

# NEEDLECHAIN: Measuring Intact Context Comprehension Capability of Large Language Models

Hyeonseok Moon, Heuseok Lim<sup>†</sup>

Department of Computer Science and Engineering, Korea University  
{glee889, limhseok}@korea.ac.kr

## Abstract

Recent reports suggest that LLMs can handle increasingly long contexts. However, many existing benchmarks for context understanding embed substantial query-irrelevant content, which shifts evaluation toward retrieving relevant snippets rather than fully integrating all provided information. Under this setting, we view that current benchmarks can overestimate true context-understanding ability of LLMs. In particular, we demonstrate that when the context consists entirely of query-relevant text, even advanced models such as GPT-4o fail to reliably integrate inputs as short as 200 tokens. To evaluate this capability more rigorously, we introduce NEEDLECHAIN, a benchmark designed to test whether models can faithfully incorporate all given evidence. NEEDLECHAIN includes three variants that differ in the required order of comprehension, along with a parallel benchmark based on the needle-in-a-haystack(NIAH) paradigm. By comparing these variants, NEEDLECHAIN enables a more comprehensive assessment of context understanding. We further propose a training-free strategy that encourages models to reflect all available information, ROPE contraction, highlighting the importance of full-context integration and pointing to new directions for improving reliable reasoning over context.

## 1 Introduction

Over the past few years, we have witnessed significant advancement of large language models (LLMs) in various aspects (Yang et al., 2025a; Hurst et al., 2024; Team et al., 2025). Enhanced reasoning abilities now allow LLMs to tackle several challenges and expand the range of tasks they can address (Liu et al., 2024a; Team, 2025). In particular, recent work has substantially extended the maximum context length of LLMs, enabling them




Figure 1: Performance comparison between the NEEDLECHAIN (Backward chain) and its parallel NIAH paradigm benchmark (NeedleStack). Reported number of tokens were estimated with Qwen2.5 tokenizer.

to handle more complex tasks (Tworkowski et al., 2023; Ge et al., 2025; Jin et al., 2024; Ding et al., 2024). For example, Llama-2 (Touvron et al., 2023) supports a 4,096-token context window, whereas Llama-4 (Meta, 2025) reportedly scales to over million tokens. These advances suggest markedly improved ability to model and utilize long-range dependencies in the input.

However, we question whether LLMs have reached to the point where they can reliably recognize and interpret all provided context. Numerous context understanding benchmarks evaluate LLMs by presenting a lengthy passage that query-relevant and query-irrelevant information coexist (Bai et al., 2024). The needle-in-a-haystack (NIAH) paradigm pushes this to an extreme by filling the context with misleading content while providing only a small amount of query-relevant evidence (Li et al., 2024b; Schuster et al., 2025; Wang et al., 2024). We argue that such settings can overstate contextual understanding: evaluation might prioritize *retrieval of the necessary snippet* rather than verifying whether the model tracks and integrates *all* information in the context.

<sup>†</sup> Corresponding Author

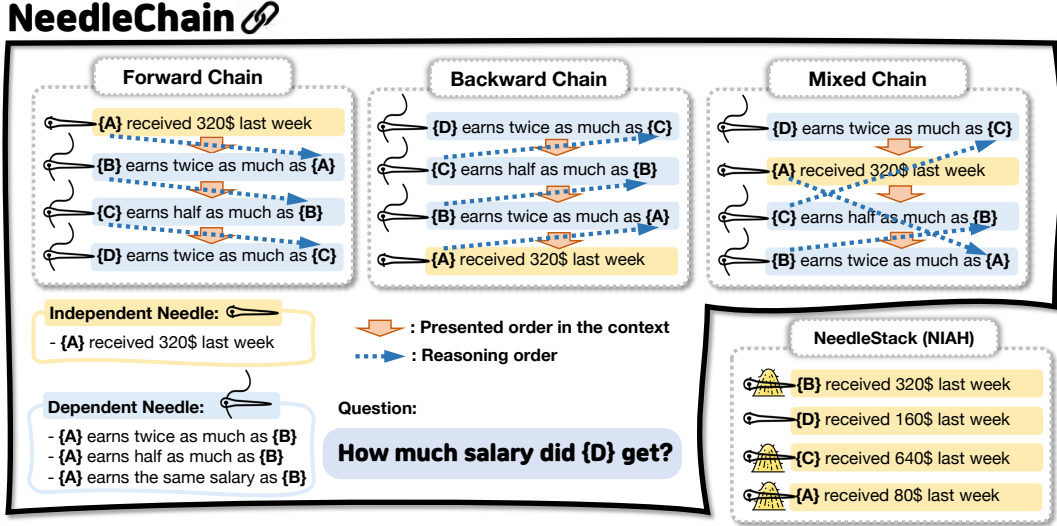


Figure 2: Performance variation with respect to the domain composition of training data

In this paper, we cast a simple yet overarching question: "When the context contains only query-relevant information, can an LLM fully incorporate and understand it?" Prior studies have notified a large gap between the reported context length and the effective length models can use, under the NIAH paradigm (Kuratov et al., 2024; Hsieh et al., 2024). We go further by characterizing the intact context-understanding length, i.e., the maximum context length over which the model preserves the ability to reflect the entire context without omissions.

In particular, we introduce **NEEDLECHAIN**, a benchmark that rigorously evaluates holistic context comprehension. In **NEEDLECHAIN**, each statement forms part of a causal chain and remains necessary for deriving the correct answer; omitting any single element prevents a correct solution. We build the benchmark from short, tightly connected statements (e.g., "A received \$1,600 last week" and "A earns twice as much as B,") that the model must jointly integrate to answer a query. Specifically, we formalize **reasoning order**, which denotes the logical sequence in which a model must process contextual information to resolve a query. We instantiate three variants: *forward chain* (requires left-to-right integration), *backward chain* (requires right-to-left integration), and *mixed chain*. Beyond these variants, we construct an NIAH dataset, **NeedleStack**, from the same benchmark component (needles). Figure 2 illustrates representative examples, and subsequent sections describe the construction process in detail. Comparing these variants exposes

previously overlooked weaknesses in LLMs' ability to preserve and integrate all the context intactly. Notably, figure 1 contrasts the performance on **NEEDLECHAIN** (backward) benchmark with the **NeedleStack** (NIAH) benchmark. Our findings reveal that even with a context of just 200 tokens, LLMs often fail to reliably retain and combine all information required to answer the query.

Through detailed error analysis, we identify factors that hinder faithful context understanding. We further propose a **ROPE** contraction strategy as a simple yet effective method to enhance context comprehension. This strategy sets a larger **ROPE** (Su et al., 2024) embedding rotation angle at inference than at training time, thereby sharpening positional distinctions and strengthening context utilization. Overall, we argue that beyond simply extending a model's context length, its ability to fully utilize the provided context constitutes a critical factor.

## 2 NEEDLECHAIN

**NEEDLECHAIN** is designed as an information-dense context comprehension task. In each context, every sentence contains essential content, requiring the model to fully understand all details for successful completion. Below, we discuss the details related to data construction.

### 2.1 Needle Design

We design our benchmark to analyze the weaknesses of LLMs in processing context by treating the entire context as a single semantic unit and

defining a reasoning order required to understand the information provided in the context. To achieve this, we introduce the concept of a "needle" to create connections between preceding and succeeding information. We define two types of needles:

- **Independent Needle** This refers to a sentence that contains independent information without relying on other context. ("*A received \$1600 last week*")
- **Dependent Needle** This refers to a sentence that provides information in conjunction with other sentences. It is designed to create an extended context while maintaining a coherent semantic unit. ("*A earns twice as much as B*" / "*A earns half as much as B*" / "*A earns the same salary as B*")

In particular, we design three types of dependent chains—"halving," "doubling," and "retaining"—that require clear yet straightforward reasoning. This design allows us to evaluate the reasoning capabilities of LLMs without underestimating their core language comprehension abilities due to their lack of precision in complex mathematical calculations, such as decimal operations. We combine these needles according to the specific objectives of our benchmark to form the final evaluation suites.

## 2.2 Chain Composition

Each data point (chain) of **NeedleChain** consists of one independent needle and  $k - 1$  dependent needles. By leveraging these needles, the combined chain is designed to maintain a single sequential reasoning order (e.g. *A is related to B, B is related to C, C is related to D*). We then design a question about the last needle in the reasoning order to ensure all context is relevant to the query (*How much salary did D get?*). Specifically, we propose three variants as follows. Figure 2 illustrates examples for each chain composition.

- **Forward Chain** The reasoning process must proceed in a **left-to-right** sequence. Accurate conclusions can only be drawn by following the presented order of the given input.
- **Backward Chain** The reasoning process must follow a **right-to-left** sequence. LLM must track the contextual information in the reversely-presented order, starting from the most recently presented data.

- **Mixed Chain** The sequence of required reasoning steps is set **arbitrarily**. The LLM must identify the randomly given reasoning order to arrive at the correct answer.

We design three chain variants using identical needle composition, distinguished solely by the sequence of provided information, thereby ensuring the same reasoning order while different presented order across chain variants. We determine our benchmark's evaluation suite by posing questions related to the needle at the end of reasoning order. This approach ensures our benchmark is structured so that answering the questions correctly requires 1) following the designated path and 2) fully understanding the given context. Comparing the performance of these variants enables clearly characterizing the limitations of LLMs in contextual understanding.

Furthermore, we use our needles to construct data under the NIAH paradigm, which concatenates multiple independent needles into a single context. In this setting, when we ask a question about one needle, all remaining needles serve as irrelevant context (i.e., the haystack). We refer to this benchmark as "NeedleStack (NS)" By comparing these variants of **NEEDLECHAIN**, we effectively demonstrate context-comprehension capabilities of LLMs and clearly distinguish performance on sparse contexts versus information-dense contexts.

## 2.3 Benchmark Details

In constructing the needle, we utilized a randomly selected name list officially released by the U.S. government<sup>1</sup>. For each data point, we first established a name list, then used this list to create three chain variants and a NeedleStack. This approach allowed us to apply the same name list across our benchmark variants, thereby minimizing unintended bias related to naming (Eloundou et al., 2025).

As our benchmark uses synthetic data, we can adjust the context length by varying the number of needles. Under this setup, we investigate context comprehension in LLMs by examining performance as the total number of needles, denoted by  $k$ , increases up to 200. While we theoretically could increase  $k$  further, even relatively short contexts (approximately 2K tokens) already revealed

<sup>1</sup><https://www.ssa.gov/oact/babynames/decades/century.html>

Model	k=5 (Token Length: 0.05K)				k=10 (Token Length: 0.1K)				k=20 (Token Length: 0.2K)			
	NS	NeedleChain			NS	NeedleChain			NS	NeedleChain		
		F	B	M		F	B	M		F	B	M
Qwen3-32B	100.0	100.0	99.5	100.0	100.0	99.5	91.5	98.0	99.5	94.0	65.0	86.0
Qwen2.5-32B	100.0	100.0	99.5	99.5	100.0	98.0	91.5	96.5	100.0	95.0	60.0	89.5
QwenLong-L1	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	98.5	92.5	99.5
Llama3.3-70B	100.0	100.0	99.5	99.5	100.0	100.0	90.5	98.5	99.5	94.0	63.5	91.0
GPT-4o	100.0	100.0	100.0	99.5	100.0	100.0	95.5	98.5	100.0	98.0	78.5	88.5

Model	k=50 (Token Length: 0.5K)				k=100 (Token Length: 1K)				k=200 (Token Length: 2K)			
	NS	NeedleChain			NS	NeedleChain			NS	NeedleChain		
		F	B	M		F	B	M		F	B	M
Qwen3-32B	99.5	80.5	13.5	46.5	99.5	68.5	10.5	10.5	99.5	47.5	1.5	4.5
Qwen2.5-32B	100.0	81.0	19.5	44.5	100.0	65.5	7.0	16.5	100.0	43.0	0.5	8.5
QwenLong-L1	100.0	88.5	46.0	76.0	99.5	86.5	21.5	39.0	99.5	67.8	6.5	11.0
Llama3.3-70B	100.0	76.5	31.0	79.0	100.0	67.5	22.5	45.0	99.5	44.0	18.0	12.5
GPT-4o	100.0	76.5	20.5	61.0	100.0	36.0	7.0	26.0	98.0	14.0	4.0	5.0

Table 1: Performance of several LLMs on NeedleChain (NS: NeedleStack, F: Forward chain, B: Backward, M: Mixed Chain).

clear trends in the models’ context understanding capacity that we aim to analyze. Thus, we conducted experiments using only 5 to 200 needles. We generated 200 test instances for each dataset.

### 3 Experiments

We conduct experiments using state-of-the-art LLMs reported to possess long context understanding capabilities. In particular, we focus our experiments on the widely used LLMs: Qwen2.5-32B (Yang et al., 2024), QwenLong-L1 (Yang et al., 2024), Qwen3-32B (Yang et al., 2025a), Llama3.3-70B (Grattafiori et al., 2024), and GPT-4o (Hurst et al., 2024). Detailed information about the models, prompts used for evaluation, and the evaluation environment is provided in the Appendix A.

#### 3.1 Main Results

We first evaluate the performance of current LLMs on our NEEDLECHAIN and NeedleStack benchmarks. The experimental results are presented in Table 1. Key insights from this study are as follows:

##### Limitations in Long-Context Understanding

All LLMs achieve near-perfect performance on NeedleStack (the NIAH paradigm). However, accuracy on our NEEDLECHAIN benchmark begins to drop once  $k$  exceeds 10 and cannot remain high when  $k$  reaches 50. This pattern suggests that models still fail to fully exploit context even when the total sequence length is only 0.5K tokens. Although the LLMs we evaluate nominally support much longer context windows (GPT-4o: 16K, Qwen2.5-32B: 32K, Qwen3-32B: 32K, QwenLong-L1: 1M, Llama3.3-70B: 128K), NeedleChain exposes clear

limitations in their practical context understanding that existing benchmarks do not reveal.

##### LLMs Struggle to find solution on Backward Chain

We observed a particularly pronounced performance decline in the backward chain. This finding indicates that the LLM’s ability to use context depends strongly on the reasoning direction, revealing a clear vulnerability in reverse reasoning. The backward chain even showed a larger performance drop than mixed chains with arbitrary reasoning paths. The higher performance in mixed chains suggests that text which merely *appears complex* may not substantially challenge LLMs. Instead, the main difficulty lies not in the seemingly intricate reasoning paths, but in the requirement to reason in the reverse direction itself.

##### Forward-direction reasoning is suitable for LLMs

Among the three chain variants in the NEEDLECHAIN benchmark, comprehension in the forward direction is notably high. This suggests that LLMs can comfortably comprehend the given context when its logical order aligns with their left-to-right processing. This observation not only reveals limitations of current LLMs, but indicates that we can maximize their reasoning capabilities by presenting information sequentially.

#### 3.2 Error Analysis

To gain a deeper understanding of the weaknesses in LLMs’ context comprehension as identified by the NEEDLECHAIN benchmark, we analyze the error cases revealed by our benchmark. We categorize the errors in LLMs into three distinct types and find that this taxonomy successfully encompasses

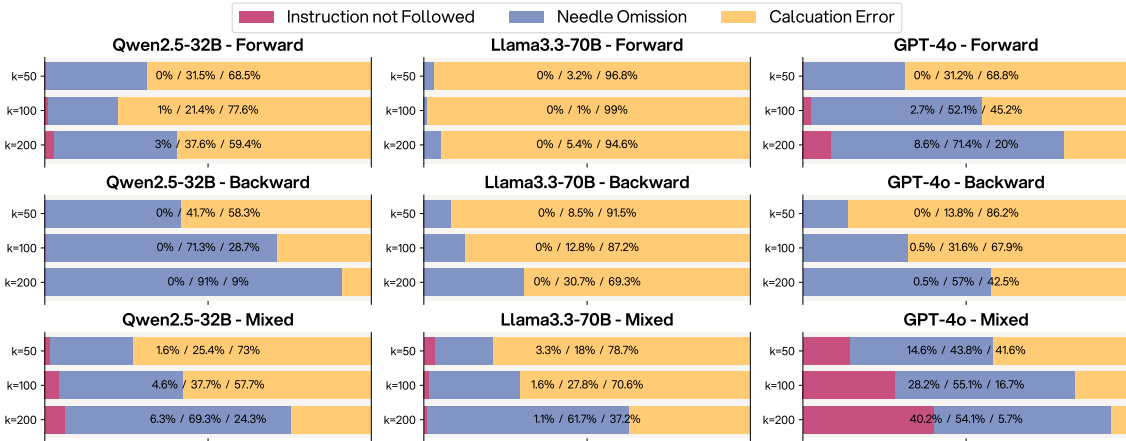


Figure 3: Error analysis on NEEDLECHAIN. We analyze errors in each category to determine which of the three predefined error types they fall into.

all observed error cases:

- **Instruction not Followed:** This refers to instances where the model fails to generate a response by not adhering to the given output format, or fails to determine the final answer.
- **Needle Omission** This refers to cases when certain "needles" are omitted in generating final answer. Specifically, it refers to cases where the name provided as input is absent in the output, resulting in an incorrect answer.
- **Calculation Error** This pertains to situations where the intermediate steps are correct, but an error occurs in the final answer computation. If an error does not fall into the first two categories, we classify such case here.

Concrete examples for each error type are contained in the Appendix B. Using this taxonomy, we analyze error cases made by LLMs in our benchmark. Figure 3 displays the results. Key takeaways from our analysis are summarized as follows:

**For small k:** We find that calculation errors are the dominant source of error, particularly when k is small (k=50, in this experiment). We can observe this phenomenon consistently across all models, regardless of the reasoning order. Considering the substantial forward-backward performance gap in previous experiments (Table 1), this result further underscores the strong influence of reasoning order on calculation capability. For Llama3.3-70B, for instance, the forward-backward performance gap exceeds 40 percentage points at k=50, and the

results in Figure 3 indicate that this gap stems entirely from reduced calculation capability induced by the change in reasoning order.

**For larger k:** As we increase k, needle omission emerges as a primary source of error, particularly in the Mixed chain. Across all models, the proportion of errors due to needle omission increases approximately linearly with k. This implies that the model fails to incorporate a clearly identifiable portion of the given context, and we observe that this failure mode already appears frequently once the input length exceeds only 1K tokens (k=100). This pattern suggests that achieving robust context understanding remains an open challenge. We provide a more detailed analysis of this error type in a subsequent section.

### 3.3 Position Heatmap

In this section, we utilize NEEDLECHAIN to identify positional weaknesses of LLMs in the context comprehension. Note that our input consists of multiple needles, each carrying key information about a specific "name." Considering the final answer can only be derived when the information from all needles is reflected, if any name is omitted from the LLM's response, it indicates that the corresponding needle was not considered during generation.

Through this approach, we identify positional weaknesses by determining whether the names given in the input are included in the generated text. We specifically analyze these positional weaknesses from two perspectives: presented order and reasoning order. We represent the ratio of names mentioned at each position relative to the total number of responses. The experimental results are

shown in Figure 4.

The results show that LLMs struggle to fully capture information from the input, as reflected by the unstable pattern of red regions in the heatmap, which highlights the “**logically lost-in-the-middle**” phenomenon (Liu et al., 2024b). We observe this most clearly in the mixed-chain setting. When we evaluate positions by their presented order in the context, we see little position-specific weakness: performance decreases fairly uniformly across all positions. In contrast, when we evaluate positions by reasoning order, the model’s ability to exploit information drops sharply at the “middle position”. This suggests a practical takeaway: LLMs tend to *get lost in the middle of the logical flow, rather than in the middle of the given context*.

### 3.4 Case Study: Question Variants

One can, in principle, construct an evaluation that requires aggregating all information in the given context simply by adjusting the query. For instance, even when using NeedleStack(NS), a question like “What is the total salary of all mentioned individuals?” requires the language model to integrate all context information. Building on this intuition, we propose two types of questions to more deeply analyze context understanding and to distinguish performance across these settings. Figure 5 reports the experimental results.

The experimental results demonstrate that when using  $q_{total}$ , the performance of NS significantly declines, with scores dropping below 10 for all LLMs at  $k = 20$ . Notably, when using  $q_{total}$ , the performance on the NEEDLECHAIN benchmark exceeds that of the NS benchmark. This indicates that **deeper semantic connections (i.e., explicit reasoning paths within the context) enhance comprehension more than less interrelated contexts**. In a subsequent section, we further show how the design of NEEDLECHAIN provides additional insight into the context comprehension abilities of LLMs.

### 3.5 Case Study: Tool Incorporation

In Figure 3, we identified calculation errors as one of the critical factors contributing to the performance degradation of our benchmark. To address such error, recent approaches have sought to mitigate the mathematical limitations of LLMs by incorporating code implementation capabilities (Liao et al., 2024; Sharma et al., 2025). In this section, we aim to test whether computational deficiencies

cause us to underestimate limitations in language comprehension by evaluating performance after code integration. Experimental results are detailed in Table 2.

Type	Question	k	Code Merging
<b>NeedleStack</b>			
NIAH	$q_{total}$	5	96.0 → 94.0 ( ▽ -2.0 )
NIAH	$q_{total}$	10	27.5 → 53.0 ( △ +25.5 )
NIAH	$q_{total}$	20	5.0 → 73.0 ( △ +68.0 )
NIAH	$q_{total}$	50	6.5 → 72.0 ( △ +65.5 )
<b>NeedleChain</b>			
Forward	$q_{single}$	50	76.5 → 74.0 ( ▽ -2.5 )
Forward	$q_{single}$	100	36.0 → 42.5 ( △ +6.5 )
Backward	$q_{single}$	50	20.5 → 23.5 ( △ +3.0 )
Backward	$q_{single}$	100	7.0 → 5.0 ( ▽ -2.0 )
Mixed	$q_{single}$	50	61.0 → 54.0 ( ▽ -7.0 )
Mixed	$q_{single}$	100	26.0 → 21.5 ( ▽ -4.5 )

Table 2: Performance variation on tool merging. We report performance of GPT-4o model.

Interestingly, tool incorporation proved to be particularly effective in NeedleStack with  $q_{total}$ , where contextual information between segments is weakly correlated. The improvement from an apparent lack of understanding to performance above 70 points indicates that tool incorporation can be partially beneficial.

However, such incorporation did not prove effective in NEEDLECHAIN. We observed performance declines in most cases. This underscores the robustness of our benchmark, demonstrating that it presents challenges that can only be resolved with comprehensive long-context reasoning abilities. This finding suggests that the low performance of LLMs in NEEDLECHAIN is not solely due to computational limitations but rather reflects a deficiency in effectively integrating contextual information.

## 4 Discussion

Based on these discussions, we conclude that LLMs do not yet fully comprehend given contexts. **We argue that rather than hastily increasing the extent of context length, it might be beneficial to enhance comprehension within a limited range.** Accordingly, we revisit one of the strategies currently used to increase context length, namely ROPE extension (Zhong et al., 2025; Ding et al., 2024; Chen et al., 2023). This approach modifies the rotary positional embeddings to theoretically extend the maximum usable context length. A com-

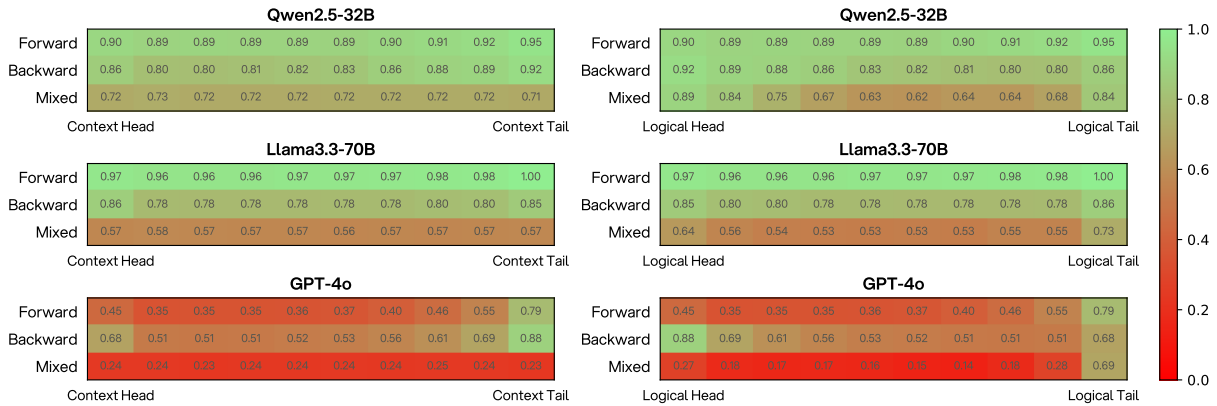


Figure 4: Heatmap to show the weaknesses for each position. Left-sided figures shows positional needle-missing heatmap with respect to the “presented order”. Right-sided figures shows those of “reasoning order”. We conducted experiments with k=200.

mon instantiation reduces the rotation angle used at inference time relative to that used during training. However, we hypothesize that such strategies might saturate positional distinctions and thereby substantially degrade the model’s ability to intactly utilize information across the entire context.

Conversely, we also hypothesize that sharpening positional distinctions can mitigate omissions of information within the given context, and we therefore propose ROPE contraction, a concept opposite to ROPE extension. A simple example is presented in Figure 6. For a simple toy experiment, we apply rotation angles that are 2x and 4x larger than those used during training to emphasize positional distinction.

We compare our results to the commonly used ROPE extension method known as Yarn (Peng et al., 2024). The results of these experiments are shown in Figure 7. The experimental results clearly support our argument. The use of Yarn significantly decreases performance on the NEEDLECHAIN benchmark, while the contraction method substantially improves the model’s ability to use the full context faithfully. This paper demonstrates that even simple methodologies can lead to meaningful performance improvements, paving the way for future research in this area.

The experimental results clearly support our argument. The use of Yarn significantly decreases performance on the NeedleChain benchmark, while the contraction method substantially improves the ability to understand context intactly. This paper demonstrates that even simple methodologies can lead to meaningful performance improvements, paving the way for future research in this area.

## 5 Related Works

There are currently numerous approaches to objectively assess the context understanding capabilities of LLMs (Bai et al., 2024; Hsieh et al., 2024; Kuratov et al., 2024; An et al., 2024; Li et al., 2024a; Zhang et al., 2024). Among these, the Needle-in-a-Haystack benchmark is a widely used tool to evaluate long-range context understanding in LLMs (Yu et al., 2025; Song et al., 2025). However, this benchmark typically uses contexts in which most of the information is irrelevant to the query, which limits its ability to assess comprehensive understanding. As a result, it tends to evaluate only partial and shallow comprehension rather than full understanding of the given contexts (Hsieh et al., 2024; Kuratov et al., 2024). While approaches such as (Li et al., 2024b) have attempted to evaluate holistic understanding, they remain rudimentary and fail to provide a thorough analysis of the insights such methods can offer. In particular, these approaches often require models to simply identify ancestors, which only demands shallow reasoning and therefore falls short of providing a complete evaluation. In response, we propose a novel benchmark that evaluates holistic context comprehension and requires models not to miss any piece of evidence in the given context. Our benchmark and experimental results show that comprehensive understanding within a given context length is a high-stakes task that current models still struggle to handle.

## 6 Conclusion

In this work, we introduced NEEDLECHAIN, a context comprehension benchmark in which every part of the context contains query-relevant information.




Figure 5: We compare the accuracy of models for different types of questions: those requires understanding the tail of the reasoning chain ( $q_{single}$ ) and those requiring comprehensive understanding of the entire context ( $q_{total}$ ).

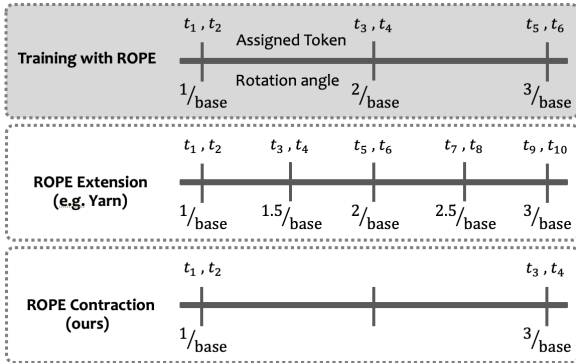


Figure 6: A simple diagram illustrating the concepts of ROPE extension and contraction

We designed NEEDLECHAIN so that the model cannot reach the correct answer if it misses even a single piece of evidence. Our experiments show that even when only 10 such evidence units exist, LLMs still struggle to fully comprehend the context. By constructing forward, backward, and mixed chain variants, we provided a comprehensive analysis of LLMs’ context understanding. Our results reveal that LLMs perform markedly worse on backward chains, where the reasoning order proceeds right-to-left, than on mixed-order contexts. We further proposed ROPE contraction, a simple yet compelling method to enhance context understanding abilities, achieving significant performance improvements on the NEEDLECHAIN benchmark. Our research indicates substantial room for improvement in the

context comprehension capability of LLMs. Additionally, our analysis offers practical advice, suggesting that designing reasoning orders in a forward direction is beneficial when establishing long contexts. For future research, we aim to design benchmarks that encompass a wider range of domains.

## 7 Limitation

Due to resource constraints, we were unable to conduct extensive experiments with reasoning models. The excessive length of the input for the reasoning model made it impossible to perform benchmark evaluations under our limited resources. For instance, in a scenario with  $k=100$ , the QwQ model (Team, 2025) required over 30 minutes to process a single query in our experimental setup. Given the need to process 200 queries for a single test, conducting a wide range of experiments with the inference model was impractical for us. Instead, we report the experimental results for  $k=50$  in the Appendix C and publicly release the data generation code to enable experiments with any higher  $k$ . We hope this will facilitate broader experimentation using our data in the future.

One limitation of our study is the exclusive use of needles requiring numerical calculation. However, through rigorously designed controlled experiments, we have clearly and robustly demonstrated our conclusions within the given setting. We plan




Figure 7: Performance variation derived by the ROPE contraction and extension methodologies

to extend our benchmark paradigm to enhance its generalizability in future research.

### Ethics Statement

All names referenced in the dataset are fictional, as noted in Section 2.3, and are merely borrowed for illustrative purposes. We affirm that all salaries and names mentioned in the data bear no connection to real-world individuals. Additionally, any potentially referenced names contain no harmful content whatsoever. The models and methodologies employed adhere to community-accepted ethical practices and are not designed for harmful, malicious, or discriminatory purposes. We believe this work aligns with responsible AI research principles and does not introduce foreseeable risks of misuse.

### Acknowledgement

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2021R1A6A1A03045425) This work was supported by the Commercialization Promotion Agency for R&D Outcomes(COMPA) grant funded by the Korea government(Ministry of Science and ICT)(2710086166) This work was supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (RS-2024-00398115, Research on the reliability and coherence of outcomes produced by Generative AI) This work was supported by Institute of Information & communications Technology Planning & Evaluation(IITP) under the Leading Generative AI Human Resources Development(IITP-2026-RS-2024-00397085) grant funded by the Korea government(MSIT)

### References

- Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2024. [L-eval: Instituting standardized evaluation for long context language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14388–14411, Bangkok, Thailand. Association for Computational Linguistics.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [LongBench: A bilingual, multi-task benchmark for long context understanding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3119–3137, Bangkok, Thailand. Association for Computational Linguistics.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. Longrope: extending llm context window beyond 2 million tokens. In *Proceedings of the 41st International Conference on Machine Learning*, pages 11091–11104.
- Tyna Eloundou, Alex Beutel, David G. Robinson, Keren Gu, Anna-Luisa Brakman, Pamela Mishkin, Meghan Shah, Johannes Heidecke, Lilian Weng, and Adam Tauman Kalai. 2025. [First-person fairness in chatbots](#). In *The Thirteenth International Conference on Learning Representations*.
- Suyu Ge, Xihui Lin, Yunan Zhang, Jiawei Han, and Hao Peng. 2025. [A little goes a long way: Efficient long context training and inference with partial contexts](#). In *The Thirteenth International Conference on Learning Representations*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and Boris Ginsburg. 2024. [RULER: What’s the real context size of your long-context language models?](#) In *First Conference on Language Modeling*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. 2024. Llm maybe longlm: Selfextending llm context window without tuning. In *Proceedings of the 41st International Conference on Machine Learning*, pages 22099–22114.
- Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. 2024. [Babilong: Testing the limits of llms with long context reasoning-in-a-haystack](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 106519–106554. Curran Associates, Inc.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024a. [LooGLE: Can long-context language models understand long contexts?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16304–16333, Bangkok, Thailand. Association for Computational Linguistics.
- Mo Li, Songyang Zhang, Yunxin Liu, and Kai Chen. 2024b. Needlebench: Can llms do retrieval and reasoning in 1 million context window? *arXiv preprint arXiv:2407.11963*.
- Minpeng Liao, Chengxi Li, Wei Luo, Wu Jing, and Kai Fan. 2024. [MARIO: MATH reasoning with code interpreter output - a reproducible pipeline](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 905–924, Bangkok, Thailand. Association for Computational Linguistics.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12.
- AI Meta. 2025. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>, checked on, 4(7):2025.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2024. [YaRN: Efficient context window extension of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Thomas Schuster, Marian Lambert, Nico Döring, and Julius Trögele. 2025. Needle-in-the-haystack testing llms with a complex reasoning task. In *International Conference on Engineering Applications of Neural Networks*, pages 254–266. Springer.
- Aditya Sharma, Aman Dalmia, Mehran Kazemi, Amal Zouaq, and Christopher Pal. 2025. [GeoCoder: Solving geometry problems by generating modular code through vision-language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 7340–7356, Albuquerque, New Mexico. Association for Computational Linguistics.
- Mingyang Song, Mao Zheng, and Xuan Luo. 2025. [Counting-stars: A multi-evidence, position-aware, and scalable benchmark for evaluating long-context large language models](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3753–3763, Abu Dhabi, UAE. Association for Computational Linguistics.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Qwen Team. 2025. Qwq-32b: Embracing the power of reinforcement learning.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Szymon Tworkowski, Konrad Staniszewski, Mikołaj Pacek, Yuhuai Wu, Henryk Michalewski, and Piotr Miłoś. 2023. Focused transformer: Contrastive training for context scaling. *Advances in neural information processing systems*, 36:42661–42688.
- Weiyun Wang, Shuibo Zhang, Yiming Ren, Yuchen Duan, Tiantong Li, Shuo Liu, Mengkang Hu, Zhe Chen, Kaipeng Zhang, Lewei Lu, Xizhou Zhu, Ping Luo, Yu Qiao, Jifeng Dai, Wenqi Shao, and Wenhui Wang. 2024. [Needle in a multimodal haystack](#). In *Advances in Neural Information Processing Systems*,

volume 37, pages 20540–20565. Curran Associates, Inc.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, et al. 2025b. Qwen2.5-1m technical report. *arXiv preprint arXiv:2501.15383*.

Yifei Yu, Qian-Wen Zhang, Lingfeng Qiao, Di Yin, Fang Li, Jie Wang, Zengxi Chen, Suncong Zheng, Xiaolong Liang, and Xing Sun. 2025. Sequential-niah: A needle-in-a-haystack benchmark for extracting sequential needles from long contexts. *arXiv preprint arXiv:2504.04713*.

Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024. [∞Bench: Extending long context evaluation beyond 100K tokens](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15262–15277, Bangkok, Thailand. Association for Computational Linguistics.

Meizhi Zhong, Chen Zhang, Yikun Lei, Xikai Liu, Yan Gao, Yao Hu, Kehai Chen, and Min Zhang. 2025. Understanding the rope extensions of long-context llms: An attention perspective. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 8955–8962.

## A Evaluation Details

We conducted all experiments using eight RTX A6000 GPUs. We performed decoding with a temperature setting of 0.6 and a top-p of 0.95, as these parameters represent the optimal prompt suggested in Qwen3<sup>2</sup>. We utilized publicly available checkpoints from HuggingFace (Wolf et al., 2020) for all models and executed inference using the vllm (Kwon et al., 2023) framework. Configuration for each model is as follows: Qwen2.5-32B (Yang et al., 2024) (Qwen/Qwen2.5-32B-Instruct), Qwen3-32B (Yang et al., 2025a) (Qwen/Qwen3-32B), QwenLong-L1 (Yang et al., 2025b) (Tongyi-Zhiwen/QwenLong-L1-32B), Llama3.3-70B (Grattafiori et al., 2024) (meta-llama/Llama-3.3-70B-Instruct), QwQ-32B (Team, 2025) (Qwen/QwQ-32B), GPT-4o (Hurst et al., 2024) (gpt-4o-2024-08-06). The prompt employed for model evaluation is as follows:

## B Qualitative Analysis

In this section, we extend the analysis from Section 3.2 and provide specific examples for each error instance. Table 4 shows examples for each error category we defined, derived from GPT-4o outputs generated using the Mixed chain. Here, we also provide examples of the chain compositions we define. Each chain consists of a common context, a chain component, and a common question, and we treat the concatenation of these three elements as a single data point.

**Instruction not followed** also includes potential needle omission error. In this case, the model receives the relevant information but treats it as a free variable when producing the final answer. For convenience, we categorize such cases as instruction not followed because the model does not produce a deterministic final answer.

**Needle omission** denotes cases where the model reaches a final answer but does not mention any needle in its intermediate reasoning. We label these responses as incorrect **regardless of the correctness of the final answer**, in order to exclude lucky guesses. We justify this choice as follows: the explanation explicitly lays out the full reasoning process, so if the reasoning omits the needle, the model has not actually used the provided context.

<sup>2</sup><https://huggingface.co/Qwen/Qwen3-32B>

---

**## System Prompt**

You are a financial assistant AI skilled in calculating wages and solving salary-related queries.  
I will give you context with the facts about salary of several people.  
You need to answer the question based only on the information from the facts.  
Before you derive the final answer, provide me a brief explanation.  
Output your final verdict by strictly following this format: '## Answer: \${your\_answer}'

**## Input Template**

There are {num\_names} workers in the office.  
Their names are as follows: {names}

Salary for each worker is as follows:  
{chain}

Now, respond to my question:  
{question}

---

Table 3: The default prompt for evaluating NeedleChain

Most **calculation errors** arise from mistakes involving powers of two. In a few rare cases, the model outputs a number that is not divisible by a power of two.

## C Performance of Reasoning Model

Reasoning models demonstrate exceptional problem-solving for a variety of tasks. In this study, we analyze the performance of reasoning models on our benchmark. The experimental results are presented in Table 5. Along with accuracy, we also report the length of the generated responses. Similar to existing models, reasoning models exhibit diminished performance in backward chain. Although the decline is less pronounced than that in traditional LLMs, we observe a clear trend of performance degradation with larger  $k$ . This indicates that achieving intact understanding remains a consistent challenge even for reasoning models.

## D LLM Usage Disclaim

We utilized an AI writer solely for polishing the English expressions in the paper. Beyond this, no assistance was received in terms of ideation or other content development; the AI writer’s role was strictly limited to correcting English expressions.

Chain Composition		
<b>Common Context:</b>	There are 5 workers in the office. Their names are as follows: <b>Cairo</b> , <b>Ramon</b> , <b>Nala</b> , <b>Kase</b> , <b>Ryker</b>	
<b>[Forward Chain]</b> Cairo received \$1600 last week. Ramon earns twice as much as Cairo. Nala earns the same salary as Ramon. Kase earns the same salary as Nala. Ryker earns twice as much as Kase.	<b>[Backward Chain]</b> Ryker earns twice as much as Kase. Kase earns the same salary as Nala. Nala earns the same salary as Ramon. Ramon earns twice as much as Cairo. Cairo received \$1600 last week.	<b>[Chaotic Chain]</b> Ramon earns twice as much as Cairo. Ryker earns twice as much as Kase. Nala earns the same salary as Ramon. Cairo received \$1600 last week. Kase earns the same salary as Nala.
<b>Common Question:</b>	Now, respond to my question: How much salary did <b>Ryker</b> get?	
Error Category		
Instruction not Followed	Needle Omission	Calculation Error
...(explanation)... ## Answer: Amount adherent context depends suggestively Jayceon’s half per explicit finalized quotas otherwise prescribed, rather undetermined under fact via customary series.	...(three person missing)... ## Therefore, since Kyren equals Justin, which equals Salma, which is equal to Ronin: \$200. ## Answer: \$200	...(explanation)... ## Answer: 1600 \$ (reference answer: 3200 \$)
...(explanation)... Thus, Miriam’s salary in terms of (x) is: ## Answer: 4x	...(five person missing) Therefore, Wrenleigh receives a salary of \$12.5. ## Answer: \$12.50	...(explanation)... ## Answer: 200 \$ (reference answer: 6400 \$)
...(explanation)... deductible as a function of referential salary "X" (Alora), as leaving Sol’s clear deduction numerically unexpressed specific to definitively concrete values. Without Alora or alternate equitably lead measure detail.	...(five person missing)... - Dayana earns twice as much as Wren. Therefore, Dayana earns \$25. 18. **Bridget and Dayana: - Bridget earns half as much as Dayana. Therefore, Bridget earns \$12.5. ## Answer: \$12.5	...(explanation)... ## Answer: 2200 \$ (reference answer: 3200 \$)

Table 4: Data samples for each chain composition, along with qualitative error analysis.

Model	k=10 (Token Length: 0.1K)				k=20 (Token Length: 0.2K)				k=50 (Token Length: 0.5K)			
	NS	F	B	M	NS	F	B	M	NS	F	B	M
<b>Qwen2.5-32B</b>	100 (36.41)	100 (283.76)	91.5 (284.09)	93.5 (306.14)	100 (33.27)	97.0 (500.44)	53.0 (505.64)	76.0 (568.13)	99.5 (31.98)	84.5 (1072.82)	13.0 (1175.93)	25.5 (1282.09)
<b>QwQ-32B</b>	100 (300.95)	100 (618.09)	96.5 (1220.63)	99.5 (934.02)	99.5 (626.59)	93.5 (1036.64)	76.0 (2950.03)	91.5 (2196.24)	100 (875.13)	86.5 (2727.29)	19.0 (7938.16)	62.5 (6822.31)
<b>Qwen3-32B</b>	100 (283.97)	100 (492.95)	99.0 (573.08)	100 (655.75)	100 (339.8)	97.0 (826.15)	88.5 (1031.43)	96.5 (1141.44)	98.5 (422.35)	77.5 (1810.63)	24.0 (3001.62)	63.0 (3368.24)

Table 5: Performance of reasoning LLMs on NeedleChain (**NIAH**: Needle Stack, **F**: Forward chain, **B**: Backward, **M**: Mixed Chain). We report both accuracy and the token length of the generated text.