inference. While the time cost is kept manageable due to the efficient KV caching mechanism, future research could explore ways to repeat only the core contents of the context, thereby minimizing unnecessary expansion. Moreover, as discussed in the analysis section, it is crucial to address the performance degradation caused by noisy documents when repeating the context. A promising avenue for future work would be to combine the inherent capabilities of LLMs with techniques that minimize the influence of noisy documents while maximizing the effectiveness of context repetition. Overall, our work lays the foundational groundwork for addressing the misordered context problem through context augmentation, and future studies could build upon this foundation to further refine and extend the approach.

## Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2022R1A3B1077720), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [NO.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)], the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University, Samsung Electronics Co., Ltd (IO240311-09242-01), and a grant from the Yang Young Foundation.

## References

Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 1877–1901.

Xinyun Chen, Ryan Andrew Chi, Xuezhi Wang, and Denny Zhou. 2024. Premise order matters in reasoning with large language models. In *Forty-first International Conference on Machine Learning*.

Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Maarek, Nicola Tonello, and Fabrizio Silvestri. 2024. [The power of noise: Redefining retrieval for rag systems](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024*. ACM.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. [Unsupervised dense information retrieval with contrastive learning](#).

Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Zhengbao Jiang, Frank Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023b. [Active retrieval augmented generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7969–7992, Singapore. Association for Computational Linguistics.

Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2022. Decomposed prompting: A modular approach for solving complex tasks. *arXiv preprint arXiv:2210.02406*.

- Jaehyung Kim, Jaehyun Nam, Sangwoo Mo, Jongjin Park, Sang-Woo Lee, Minjoon Seo, Jung-Woo Ha, and Jinwoo Shin. 2024. Sure: Summarizing retrievals using answer candidates for open-domain qa of llms. In *The Twelfth International Conference on Learning Representations*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. 2022. Internet-augmented language models through few-shot prompting for open-domain question answering. *arXiv preprint arXiv:2203.05115*.
- Jinhyuk Lee, Anthony Chen, Zhuyun Dai, Dheeru Dua, Devendra Singh Sachan, Michael Boratko, Yi Luan, Sébastien MR Arnold, Vincent Perot, Siddharth Dalmia, et al. 2024. Can long-context language models subsume retrieval, rag, sql, and more? *arXiv preprint arXiv:2406.13121*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Yanyang Li, Shuo Liang, Michael Lyu, and Liwei Wang. 2024. Making long-context language models better multi-hop reasoners. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2462–2475.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. *Lost in the middle: How language models use long contexts*. *Preprint*, arXiv:2307.03172.
- Christopher Malon and Bing Bai. 2020. Generating followup questions for interpretable multi-hop question answering. *arXiv preprint arXiv:2002.12344*.
- Vaibhav Mavi, Anubhav Jangra, Adam Jatowt, et al. 2024. Multi-hop question answering. *Foundations and Trends® in Information Retrieval*, 17(5):457–586.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hananeh Hajishirzi. 2019. Multi-hop reading comprehension through question decomposition and rescoring. In *ACL*.
- OpenAI. 2024. *Gpt-4o mini: advancing cost-efficient intelligence*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Pouya Pezeshkpour and Estevam Hruschka. 2024. Large language models sensitivity to the order of options in multiple-choice questions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. *Measuring and narrowing the compositionality gap in language models*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. 2023. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024a. Replug: Retrieval-augmented black-box language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8364–8377.
- Zhengliang Shi, Shuo Zhang, Weiwei Sun, Shen Gao, Pengjie Ren, Zhumin Chen, and Zhaochun Ren. 2024b. *Generate-then-ground in retrieval-augmented generation for multi-hop question answering*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7339–7353, Bangkok, Thailand. Association for Computational Linguistics.
- Jacob Mitchell Springer, Suhas Kotha, Daniel Fried, Graham Neubig, and Aditi Raghunathan. 2024. Repetition improves language model embeddings. *arXiv preprint arXiv:2402.15449*.
- Qwen Team. 2024. *Qwen2.5: A party of foundation models*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. MuSiQue: Multi-hop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. *Interleaving retrieval*

- with chain-of-thought reasoning for knowledge-intensive multi-step questions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Siye Wu, Jian Xie, Jiangjie Chen, Tinghui Zhu, Kai Zhang, and Yanghua Xiao. 2024. How easily do irrelevant inputs skew the responses of large language models? In *First Conference on Language Modeling*.
- Zheyang Xiong, Vasilis Papageorgiou, Kangwook Lee, and Dimitris Papailiopoulos. 2024. From artificial needles to real haystacks: Improving retrieval capabilities in llms by finetuning on synthetic data. *arXiv preprint arXiv:2406.19292*.
- Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2024. Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks. In *Proceedings of the ACM on Web Conference 2024*, pages 1362–1373.
- Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu, Hongbo Xu, Guodong Long, and Jian-guang Lou. 2023. Re-reading improves reasoning in language models. *arXiv preprint arXiv:2309.06275*.
- Vikas Yadav, Steven Bethard, and Mihai Surdeanu. 2019. Quick and (not so) dirty: Unsupervised selection of justification sentences for multi-hop question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2578–2589.
- Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*.
- 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 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2023. Chain-of-note: Enhancing robustness in retrieval-augmented language models. *arXiv preprint arXiv:2311.09210*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*.
- Wenhu Chen Ziyang Jiang, Xueguang Ma. 2024. Longrag: Enhancing retrieval-augmented generation with long-context llms. *arXiv preprint arXiv:2406.15319*.



## A Proof of Theorem 1

We will prove this theorem by demonstrating that the augmented context  $f_{\text{rep}}^{(k)}(C)$ , constructed by repeating the original context  $C$   $k$  times, inherently arranges the supporting documents in the order specified by the permutation  $\sigma$ .

By applying the augmentation function, we obtain:

$$f_{\text{rep}}^{(k)}(C) = \underbrace{(n_0, d_{\tau(1)}, n_1, \dots, d_{\tau(k)}, n_k)}_{\text{1st repetition}} \oplus \underbrace{(n_0, d_{\tau(1)}, n_1, \dots, d_{\tau(k)}, n_k)}_{\text{2nd repetition}} \oplus \dots \oplus \underbrace{(n_0, d_{\tau(1)}, n_1, \dots, d_{\tau(k)}, n_k)}_{\text{k-th repetition}}.$$

We can extract the supporting documents in the order specified by  $\sigma$  by selecting the  $d_{\sigma(i)}$  from the  $i$ -th repetition. All other documents in each repetition of  $C$  (i.e., all  $d_j$  where  $j \neq \sigma(i)$  in the  $i$ -th repetition and all  $n_j$ ) are considered noisy documents. This means they are treated as  $n'_j$  in the definition of  $\mathcal{C}_\sigma$ . The augmented context  $f_{\text{rep}}^{(k)}(C)$  can thus be viewed as:

$$f_{\text{rep}}^{(k)}(C) = (n'_0, d_{\sigma(1)}, n'_1, d_{\sigma(2)}, \dots, d_{\sigma(k)}, n'_k),$$

where each  $d_{\sigma(i)}$  is extracted from the  $i$ -th repetition, and all  $n'_j$  are noisy documents. This directly matches the structure of some  $C_\sigma \in \mathcal{C}_\sigma$ , satisfying:

$$f_{\text{rep}}^{(k)}(C) \in \mathcal{C}_\sigma.$$

Since  $\sigma$  was arbitrary, this holds for any permutation  $\sigma$ .

## B Data Statistics

Table 3 shows the statistics of the evaluation sets in the multi-hop QA tasks conducted in the main experiments.

## C Prompts for Experiments

Table 4 and 5 show the prompt template of CoRe for the multi-hop QA tasks and the synthetic task conducted in the section 4, respectively.

## D Additional Results

In this section, we describe additional experimental results not reported in the main text. Specifically, we extend the analysis from Section 4.3, where the impact of applying CoRe to each permuted context was only evaluated on Llama, to additional models, including Mistral, Qwen, and Phi. Figures 8 to 10 present the analysis results for each model. Similar to the findings for Llama, we observe that the performance improvement from CoRe increases as the context be close to the worst-order cases for all models. This further solidifies that the performance gains of CoRe stem from augmenting the original context into an optimal context. Meanwhile, we notice a slight performance drop in contexts close to the best order as repetition increases, particularly for models like Mistral and Phi. This trend is more pronounced compared to Llama and can be interpreted as a result of these models being less capable of handling noisy documents.

Additionally, we evaluate the performance of CoRe when applied to Chain-of-Thought (CoT) reasoning in multi-hop QA tasks. As shown in Table 6, CoRe consistently leads to significant improvements, particularly in bridge and compositional question types, aligning with the main results presented earlier. In the case of GPT-4o-mini, some performance metrics remain stable, which we attribute to an observed trade-off: while recall increases, precision decreases. This suggests that when CoT reasoning is combined with CoRe, the model is better at deducing the correct answer, but the length of the generated response tends to increase, resulting in a slight drop in precision. In conclusion, our results demonstrate that CoRe can meaningfully enhance performance even in the context of CoT reasoning. Table 7 shows the prompt template of CoRe for CoT reasoning in the multi-hop QA tasks.

## E Experimental Details of Retrieve-and-Reason Task

In this section, we provide the experimental details of the retrieve-and-reasoning task conducted in Section 5.2. The retrieval database consists of 101,962 documents, constructed by ensuring no overlap between the documents assigned to each sample from the MuSiQue training and validation sets. The evaluation set is built by randomly sampling 480 samples from the MuSiQue validation set, comprising 249 2-hop samples, 149 3-hop sam-

ples, and 82 4-hop samples.

The standard RAG approach retrieves the top 100 documents at once, which are then used as the context for prompting the LLM. IRCOT operates by having the LLM perform Chain-of-Thought (CoT) reasoning while generating sentences. After each sentence is generated, the top 10 documents related to the generated sentence are retrieved and concatenated with the previously retrieved documents, updating the context. If the LLM concludes its reasoning with an expression such as "So the answer is," the subsequent output is taken as the final answer. If the context update process exceeds 10 iterations without reaching a conclusion, the entire output generated during the CoT process is considered the final answer. This setup follows the one described in [Trivedi et al. 2023](#). When IRCOT is combined with CoRe, after each CoT sentence is generated, the context of the retrieved documents and previously generated CoT responses is repeated in the prompt.

In the question decomposition setting, LLM is instructed to decompose the query into several sub-queries before answering the query. Once the sub-queries are generated, the LLM is directed to perform step-by-step reasoning by following each sub-query, with the process guided by an instruction that ensures the sub-queries are considered during the answering phase. This setting aligns with the prompt design of [Press et al. \(2023\)](#). Table 8 shows the prompt template of query decomposition.

## F Ablation Studies

### F.1 Prompt Style Ablation

In this section, we investigate the impact of changing the position where CoRe is applied within the prompt. Specifically, we conduct an ablation study for the role where CoRe is applied in the chat template prompt. In our default setup, CoRe is applied through the assistant role, where the context is repeated in the responses generated by the assistant. However, we also consider an alternative setup where the context is repeated in the user role. We evaluate the performance of Llama, Mistral, Qwen, and Phi on 2WikiMultihopQA using F1 scores.  $\hat{k} = 2$  for CoRe in this study.

The results, presented in Table 9, reveal that in some cases, applying CoRe in the user role leads to significant performance improvements. Notably, except for Qwen, the other models show substantial gains when applying CoRe in the user role.

For instance, Mistral demonstrates a 17-point improvement in compositional-type questions and a 15-point gain in comparison-type questions. Phi also shows an 18-point increase in F1 score on bridge-comparison-type questions. We hypothesize that these improvements are due to the task being framed as a single-turn chat when CoRe is applied in the user role. Some models are predominantly tuned on single-turn chat tasks rather than multi-turn interactions, making it easier for them to reason when CoRe is applied in the user role. By adapting the CoRe methodology to align with the model's training environment, we would create a more familiar reasoning context, leading to larger performance gains. Table 11 shows the prompt template of CoRe applied in the user role position within the prompt.

### F.2 Repetition Style Ablation

Additionally, we conduct an ablation study to explore the effects of modifying the repeated context. Specifically, we examine 4 cases where the content of the context is altered during repetition: (1) each document is paraphrased, (2) each document is summarized, (3) the order of the documents is shuffled randomly, and (4) the order of the documents is reversed. We evaluate the performance of each of these cases. The experiments are conducted on 2WikiMultihopQA using Llama, where the paraphrasing and summarizing are directly performed by Llama responsible for reasoning.  $\hat{k} = 2$  for CoRe in this study.

Table 10 presents the results for each case when repetition is applied once. Interestingly, in comparison tasks, all ablation settings yield higher performance than the baseline setting, which involves simply repeating the context as is. This suggests that the process of recognizing the same content in different forms is particularly beneficial for comparison tasks. For the cases of shuffling and reversing the order of documents, performance improvements are observed in three out of four types, excluding inference. This indicates that varying the order of documents during repetition enables the model to consider a wider range of reasoning chains. As discussed in Section 4.3, future research could explore applying CoRe while reducing the noise power within the context, and actively modifying the form of the context, as in this study, could be another potential direction to consider. Table 12 and 13 show the prompt template of paraphrasing and summarizing, respectively.

## **G Computational Environment**

For the main experiments, we utilize NVIDIA A40 GPU. For the experiments for the retrieve-and-reason task, we utilize NVIDIA H100 80GB HBM3 GPU. All the models except GPT-4o-mini in our work are from huggingface. We utilize GPT-4o-mini through OpenAI API package.

## **H Licenses**

HotpotQA, 2WikiMultihopQA, and MuSiQue are licensed under the CC BY-SA 4.0, CC-BY-4.0, and Apache-2.0 licenses, respectively. Llama-3.1, Mistral, Qwen, Phi, and GPT-4o-mini are licensed under the Llama 3.1 COMMUNITY LICENSE, Apache-2.0, Apache-2.0, MIT, and OpenAI.

## **I Usage of AI Writing Assistant**

This paper benefited from linguistic support provided by the AI assistant ChatGPT-4o, which contributed by offering paraphrasing and enhancing the original content. No other assistance was sought beyond these services.




Figure 8: Performance of Mistral-7B-Instruct-v0.3 with permuted contexts of MuSiQue during repetitions. The red line denotes the context in the worst order, and the purple line denotes the context in the best order.




Figure 9: Performance of Qwen2.5-7B-Instruct with permuted contexts of MuSiQue during repetitions. The red line denotes the context in the worst order, and the purple line denotes the context in the best order.




Figure 10: Performance of Phi-3.5-mini-instruct with permuted contexts of MuSiQue during repetitions. The red line denotes the context in the worst order, and the purple line denotes the context in the best order.



Datasets	HotpotQA		2WikiMultihopQA				MuSiQue		
	Comparison	Bridge	Compositional	Comparison	Inference	Bridge-Comparison	2-hop	3-hop	4-hop
# of samples	1000	1000	500	500	500	500	1252	760	405
Context length	869.17	977.33	701.15	567.84	772.21	593.87	1694.4	1827.92	1778.8

Table 3: Data statistics of evaluation sets for each dataset of multi-hop QA tasks in the main experiments. Context length denotes the mean number of words included in the context of each type of sample.

---

**[User role]**

Question: {*Question*}

Documents:

Document [0] {*Document 0*}

Document [1] {*Document 1*}

...

Document [N] {*Document N*}

Answer the question based on the given documents. Respond only the answer within a few words after 'Answer:'.

**[Assistant role]**

Sure. Before answering the question, I'll reconsider the question and the documents once more.

Question: {*Question*}

Documents:

Document [0] {*Document 0*}

Document [1] {*Document 1*}

...

Document [N] {*Document N*}

**[User role]**

Now answer the question based on the documents. Respond only the answer within a few words after 'Answer:'.

**[Assistant role]**

Answer:

---

Table 4: Prompt template of CoRe used in main experiments for the multi-hop QA tasks.

---

**[User role]**  
Information:  
All the {k} lists described in the below contain exactly {n} elements.  
In the list 0, {lists[0][0]} is positioned immediately before {lists[0][1]}.  
...  
In the list {k}, {lists[k][0]} is positioned immediately before {lists[k][1]}.  
...  
In the list 0, {lists[0][n-2]} is positioned immediately before {lists[0][n-1]}.  
...  
In the list {k}, {lists[k][n-2]} is positioned immediately before {lists[k][n-1]}.

Question: What is the first element of the list that contains {lists[0][n-1]}?  
Answer the question based on the given information. Respond only the answer without any explanation after 'Answer:'.

**[Assistant role]**  
Sure. Before answering the question, I'll reconsider the question and the documents {t} times more.

Information:  
{Information}

Question: {Question}  
...  
(repeat t-times)

**[User role]**  
Now answer the question based on the documents. Respond only the answer without any explanation after 'Answer:'.

**[Assistant role]**  
Answer:

---

Table 5: Prompt template of CoRe used in main experiments for the synthetic task.

Models	HotpotQA		2WikiMultihopQA				MusiQue		
	Comparison	Bridge	Compositional	Comparison	Inference	Bridge-Comparison	2-hop	3-hop	4-hop
Llama-3.1-8B	<b>54.07</b>	67.74	52.36	70.46	61.21	42.23	49.71	43.03	31.28
+ CoRe ( $\hat{k} = 2$ )	53.07	<b>71.93</b>	<b>53.36</b>	<b>72.12</b>	<b>63.54</b>	<b>77.55</b>	<b>51.94</b>	<b>44.7</b>	<b>37.98</b>
Mistral-7B	<b>32.19</b>	43.68	27.31	39.74	37.22	26.47	21.87	14.18	9.41
+ CoRe ( $\hat{k} = 2$ )	36.04	<b>52.74</b>	<b>37.37</b>	<b>45.94</b>	<b>48.51</b>	<b>53.47</b>	<b>25.99</b>	<b>15.96</b>	<b>11.22</b>
Qwen2.5-7B	65.93	68.01	49.29	84.83	55.37	77.38	45.33	38.18	36.53
+ CoRe ( $\hat{k} = 2$ )	<b>65.67</b>	<b>72.03</b>	<b>51.09</b>	<b>85.59</b>	<b>59.97</b>	<b>83.96</b>	<b>49.94</b>	<b>45.21</b>	<b>40.25</b>
Phi-3.5-mini	<b>44.58</b>	59.76	44.92	81.44	48.48	81.14	28.53	23.39	19.49
+ CoRe ( $\hat{k} = 2$ )	46.23	<b>62.4</b>	<b>46.01</b>	<b>83.26</b>	<b>51.23</b>	<b>82.83</b>	<b>34.98</b>	<b>31.24</b>	<b>30.61</b>
GPT-4o-mini	<b>59.29</b>	71.74	<b>55.75</b>	<b>78.91</b>	76.44	<b>88.46</b>	55.06	49.68	41.84
+ CoRe ( $\hat{k} = 2$ )	57.68	<b>74.04</b>	54.93	77.53	<b>77.65</b>	88.38	<b>56.8</b>	<b>51.41</b>	<b>44.35</b>

Table 6: Chain-of-Thought results of F1 score in the multi-hop QA tasks. All models are instruction-tuned LLMs.

---

**Chain-of-Thought reasoning****[User role]**

Question: {*Question*}

Documents:

Document [0] {*Document 0*}

Document [1] {*Document 1*}

...

Document [N] {*Document N*}

Answer the question based on the given documents.

**[Assistant role]**

Sure. Before answering the question, I'll reconsider the question and the documents once more.

Question: {*Question*}

Documents:

Document [0] {*Document 0*}

Document [1] {*Document 1*}

...

Document [N] {*Document N*}

**[User role]**

Now answer the question based on the documents.

**[Assistant role]**

Let's think step by step.

---

**Extract short answer from CoT response**

(After same prompt of user-assistant-user as above prompt for chain-of-thought reasoning ...)

**[Assistant role]**

Let's think step by step. {*CoT response*}

**[User role]**

Respond only the answer in a few words after 'Answer:'.

**[Assistant role]**

Answer:

---

Table 7: Prompt template of CoRe used in Chain-of-Thought reasoning for the multi-hop QA tasks.

---

**[User role]**

Decompose the following question into several sub-questions.

Question: {*Question*}

**[Assistant role]**

1.

---

Table 8: Prompt template of query decomposition for the retrieve-and-reason task.

Models	Compositional	Comparison	Inference	Bridge-Comparison
Llama-3.1-8B	38.1	55.77	40.75	34.86
+ CoRe in assistant	48.15	58.13	44.81	65.11
+ CoRe in user	<b>53.45</b>	<b>69.36</b>	<b>47.59</b>	<b>68.76</b>
Mistral-7B	18.82	32.9	29.77	15.43
+ CoRe in assistant	27.78	35.67	36.39	<b>29.09</b>
+ CoRe in user	<b>44.07</b>	<b>50.7</b>	<b>39.64</b>	25.29
Qwen2.5-7B	41.36	65.47	37.26	39.45
+ CoRe in assistant	<b>54.2</b>	<b>70.3</b>	<b>42.4</b>	64.54
+ CoRe in user	54.05	69.73	41.64	<b>68.3</b>
Phi-3.5-mini	27.7	45.26	27	35.19
+ CoRe in assistant	37.64	50.92	32.34	36.71
+ CoRe in user	<b>42.56</b>	<b>54.89</b>	<b>34.51</b>	<b>54.37</b>

Table 9: Results in the ablation study about the position of context repetition

Methods	Compositional	Comparison	Inference	Bridge-Comparison
Baseline (CoRe)	48.15	58.13	<b>44.81</b>	65.11
Paraphrase	48.23	61.17	42.32	57.19
Summary	48.44	61.68	40.35	64.79
Shuffle	<b>51.25</b>	<b>63.4</b>	43.33	65.97
Reverse	50.65	63.35	42.9	<b>67.11</b>

Table 10: Results for Llama in the ablation study about the content of repeated contexts.

---

**[User role]**  
Question: {*Question*}

Documents:  
Document [0] {*Document 0*}  
Document [1] {*Document 1*}  
...  
Document [N] {*Document N*}

Answer the question based on the given documents. Respond only the answer within a few words after 'Answer:'.

Look again the input prompt:

Question: {*Question*}

Documents:  
Document [0] {*Document 0*}  
Document [1] {*Document 1*}  
...  
Document [N] {*Document N*}

Answer the question based on the given documents. Respond only the answer within a few words after 'Answer:'.

**[Assistant role]**  
Answer:

---

Table 11: Prompt template of CoRe applied in the user role position.



---

**[System role]**

You are a professional paraphraser. Your task is to paraphrase the given text based on the below instructions. Follow the instructions to achieve the desired output.

- Objective: Rewrite the text more thoroughly, changing both vocabulary and sentence structure while preserving the original meaning.
- Instructions: (MOST IMPORTANT) Use your own style for natural paraphrase. Introduce new expressions, alter sentence structure, and rearrange clauses. Use synonyms and change the form of words (e.g., verbs to nouns, or active to passive voice). Retain the original message but express it in a noticeably different way.

- Example:

Original: "The quick brown fox jumps over the lazy dog."

Paraphrase: "With swift movements, the brown fox leaps over the dog lying lazily."

# Format of the paraphrasing task

- Original: The original text.
- Paraphrase: The paraphrased version of the text based on the above guideline. Provide the output text immediately.

**[User role]**

Paraphrase the original text below.

Original: {input text}

**[Assistant role]**

Paraphrase:

---

Table 12: Prompt template of paraphrasing for the ablation study.

---

**[System role]**

Summarize the following text while ensuring that no key information, factual accuracy, or essential meaning is lost. Follow these guidelines:

- Maintain Key Details: All critical points, facts, and arguments from the original text must be preserved.
- Conciseness: The summary should be significantly shorter than the original text while capturing its essence.
- Clarity and Precision: Use clear, professional language. Avoid vague phrasing.
- No Alteration of Meaning: Do not add, alter, or infer information that is not present in the original text.

Please follow below pattern of example. Provide the output text immediately.

- Example:

Original Text: The rapid development of artificial intelligence over the last decade has led to significant breakthroughs in various fields, including healthcare, finance, and transportation. However, these advancements also raise concerns about data privacy, job displacement, and the ethical use of AI technologies.

Summary: AI advancements in healthcare, finance, and transportation have been substantial, though concerns about data privacy, job displacement, and ethical issues have emerged.

**[User role]**

Summarize the following text below.

Original Text: {input text}

**[Assistant role]**

Summary:

---

Table 13: Prompt template of summarizing for the ablation study.