

# What are the Essential Factors in Crafting Effective Long Context Multi-Hop Instruction Datasets? Insights and Best Practices

Zhi Chen<sup>\*</sup> Qiguang Chen<sup>♡\*</sup> Libo Qin<sup>◇†</sup> Qipeng Guo<sup>♣</sup> Haijun Lv<sup>♣</sup>

Yicheng Zou<sup>♣</sup> Hang Yan<sup>♣</sup> Kai Chen<sup>♣</sup> Dahua Lin<sup>♣</sup>

<sup>♣</sup> Shanghai Artificial Intelligence Laboratory

<sup>‡</sup> Research Center for Social Computing and Interactive Robotics

<sup>‡</sup> Harbin Institute of Technology

<sup>◇</sup> School of Computer Science and Engineering, Central South University

chenzhi@pjlab.org.cn, qgchen@ir.hit.edu.cn

## Abstract

Recent advancements in large language models (LLMs) with extended context windows have significantly improved various tasks. To improve long-context capabilities, much work focuses on augmenting LLM’s capabilities with synthetic data. Existing methods often leverage the Self-Instruct framework to generate long-context instruction-tuning data. However, our preliminary experiments show that fewer than 35% of samples generated by Qwen-2<sub>72B</sub> are multi-hop, and over 40% exhibit poor quality, limiting comprehensive understanding and further research. To address this, we propose the Multi-agent Interactive Multi-hop Generation (MIMG) framework, which integrates a quality verification agent, a single-hop question generation agent, a multiple question sampling strategy, and a multi-hop question merger agent. This framework significantly improves data quality, with high-quality, multi-hop, and diverse data. Furthermore, we conduct a thorough analysis of document selection, question merging, and validation techniques through extensive experiments across various models. Our results demonstrate that synthetic high-quality long-context instruction data can enhance model performance, surpassing even models trained on larger amounts of human-annotated data. Our code and relevant data are available at: <https://github.com/WowCZ/LongMIT>.

## 1 Introduction

Recently, large language models (LLMs) with long-context windows have significantly improved tasks such as information extraction, question answering, and even complex planning scenarios (Liu et al., 2024a; Bai et al., 2024b; Hu et al., 2025; Xu et al., 2024b). Research on developing long-context LLMs has predominantly focused on extending

<sup>\*</sup> Equal Contribution

<sup>†</sup> Corresponding Author

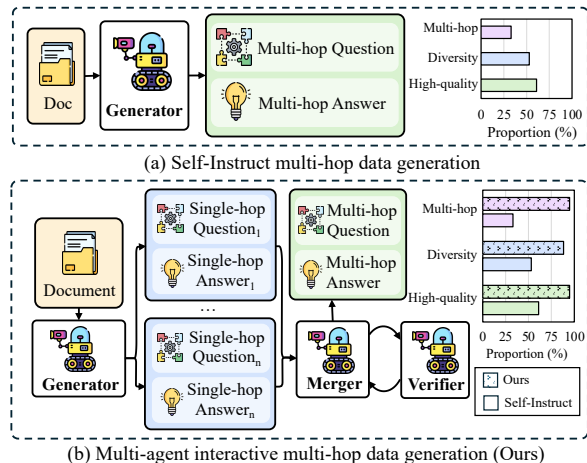


Figure 1: Comparison between traditional self-instruct-based data synthesis method and our Multi-agent Interactive Multi-hop Generation (MIMG) framework, where all data are generated by Qwen-2<sub>72B</sub> (Yang et al., 2024).

the context window (Ding et al., 2024; Jin et al., 2024; Peng et al., 2024). Nevertheless, in practical applications, merely expanding the context window is insufficient for effectively utilizing long-context (Hsieh et al., 2024; Huang, 2024), which presses a need for training to optimize utilization of long-context (Zhang et al., 2024a), especially in instruction-tuning (IT) (Fu et al., 2024b). In the IT phase, a large amount of high-quality long-context IT data is required. However, acquiring such data is challenging, with annotation costs significantly higher than those for short-context data (Bai et al., 2024b; Xiong et al., 2024). To address this, Xiong et al. (2023) and Bai et al. (2024a) have explored leveraging LLMs to generate IT data using the Self-Instruct framework (Wang et al., 2023b), thereby mitigating the scarcity of long-context IT data.

Moreover, the challenge often lies not in extracting single-hop information, but in integrating multiple hops of data from long contexts to derive complex conclusions. However, existing studies struggle to generate high-quality, multi-hop IT

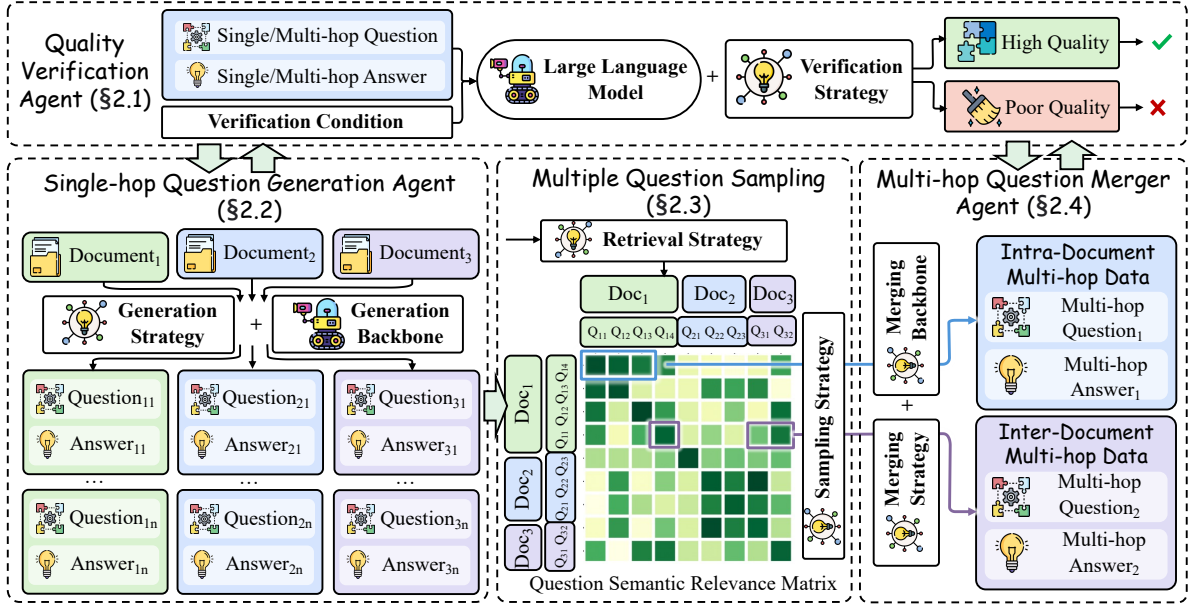


Figure 2: The overall process of Multi-agent Interactive Multi-hop Generation (MIMG) data synthesis framework.

data, primarily due to insufficient focus on the data synthesis process and the factors influencing data effectiveness. As illustrated in Figure 1 (a), our preliminary manual annotation experiments show that direct self-instruction yields less than 35% multi-hop samples, with high-quality examples representing only 60%. Additionally, sample diversity remains problematic, with over 45% of the samples exhibiting semantic duplication. These issues hinder comprehensive understanding and further advancement in this domain<sup>1</sup>.

Motivated by these challenges, this paper systematically investigates the research question: *What are the essential factors in crafting effective long-context multi-hop instruction datasets?* To address this, inspired by recent advancement on Agent (Hu et al., 2023; Wang et al., 2024), we propose a Multi-agent Interactive Multi-hop Generation (MIMG) framework. First, to ensure data quality, a Quality Verification Agent is introduced to evaluate the quality of long-context samples throughout the process. Second, for multi-hop reasoning, a Single-hop Question Generation Agent will be followed by a Multi-hop Question Merging Agent for stepwise synthesis of multi-hop queries. Finally, to ensure diversity, Multiple Question Sampling strategies are proposed to reduce redundancy and promote variety. To comprehensively examine the factors in long-context multi-hop data creation, we conduct extensive experiments, applying

17 strategies across 10 domains and 5 LLMs. As shown in Figure 1 (b), our method significantly improves data quality, yielding over 85% multi-hop, high-quality, and non-duplicative samples. Notably, LLMs trained on the synthetic high-quality data show an average improvement of 7.54%, even surpassing those LLMs trained on larger human-annotated datasets.

Overall, the main contributions are as follows:

- We systematically explore strategies for generating high-quality multi-hop instruction data to identify unexplored but critical factors, that influence the quality of long-context data. These factors include scoring verifiers, question-then-answer generation, question-based sampling, and question-answer merging strategies.
- We introduce the Multi-Agent Interactive Multi-hop Generation (MIMG) framework, which enhances the quality and relevance of synthesized data through multiple agent interactions.
- Our synthetic dataset, LongMIT, has shown superior performance across various long-context datasets. It not only improves long-context utilization but also surpasses larger human-labeled datasets, demonstrating its practical impact on advancing long-context LLMs.

## 2 Framework

Our framework consists of 4 main components: quality verification agent (QVA; § 2.1), single-hop question generation agent (SQGA; § 2.2), multiple

<sup>1</sup>See Appendix A for all detailed metric description.

question sampling (MQS; § 2.3), and multi-hop question merging agent (MQMA; § 2.4). Specifically, the QVA is first designed as a validator to control and supervise the data quality at each stage. The SQGA then generates simple and direct single-hop questions. Next, MQS strategies expand on this by sampling questions that cover various documents, enhancing multi-hop instruction generation. Finally, the MQMA integrates these single-hop questions into the multi-hop questions, requiring multiple document information. The detailed architecture is illustrated in Figure 2.

## 2.1 Quality Verification Agent

The first module is the Quality Verification Agent, which globally supervises and ensures the quality of generated samples. This component involves two main processes:

**Verification Strategy:** This includes additional heuristic strategies to judge which samples should be contained as high-quality data. Specifically, we explore two widely-used verification strategies:

- **Scoring:** We prompt LLMs to assign continuous scores and determine a threshold using the validation set to filter high-quality data. Formally, given a sample  $s$ , the selection criterion is:

$$\mathcal{V}(s|\mathcal{M}) = \begin{cases} \text{Approved} & \mathcal{F}_S(s|\mathcal{M}) > \theta; \\ \text{Rejected} & \mathcal{F}_S(s|\mathcal{M}) \leq \theta, \end{cases} \quad (1)$$

where  $\mathcal{F}_S(s|\mathcal{M})$  represents the score of sample  $s$  based on model  $\mathcal{M}$ , and  $\theta$  is the threshold.

- **Classification:** We prompt LLMs to perform binary classification and retain only samples classified as high-quality. Formally, given a sample  $s$ , the selection criterion is:

$$\mathcal{V}(s|\mathcal{M}) = \begin{cases} \text{Approved} & \mathcal{F}_C(s|\mathcal{M}) = 1; \\ \text{Rejected} & \mathcal{F}_C(s|\mathcal{M}) = 0, \end{cases} \quad (2)$$

where  $\mathcal{F}_C(s|\mathcal{M})$  represents the binary classification process of sample  $s$ .

**Verification Condition:** This involves setting specific conditions  $\mathcal{C}$  that both questions and answers must meet to be considered high-quality verification ( $\mathcal{V}(s|\mathcal{M}, \mathcal{C})$ ). The process includes:

- **Criteria Perspectives:** Criteria include relevance to the document, clarity, factual accuracy, logical coherence, and complexity of the question and answer. Formally, these perspectives

can be formulated as:

$$\mathcal{C} = \{c_1, \dots, c_n\}, \quad (3)$$

where  $c_i$  denotes the  $i$ -th criteria instruction.  $n$  denotes the number of criteria perspectives.

- **Auxiliary Context Information:** We integrate additional contextual instructions to enhance the model’s accuracy and robustness, like guidelines. These conditions are formally represented as:

$$\mathcal{C} = \{c_1, \dots, c_n\} \oplus \text{Context}, \quad (4)$$

where the `Context` denotes the context including auxiliary guidelines.

- **Auxiliary Generation Information:** We enable the model to provide reasoning rationale during output generation and observe its effectiveness.

$$\mathcal{C} = \{c_1, \dots, c_n\} \oplus I_R, \quad (5)$$

where the  $I_R$  denotes the instruction that can prompt LLM to generate rationales.

## 2.2 Single-hop Question Generation Agent

This phase generates single-hop questions and answers from individual documents through the following components:

**Generation Backbone:** We need a robust LLM to generate multiple single-hop questions and answers per document, ensuring diversity for data synthesis. Therefore, we evaluate various LLMs, both open- and closed-source, across different scales.

**Generation Strategy:** Our strategy employs structured techniques to extract potential questions:

- **Rationale-based Question Generation:** Chain-of-Thought (CoT) can effectively enhance performance on long-context tasks (Li et al., 2024). Hence, we explore whether generating rationale-supported questions improves the model’s grasp of a document’s reasoning structure.
- **Question-Answering Generation Order:** Furthermore, we assess whether order affects output quality. Generating questions before answers may simplify reasoning and enhance output clarity compared to simultaneous generation.

## 2.3 Multiple Question Sampling

To further optimize the diversity of generated samples, we introduce MQS strategy, which constructs multi-hop questions by sampling and combining questions from various documents. This approach consists of two key strategies:

**Retrieval Strategy:** This strategy selects relevant questions and documents for multi-hop question generation. By leveraging relevance sampling, it constructs a question-semantic relevance matrix to assess semantic connections across different documents, guiding the sampling process. The strategy comprises:

- **Probability-Based Sampling:** This method samples data based on probability-based document relevance, which is calculated as the frequency of keywords related to the questions, like BM25 (Robertson et al., 1995, 2009), and LDA (Hoffman et al., 2010).
- **Semantic-Based Sampling:** This approach assesses the relevance by analyzing the semantic similarity between questions and documents, like embedding similarity.

**Sampling Strategy:** Based on the relevance matrix, the most related questions should be selected to form a coherent, contextually rich multi-hop question. The strategy includes:

- **Intra-Document Sampling:** It focuses on selecting questions within the same document to ensure internal coherent multi-hop data.
- **Inter-Document Sampling:** This strategy involves selecting questions from different documents to ensure a broader contextual coverage.

## 2.4 Multi-hop Question Merging Agent

The final step merges sampled questions into coherent multi-hop questions, involving two modules:




**Merging Backbone:** We should employ LLMs to synthesize sampled questions and answers into meaningful multi-hop queries. To investigate this, we conduct the following exploration:

**Merging Strategy:** This strategy applies rules and heuristics to maintain logical and contextual consistency. It includes:

- **Document-Based Merging:** To further reduce input tokens, we explore whether incorporating long documents enhances merging performance. Formally, the process is:

$$Q_m = \mathcal{M}(Q_1, Q_2, \dots, Q_n | C), \quad (6)$$

where  $Q_1, Q_2, \dots, Q_n$  are sampled single-hop questions,  $Q_m$  is the merged multi-hop query, and  $C$  denotes whether documents are used.



(c) Analysis of the impact of precision on context length for GPT-4o.

Figure 3: The analysis of different verification strategies in quality verification, where includes 5 models: Qwen2-72B-Instruct (Yang et al., 2024); InternLM2-20B (Cai et al., 2024); Gemini-1.5-Pro (Reid et al., 2024); GPT-4o-mini and GPT-4o (Achiam et al., 2023).

- **Rationale-Based Merging:** To preserve the meaning and context of the individual questions, this method leverages the reasoning rationale behind the original questions to guide their merging process, which can be expressed as:

$$R \oplus Q_m = \mathcal{M}(Q_1, Q_2, \dots, Q_n), \quad (7)$$

where  $R$  represents the underlying rationale, and  $\oplus$  denotes the connecting elements in the generated response.

Additionally, we explore intra-document and inter-document multi-hop instruction samples for diverse scenarios.

## 3 Exploration

This section examines the framework components aimed at enhancing data quality, including verification strategies in QVA (§3.1), generation strategies in SQGA (§3.2), sampling approaches in MQS (§3.3), and merging strategies in MQMA (§3.4).

### 3.1 Quality Verification Agent

#### 3.1.1 Verification Strategy

Currently, the most commonly used model verification strategies are scoring and direct classification.



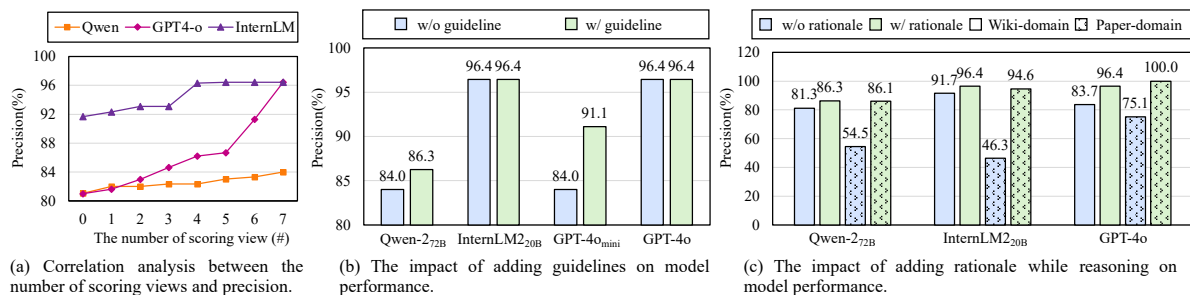


Figure 4: The analysis of different verification conditions on quality verification.

We assessed the consistency and accuracy of both methods by comparing them with human annotations in a sample analysis of long-context data.

**Scoring is a better verification strategy compared with classification.** As shown in Figure 3 (a), the scoring method significantly outperforms binary classification, yielding higher kappa and precision scores. This statistical improvement indicates that scoring better captures the subtleties of human judgment. These findings align with research in short-context scenarios (Fu et al., 2024a), highlighting the broader applicability of scoring methods across different context lengths.

**LLM is not a long-context annotator but a good selector.** As depicted in Figure 3 (a), unlike their performance in short-context verification (Wang et al., 2023a; Fu et al., 2024a), LLMs show minimal agreement with human annotators in long-context situations, reflected in low kappa scores. This suggests challenges in maintaining consistent annotations due to cognitive load and interpretive variations over extended data. However, as Figure 3 (b) shows, LLMs maintain near-perfect precision, demonstrating their strong ability to identify and select relevant information. This highlights LLMs’ potential as effective tools for data filtering and prioritization in long-context environments, in contrast to their role as accurate annotators in short-context settings.

**Scoring alleviates the long context bias but classification does not.** We further examine why classification performs poorly in long-context scenarios by analyzing precision across different context lengths. Figure 3 (c) illustrates that scoring consistently achieves higher precision and robustness in longer contexts, explaining the suboptimal performance of classification in these cases. Based on these findings, subsequent experiments will adopt the scoring strategy, using verifier precision and

the data retention ratio to assess data quality. More discussion can be seen in Appendix B.

### 3.1.2 Verification Conditions

To deeply understand what factors affect the verification of long text data quality, we further explored from three perspectives: scoring perspective, guidelines, and whether rationale is included for scoring.

#### More scoring criteria reduce long-context bias.

As shown in Figure 4 (a), incorporating more scoring criteria enhances the accuracy and robustness of filtering long-context data. Unlike short contexts, long contexts are prone to judgment bias. When fewer than three criteria are used, performance improvements are limited, with models often overestimating irrelevant samples. Increasing the number of criteria improves labeling accuracy, reducing biases linked to longer contexts (see Appendix C.4.2 for further details).

#### Effective verifiers adhere to annotation standards aligned with human judgment without guidelines.

To assess whether incorporating additional guidelines improves verification performance, we analyze the effectiveness of the method of integrating guidelines. Figure 4 (b) reveals that advanced verifiers do not require supplementary guideline information during annotation. This suggests that effective verifiers naturally adhere to annotation standards that align with human judgment.

#### Incorporating rationale enhances robustness in diverse long contexts.

By examining CoT (Wei et al., 2022; Qin et al., 2023), in various domains like wiki knowledge and paper analysis, we observe that incorporating rationale improves model performance across diverse long-context scenarios. As shown in Figure 4 (c), without rationale, performance drops by more than 8.6% across different domains. However, adding rationale during validation minimizes performance variation, with

precision fluctuations limited to 1.8%.

## 3.2 Single-hop Question Generation Agent

### 3.2.1 Generation Backbone

In practice, effective LLMs must be able to synthesize high-quality data. Thus, we evaluated several widely used LLMs for single-hop data synthesis.

**Open-source LLMs effectively generate single-hop questions.** As shown in Figure 5 (a), smaller open-source LLMs exhibit high retention rates and cost-effectiveness, demonstrating their ability to generate single-hop questions from a given context.

**Stronger LLMs can generate better single-hop question generation but higher cost.** As shown in Figure 5 (a), more advanced LLMs enhance data retention and question quality. However, these improvements are not cost-proportional, raising concerns about their economic feasibility for single-hop question generation.

### 3.2.2 Generation Strategy

Furthermore, we explore whether a question-then-answer approach, supported by rationale, improves the quality of synthetic single-hop questions.

**Question-then-answering works better than generating data from scratch.** To evaluate the effectiveness of single versus multi-stage generation, we compare two strategies: unified question-answer and question-then-answer. As shown in Figure 5 (b), generating the question first significantly improves data quality, particularly for open-sourced LLMs, increasing retention and quality scores. For implementation details, refer to Appendix C.4.3.

**Generating with rationale can improve the generated quality but much higher token cost.** As illustrated in Figure 5 (c), adding rationale makes questions more relevant and insightful with higher quality. However, the improvement brought by the rationale is minimal, while the token consumption triples, making it economically inefficient.

## 3.3 Multiple Question Sampling

### 3.3.1 Retrieval Strategy

This strategy focuses on identifying relevant documents and constructing a semantic relevance matrix to guide sampling. Observations include:

**Embedding similarity is critical for multi-question sampling.** We evaluate the effectiveness of various similarity measures using three

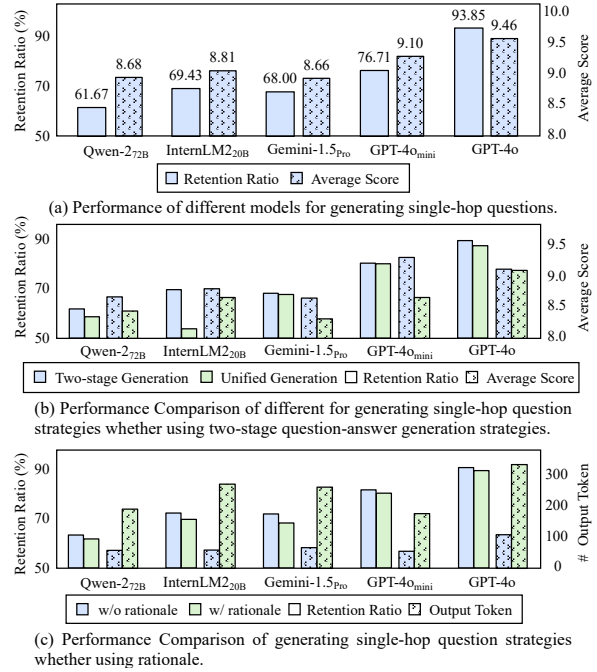


Figure 5: The analysis of generation backbone and generation strategies in SQGA.

metrics: embedding similarity (with BGE embeddings (Xiao et al., 2023)), BM25, and LDA. As shown in Figure 6 (a), BGE embeddings enable the model to select more relevant questions, thereby improving sample quality.


**Question similarity outweighs document similarity.** We further investigate the factors affecting sample quality. Figure 6 (b) reveals that question-based sampling significantly outperforms document-based approaches, as questions provide greater contextual relevance.

### 3.3.2 Sampling Strategy


This approach selects semantically related and complementary questions from within and across documents to create coherent, contextually rich multi-hop questions.

**Intra-Document generates better quality but less diversity.** As shown in Figure 6 (c), sampling questions within a single document yields more coherent and contextually aligned questions. However, this method may reduce diversity, as the questions are all drawn from the same source.


**Inter-Document generates less quality but more diversity.** As demonstrated in Figure 6 (c), sampling from multiple documents introduces a wider range of topics, enhancing diversity. However, this



(a) Performance Comparison of multiple question sampling based on different retrieval strategy.



(b) Performance Comparison of multiple question sampling based on document and question embedding.



(c) Performance Comparison of multiple question sampling based on different sampling strategies.

Figure 6: The analysis of multiple question sampling.

broader scope may diminish coherence and relevance due to the larger topic gaps.

### 3.4 Multi-hop Question Merging Agent

#### 3.4.1 Merging Backbone


We use LLMs to merge sampled questions and answers into coherent multi-hop versions, ensuring logical consistency and contextual accuracy with the aid of 5 classic LLMs. The key observations are as follows:

**Open-sourced LLMs can well merge multi-hop question generation.** As shown in Figure 7 (a), all models demonstrate strong capabilities in handling complex question-generation tasks requiring multi-step reasoning or information integration.


#### 3.4.2 Merging Strategy

**Question-answer pairs are enough for multi-hop instruction merging compared with documents.** To minimize input tokens, we examine whether long documents are necessary for improving merging performance. Figure 7 (b) shows that adding documents does not consistently enhance performance and instead increases input tokens. Thus, simple question-answer pairs are sufficient for effective multi-hop merging.


**Merging with rationale can not improve the merging quality.** Although generating content with rationales generally enhances quality (Qin



(a) Performance of different models for merging multi-hop questions.



(b) Performance Comparison of different for generating multi-hop question merger agent whether using document when question merging.



(c) Performance Comparison of generating multi-hop question merger agent whether using rationale.

Figure 7: The analysis of multi-hop question merging agent.




Figure 8: Comparison of the quality and token consumption on different generation strategies.

et al., 2023, 2024), Figure 7 (c) demonstrates that, unlike single-hop generation, rationales in multi-hop generation fail to contribute to coherent and logical question formation. Our analysis further reveals that large models often misinterpret rationales in queries and merging strategies, leading to frequent failures in CoT reasoning. Therefore, multi-hop synthesis should avoid using rationales.

## 4 Data Utilization

### 4.1 Data Synthesis Efficiency

Given the high cost of data generation, we balance cost and data quality in synthesizing long multi-hop instruction tuning (LongMIT) datasets (See Appendix C for more data details). To assess this balance, we compare the proportion of high-quality data and token cost for 200 samples generated under different strategies. At the input level, our inter-

Model	NarrativeQA	2WikiMQA	DuReader	HotpotQA	MultifieldQA <sub>en</sub>	MultifieldQA <sub>zh</sub>	MuSiQue	Qasper	AVG
InternLM2-1.8B (Cai et al., 2024)									
+ChatQA2	18.50	35.00	29.00	46.00	64.00	58.00	19.50	38.50	38.56
+LongAlign	25.00	33.00	25.00	49.50	<b>76.00</b>	67.50	24.50	44.00	43.06
+LongAlpaca	25.00	23.50	29.00	49.50	70.00	67.00	24.50	45.00	41.69
+NQ	17.00	25.50	33.50	35.00	60.00	67.00	14.50	44.00	37.06
+LongMIT	<b>26.00</b>	<b>35.50</b>	<b>60.00</b>	<b>56.00</b>	75.33	<b>75.50</b>	<b>29.00</b>	<b>47.50</b>	<b>50.60</b>
LLaMA3-8B (Dubey et al., 2024)									
+ChatQA2	24.00	41.00	50.00	49.00	64.00	69.00	26.00	51.50	46.81
+LongAlign	29.00	44.50	56.50	56.50	79.33	80.50	21.50	55.50	52.92
+LongAlpaca	18.00	50.00	48.00	55.50	76.67	80.00	27.50	<b>60.50</b>	52.02
+NQ	21.00	42.00	63.00	59.50	78.00	74.00	29.00	54.00	52.56
+LongMIT	<b>36.50</b>	<b>67.50</b>	<b>74.00</b>	<b>71.00</b>	<b>87.33</b>	<b>84.50</b>	<b>39.50</b>	54.00	<b>64.29</b>
InternLM2-7B (Cai et al., 2024)									
+ChatQA2	31.00	42.00	38.50	61.00	70.67	33.00	28.50	53.00	44.71
+LongAlign	45.00	40.00	60.00	65.50	74.67	86.00	34.00	56.50	57.71
+LongAlpaca	45.00	50.50	44.00	64.50	75.33	47.50	35.50	56.50	52.35
+NQ	12.50	37.50	61.50	45.50	75.33	77.00	21.00	57.50	48.47
+LongMIT	<b>46.50</b>	<b>57.00</b>	<b>74.00</b>	<b>73.00</b>	<b>91.33</b>	<b>91.00</b>	<b>45.00</b>	<b>62.00</b>	<b>67.48</b>

Table 1: Main accuracy results by evaluation by GPT-4o, where all benchmarks come from the LongBench (Bai et al., 2024b). More evaluation results on Ruler (Hsieh et al., 2024) are shown in Table 4.

actions increase the token count by approximately 3k tokens (roughly 2.5 times improvement), which is negligible in comparison to the substantial volume of long-context documents provided, averaging 70k tokens as shown in Figure 9. At the output level, the increase in token consumption is almost no noticeable increase, with an average increase of fewer than 0.5k tokens. Despite this minimal change in token usage, the quality of our results has improved nearly fourfold. As shown in Figure 8, strategies using open-source LLMs achieve a high-quality proportion comparable to the best strategies, at only one-third of the token cost. Additionally, our approach improves data quality with minimal token expense compared to traditional methods. Then, we will pad the context with additional documents to the target length to create sufficiently long samples. For further experimental results and details, see Appendix D. Moreover, additional exploration for long code instruction data generation can be found in Appendix G.

## 4.2 The Results of Instruction-Tuning

We conduct instruction-tuning on the synthesized data to evaluate its utility. As shown in Table 1, our data significantly improves long-context QA capabilities across various LLMs, with an average gain of at least 7.54%. Multi-hop benchmarks such as 2WikiMQA (Ho et al., 2020), MuSiQue (Trivedi et al., 2022), and HotpotQA (Yang et al., 2018) show more notable improvements. Detailed training procedures are in Appendix E. Furthermore, in

Appendix F, the high quality and logical complexity of this data enable the model to generalize to single-hop tasks not encountered during instruction tuning, confirming the reliability of MIMG.

## 4.3 Scaling Analysis

### 4.3.1 Data Scaling Analysis

To evaluate how the size of high-quality data affects model performance, we experiment on LLaMA3-8B (Dubey et al., 2024) by varying the training data volume. The results, depicted in Figure 10 in Appendix, show a clear relationship between the size of data and the performance. As the dataset size increases, model performance adjusts accordingly, demonstrating the significance of high-quality data scaling in enhancing the model efficacy.

### 4.3.2 Hop Scaling Analysis

To assess the impact of multi-hop data on model performance, we increased the number of hops in the dataset while keeping the training data volume constant. This approach isolated the effect of multi-hop reasoning on model outcomes. As indicated in Figure 11 in Appendix, there is a clear positive correlation between the number of hops and model performance. The data demonstrate that with more hops, the model achieves higher accuracy and robustness, demonstrating the effectiveness of using high-quality multi-hop data to enhance the model’s capability for complex reasoning tasks.



## 5 Related Work

Recent efforts have aimed to enhance the performance of LLMs in handling longer contexts (Hu et al., 2024; Liu et al., 2025). LongLLaMA (Xiong et al., 2023) demonstrates the impact of incorporating long text data during various pre-training stages. Then, LLaMA2-80K (Fu et al., 2024b) highlights the significance of using a domain-balanced, up-sampled long text corpus to improve long text capabilities, requiring only a 5B-token corpus for effective comprehension. Furthermore, ICLM (Shi et al., 2024) enhances long-text reasoning by transforming pre-training data into knowledge graphs and splicing adjacent documents. To improve the model’s ability to follow long text instructions, LongAlpaca (Chen et al., 2024c) combines a 9K paper question-answering (QA) corpus with 3K short instruction samples. In contrast, LongAlign (Bai et al., 2024a) utilizes Claude (Anthropic, 2023) to produce 10K QA pairs for training. Additionally, ChatQA (Liu et al., 2024b) enhances long-context QA performance by incorporating manually annotated data. Building on these approaches, ChatQA2 (Xu et al., 2024a) further incorporate existing long-text datasets, such as Natural Questions (NQ) (Kwiatkowski et al., 2019).

The method closest dataset is Quest (Gao et al., 2024), which constructs QA pairs from spliced documents, resulting in a close-sourced single-hop QA corpus. In contrast, our approach models document correlations first, then create multi-hop QA pairs using related intra-document data. Additionally, we offer systematic analysis, open-source datasets, and significantly improved models.

## 6 Conclusion

In conclusion, our proposed Multi-agent Interactive Multi-hop Generation (MIMG) framework, including a quality verification agent, a single-hop question generation agent, a multiple question sampling strategy, and a multi-hop question merger agent, achieves high-quality, diverse instruction data. Our experiments show that this synthetic data notably enhances performance, even surpassing larger human-annotated data, highlighting the effectiveness of our approaches.

## Limitations

Due to the high costs associated with large-scale distillation training experiments on GPT-4-based

methods, we focus our evaluation efforts on a small-scale assessment conducted through manual evaluation. To ensure robustness, the manual annotation is performed on a sample of 200 items. While the sample size is manageable, we acknowledge the potential for minor unavoidable biases inherent in random sampling. Moreover, considering the nature of randomization for LLMs, it is quite hard for us to strictly control LLMs to generate the same question for quality comparison.

## Ethical Considerations

Participants are recruited from universities across China, and all must have passed the CET-6 exam or achieved an IELTS score of 6 or higher. While participants come from diverse regions, we minimize the impact of national biases by focusing primarily on long context data. All annotators provided informed consent and were compensated above the local minimum wage. Additionally, no IRB review was required for the study.

The annotation process starts with an onboarding test, where participants answer 20 example questions. They receive \$42 for this phase, designed to familiarize them with the task. Annotators are then paid \$5 per hour, with a total of approximately 120 human hours dedicated to manual annotations. Overall, three experts are involved in the annotation and verification stages.

## Acknowledgments

Thanks to Hanqi Li, Yifei Yang and Da Ma for their constructive feedbacks.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2023. [Model card and evaluations for claude models](#).
- Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. 2024a. [Longalign: A recipe for long context alignment of large language models](#). *Preprint*, arXiv:2401.18058.
- 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. 2024b. [Longbench: A bilingual, multitask benchmark for long context understanding](#). *Preprint*, arXiv:2308.14508.
- Giovanni Maria Biancofiore, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, and Fedelucio

- Narducci. 2024. Interactive question answering systems: Literature review. *ACM Computing Surveys*, 56(9):1–38.
- Zheng Cai, Maosong Cao, Haojiong Chen, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. 2025a. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiaqi Wang, Mengkang Hu, Zhi Chen, Wanxiang Che, and Ting Liu. 2025b. Ecm: A unified electronic circuit model for explaining the emergence of in-context learning and chain-of-thought in large language model. *arXiv preprint arXiv:2502.03325*.
- Qiguang Chen, Libo Qin, Jiaqi WANG, Jingxuan Zhou, and Wanxiang Che. 2024a. [Unlocking the capabilities of thought: A reasoning boundary framework to quantify and optimize chain-of-thought](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu, and Wanxiang Che. 2024b. [M<sup>3</sup>CoT: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8199–8221, Bangkok, Thailand. Association for Computational Linguistics.
- Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. 2023. Exploring the use of large language models for reference-free text quality evaluation: An empirical study. In *Findings of the Association for Computational Linguistics: IJCNLP-AACL 2023 (Findings)*, pages 361–374.
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2024c. [LongLoRA: Efficient fine-tuning of long-context large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Tri Dao. 2024. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). In *The Twelfth International Conference on Learning Representations*.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [Flashattention: Fast and memory-efficient exact attention with io-awareness](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 16344–16359. Curran Associates, Inc.
- Yiran Ding, Li Lina Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [Longrope: Extending llm context window beyond 2 million tokens](#). *Preprint*, arXiv:2402.13753.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2024a. [GPTScore: Evaluate as you desire](#). 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 6556–6576, Mexico City, Mexico. Association for Computational Linguistics.
- Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hananeh Hajishirzi, Yoon Kim, and Hao Peng. 2024b. [Data engineering for scaling language models to 128k context](#). In *Proc. of ICML*.
- Chaochen Gao, Xing Wu, Qi Fu, and Songlin Hu. 2024. [Quest: Query-centric data synthesis approach for long-context scaling of large language model](#). *arXiv preprint arXiv:2405.19846*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Gal Ben Haim, Adi Braun, Haggai Eden, Livnat Burstein, Yiftach Barash, Avinoah Irony, and Eyal Klang. 2024. Ai in the ed: Assessing the efficacy of gpt models vs. physicians in medical score calculation. *The American Journal of Emergency Medicine*, 79:161–166.
- 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 *Proc. of COLING*, pages 6609–6625.
- Matthew Hoffman, Francis Bach, and David Blei. 2010. Online learning for latent dirichlet allocation. *advances in neural information processing systems*, 23.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, Yang Zhang, and Boris Ginsburg. 2024. [Ruler: What’s the real context size of your long-context language models?](#) *Preprint*, arXiv:2404.06654.
- Mengkang Hu, Tianxing Chen, Qiguang Chen, Yao Mu, Wenqi Shao, and Ping Luo. 2024. Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model. *arXiv preprint arXiv:2408.09559*.

- Mengkang Hu, Tianxing Chen, Yude Zou, Yuheng Lei, Qiguang Chen, Ming Li, Yao Mu, Hongyuan Zhang, Wenqi Shao, and Ping Luo. 2025. Text2world: Benchmarking large language models for symbolic world model generation. *arXiv preprint arXiv:2502.13092*.
- Mengkang Hu, Yao Mu, Xinmiao Yu, et al. 2023. [Tree-planner: Efficient close-loop task planning with large language models](#). *Preprint*, arXiv:2310.08582.
- Jerry Huang. 2024. How well can a long sequence model model long sequences? comparing architectural inductive biases on long-context abilities. *arXiv preprint arXiv:2407.08112*.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. 2024. [Llm maybe longlm: Self-extend llm context window without tuning](#). *Preprint*, arXiv:2401.01325.
- Seungone Kim, Se Joo, Doyoung Kim, Joel Jang, Seonghyeon Ye, Jamin Shin, and Minjoon Seo. 2023. [The CoT collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12685–12708, Singapore. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Sanwoo Lee, Yida Cai, Desong Meng, Ziyang Wang, and Yunfang Wu. 2024. Unleashing large language models’ proficiency in zero-shot essay scoring. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 181–198.
- Yanyang Li, Shuo Liang, Michael Lyu, and Liwei Wang. 2024. [Making long-context language models better multi-hop reasoners](#). In *Proc. of ACL*, pages 2462–2475.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024a. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Yijun Liu, Jinzheng Yu, Yang Xu, Zhongyang Li, and Qingfu Zhu. 2025. A survey on transformer context extension: Approaches and evaluation. *arXiv preprint arXiv:2503.13299*.
- Zihan Liu, Wei Ping, Rajarshi Roy, Peng Xu, Mohammad Shoeybi, and Bryan Catanzaro. 2024b. [Chatqa: Building gpt-4 level conversational qa models](#). *arXiv preprint arXiv:2401.10225*.
- 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*.
- Libo Qin, Qiguang Chen, Xiachong Feng, Yang Wu, Yongheng Zhang, Yinghui Li, Min Li, Wanxiang Che, and Philip S Yu. 2024. Large language models meet nlp: A survey. *arXiv preprint arXiv:2405.12819*.
- Libo Qin, Qiguang Chen, Fuxuan Wei, Shijue Huang, and Wanxiang Che. 2023. [Cross-lingual prompting: Improving zero-shot chain-of-thought reasoning across languages](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2695–2709, Singapore. Association for Computational Linguistics.
- Libo Qin, Qiguang Chen, Yuhang Zhou, Zhi Chen, Yinghui Li, Lizi Liao, Min Li, Wanxiang Che, and Philip S Yu. 2025. A survey of multilingual large language models. *Patterns*, 6(1).
- Machel Reid, Nikolay Savinov, Denis Teplyashin, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Weijia Shi, Sewon Min, Maria Lomeli, Chunting Zhou, Margaret Li, Xi Victoria Lin, Noah A. Smith, Luke Zettlemoyer, Wen tau Yih, and Mike Lewis. 2024. [In-context pretraining: Language modeling beyond document boundaries](#). In *The Twelfth International Conference on Learning Representations*.
- 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*, 10:539–554.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. [Is ChatGPT a good NLG evaluator? a preliminary study](#). In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 1–11, Singapore. Association for Computational Linguistics.
- Peng Wang, Yongheng Zhang, Hao Fei, Qiguang Chen, Yukai Wang, Jiasheng Si, Wenpeng Lu, Min Li, and Libo Qin. 2024. [S3 agent: Unlocking the power of vllm for zero-shot multi-modal sarcasm detection](#). *ACM Trans. Multimedia Comput. Commun. Appl.* Just Accepted.



- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023b. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Proc. of NeurIPS*.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. [C-pack: Packaged resources to advance general chinese embedding](#). *Preprint*, arXiv:2309.07597.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, et al. 2023. [Effective long-context scaling of foundation models](#). *Preprint*, arXiv:2309.16039.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, et al. 2024. [Effective long-context scaling of foundation models](#). In *Proc. of the NAACL*.
- Peng Xu, Wei Ping, Xianchao Wu, et al. 2024a. [Chatqa 2: Bridging the gap to proprietary llms in long context and rag capabilities](#). *arXiv preprint arXiv:2407.14482*.
- Yang Xu, Yunlong Feng, Honglin Mu, Yutai Hou, Yitong Li, Xinghao Wang, Wanjun Zhong, Zhongyang Li, Dandan Tu, Qingfu Zhu, et al. 2024b. [Concise and precise context compression for tool-using language models](#). *arXiv preprint arXiv:2407.02043*.
- An Yang, Baosong Yang, Binyuan Hui, et al. 2024. [Qwen2 technical report](#). *arXiv preprint arXiv:2407.10671*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proc. of EMNLP*, pages 2369–2380.
- Yikai Zhang, Junlong Li, and Pengfei Liu. 2024a. [Extending llms' context window with 100 samples](#). *Preprint*, arXiv:2401.07004.
- Yongheng Zhang, Qiguang Chen, Min Li, Wanxiang Che, and Libo Qin. 2024b. [Autocap: Towards automatic cross-lingual alignment planning for zero-shot chain-of-thought](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 9191–9200.
- Ziyu Zhuang, Qiguang Chen, Longxuan Ma, Mingda Li, Yi Han, Yushan Qian, Haopeng Bai, Zixian Feng, Weinan Zhang, and Ting Liu. 2023. [Through the lens of core competency: Survey on evaluation of large language models](#). *arXiv preprint arXiv:2308.07902*.

## Appendix

### A Metrics Utilized in Exploration

#### A.1 Metric Definition & Implementation

##### A.1.1 Diversity

Diversity measures the frequency of samples with the same semantics appearing in the data. In annotation, annotators sequentially identify whether new samples are semantically equivalent to previously annotated ones.

High diversity indicates a broad range of samples, ensuring that annotations do not repeat similar or identical meanings. This is essential for creating a dataset that represents various use cases and scenarios, covering a wide array of semantic topics. A diverse dataset contributes to a more robust LLM by capturing the nuances of language, context, and conceptual meaning across different long-context samples.

##### A.1.2 Multi-Hop

Multi-hop refers to the need for multiple information connections and integrations from various sources when handling complex samples. Here, annotation tasks require annotators to assess a query's needs by utilizing information from several contextual documents.

Effective multi-hop datasets pose questions that cannot be answered with a single data point but demand the combination of multiple facts or steps to reach the correct answer. Such reasoning is vital in real-world long-context tasks, such as answering questions that require complex deductions or understanding interconnected pieces of information.

##### A.1.3 High-Quality

High-quality annotations denote the accuracy, consistency, and relevance of the synthesis data. In high-quality datasets, annotators are required to judge whether each sample is precise and reliable, minimizing errors or inconsistencies.

Annotators must ensure that the data they provide accurately reflects the meaning or intent of the task at hand. In long-context NLP, high-quality data is crucial for developing models that make accurate predictions, recognize subtle patterns, and perform effectively in real-world scenarios.

#### A.2 The Impact of Different Metric

Furthermore, we conduct sampling and annotation of several instruction-tuning datasets based on three metrics. A comparison with Table 2 shows



Dataset Name	High-quality	Diversity	Multi-hop
LongAlpaca (Chen et al., 2024c)	70.3	87.9	50.7
LongAlign (Bai et al., 2024a)	87.7	83.8	52.6
ChatQA2 (Xu et al., 2024a)	83.1	80.3	30.4
NQ (Kwiatkowski et al., 2019)	90.2	63.2	10.3
LongMIT (Ours)	94.8	88.2	94.8

Table 2: Results of the data quality of different instruction datasets, reporting the data quality score as the average of three independent manual evaluations.

that the results require considering multiple factors and the specific needs of each task: NQ provides high-quality data, enhancing its effectiveness for in-domain tasks such as DuReader. However, its lack of diversity limits its performance in out-of-domain scenarios, which lowers its overall accuracy. In contrast, while LongAlign offers greater diversity, its quality is relatively lower, weakening its fundamental comprehension abilities. This is reflected in its significantly poorer performance on basic tasks like DuReader, resulting in suboptimal overall performance.

For the metric view, the conclusions are as follows: (1) High-quality annotations notably improve model performance in foundational long-context comprehension tasks, such as DuReader. (2) Annotation diversity is essential for enhancing model performance across a wide range of tasks. (3) The multi-hop property is particularly important for complex reasoning tasks that require integrating multiple long-context clues, as seen in datasets like HotpotQA and MusiQue.

In summary, we argue that quality ensures strong performance in fundamental long-context tasks, multi-hop promotes the complex multi-hop capabilities, while diversity improves performance across various fields by leveraging these basic capabilities.

### A.3 Automatic Metrics

#### A.3.1 Quality Score

Several studies have shown that large-scale models align with human judgment when generating quality scores (Chen et al., 2023; Chang et al., 2024; Lee et al., 2024). These models also perform well in specialized domains, such as medicine, where they show strong consistency with expert evaluations (Haim et al., 2024). Motivated by these findings, we prompt LLMs to automatically generate quality scores for scalable and effective long-context data quality verification.

To analyze the effectiveness of quality scores in long-context scenarios, we analyze the consistency between quality scores and human assessments.

Specifically, Figure 3 (a) demonstrates that the Kappa coefficient for the agreement between the scoring mechanism and human evaluators exceeds 0.50, significantly surpassing the performance of classification strategies that directly label items as high or low quality. Furthermore, Figure 3 (b) shows that the scoring mechanism achieves a precision level of 96.43% relative to human evaluators, underscoring its effectiveness as a robust data screening strategy.

#### A.3.2 Retention Rate

Retention rate refers to the proportion of data retained after being filtered by a quality verification agent based on a specified threshold. This concept arises from our observation that LLMs excel more as selectors than as annotators in Sec. 3.1.1, demonstrating higher retention rates and better alignment with human judgment. In complex or subtle situations, LLM judgments may vary, which can impact data interpretability and reliability. Therefore, we select models and strategies that closely align with human consistency, which achieves an almost perfect precision rate of 96.43%.

For implementation, we use a threshold of 8.5, determined through an internally annotated verification set of 200 items. The threshold yielding the highest precision in this set was chosen as the standard parameter for future evaluations.

Regarding the “average score”, our strategy assigns a quality score from 0 to 10 for each data point. After evaluating the entire dataset, we calculate the average score by averaging all individual quality scores.

## B Discussion of Scoring versus Classification

Although the document inputs are relatively long, averaging approximately 3K tokens, the classifier still assigns a “high-quality” label to nearly every sample of this length, even when we introduce multiple reference criteria. We believe that, once inputs exceed this length threshold, the classifier’s baseline score effectively rises above five, so any score over five is automatically deemed high quality. As a result, the classifier’s judgments lose discriminative power: every long sample is labeled “high-quality”.

By contrast, the scoring approach provides a more nuanced ranking. Even if all samples are of generally high quality, a score of 0.9 is still recognized as slightly better than 0.8. This finer distinction more closely reflects human evaluation.

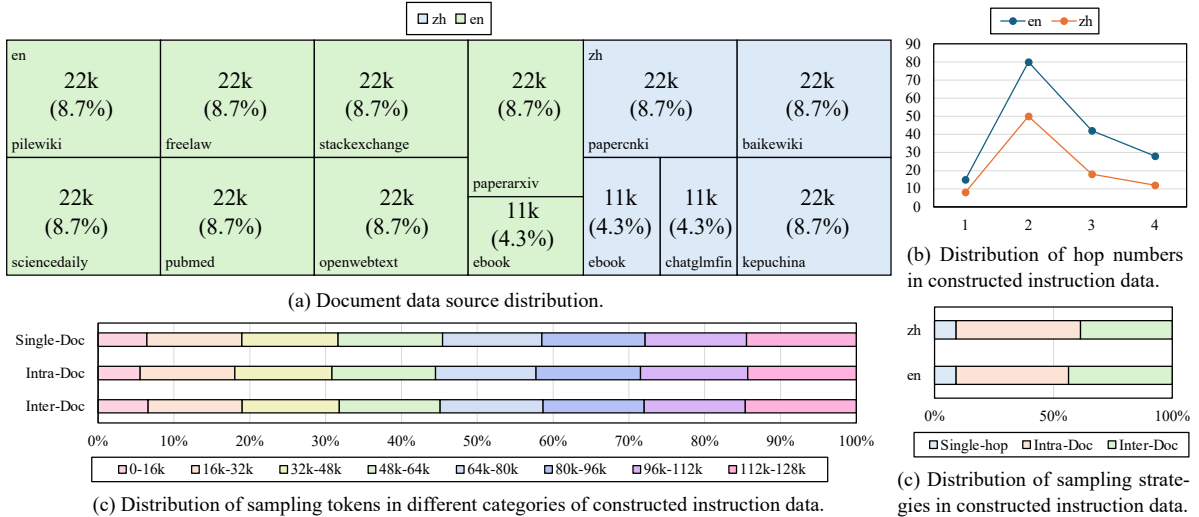


Figure 9: The analysis of constructed dataset distribution.

## C Data Construction Details

### C.1 Dataset Construction Pipeline Discussion

The current data in this field relies on Self-Instruct but lacks a systematic framework. Notably, many existing long-text works are essentially subsets of our MIMG. For instance, the prompts of Self-Instruction and those of the Single-hop Question Generation Agent are nearly identical. Additionally, Self-Instruction (+ *LLM recheck*) aligns with the architecture integrated into our Quality Verification Agent. Therefore, we assert that our approach offers a more thorough solution.

The construction of long-context multi-hop question-and-answer datasets is based on a structured approach leveraging pre-trained document corpora. This section outlines the methodology used for data collection, processing, and validation across multiple domains and languages.

### C.2 Source Data Overview

The primary source of long-text data is a pre-trained document corpus that spans nine distinct and most widely-used domains (Biancofiore et al., 2024). Inspired by Qin et al. (2025); Zhang et al. (2024b), the corpus includes data from the most widely-used bilingual sources (viz. Chinese and English), ensuring a comprehensive multilingual dataset. The domains covered are:

- **Books (eBooks):** A collection of various eBook formats that provide diverse literary content. Academic Papers: Scholarly articles sourced from repositories such as arXiv and CNKI. These datasets reflect cutting-edge research

across multiple disciplines.

- **Finance:** Data from financial documents and discussions, including the ChatGLM-fin dataset, which encompasses various financial reports and conversational data related to financial analysis.
- **Knowledge:** Information extracted from online encyclopedic sources, including Baike-Wiki and Pile-Wikipedia, covering a broad range of general knowledge.
- **Science:** Data from reputable scientific sources, including Kepuchina and ScienceDaily, that focus on advancements in various scientific fields.
- **Law:** Legal documents and case law from the Pile-Freelaw dataset, providing insight into legal precedents and interpretations.
- **Medicine:** Medical literature, including publications from Pile-PubMed Central, which includes peer-reviewed medical research and case studies.
- **Technology:** Content derived from technical discussions and knowledge-sharing platforms such as Pile-StackExchange.
- **Web Resources:** Web data extracted from open-source platforms, specifically the Pile-OpenWebText2 dataset, reflecting general web-based information.

Each domain was selected to ensure the inclusion of diverse, domain-specific content that could support the generation of robust and accurate multi-hop question-and-answer sequences. A more fine-grained analysis can be seen in Figure 9 (a).

Additionally, chain-of-thought (CoT) has be-

Dataset Name	Sample Size	Token Size
LongAlpaca (Chen et al., 2024c)	12.00k	0.11B
LongAlign (Bai et al., 2024a)	9.89k	0.17B
ChatQA2 (Xu et al., 2024a)	128.00k	1.22B
NQ (Kwiatkowski et al., 2019)	315.20k	4.56B
LongMIT (Ours)	64.40k	5.07B

Table 3: The statistics results of different datasets.

come a widely-used effective technique for various tasks (Kim et al., 2023; Chen et al., 2024b). Inspired by this, CoT has the ability to bring powerful performance improvements to the instruction tuning. What’s more, as shown in Figure 13, after adding CoT, the performance of the model has indeed improved, which following the recent reasoning conclusion (Guo et al., 2025; Zhuang et al., 2023; Chen et al., 2024a, 2025b). Therefore, in all our data synthesis processes, the answer contains a reasoning path. Furthermore, since LLMs often cannot fit all the document information that is extremely long documents, we perform truncation segmentation on the documents input to the model. After generating the sample, refill the document with other documents to a fixed length.

### C.3 Instruction Dataset Construction

To expand the domain coverage and handle longer contexts, we extend the instruction fine-tuning data across 9 domains and 2 languages. All base documents are sourced from pre-trained datasets to prevent data leakage. Our Long Multi-hop Instruction-Tuning dataset (LongMIT) results in a retention rate of over 90% in GPT-4o verification in 200 sampled samples, confirming the high quality and generalizability of our pipeline. To balance the cost and effectiveness of generating data, LongMIT are generated based on Qwen2-72B-Instruct and verified based on InternLM2-20B. We conduct a detailed statistical analysis of sample size and token consumption across various datasets.

Moreover, Table 3 presents the sample and token sizes. Our token count is comparable to that of NQ, but our dataset, while containing fewer samples, outperforms NQ by over 10%. Notably, despite NQ’s larger sample size, its performance lags behind datasets with fewer samples, such as LongAlpaca and LongAlign.

### C.4 Multi-hop Question and Answer Data Construction

The construction of multi-hop question-and-answer datasets involved a rigorous process to ensure both

linguistic accuracy and domain relevance. The methodology is as follows:

#### C.4.1 Dataset Curation

For each domain, data was independently curated to maintain a clear distinction between different knowledge sources. This allows for more focused and accurate multi-hop questions that are relevant to the particular field of study.

#### C.4.2 Quality Verification Agent

The first module in our framework is the Quality Verification Agent, which ensures that the generated questions and answers meet a certain standard of quality. We use InternLM2-20B (Cai et al., 2024) as the backbone and set the quality score threshold to 8.5. Moreover, the prompts are as follows:

Suppose you are a professional annotator, and you need to annotate the generated questions, rationales, and answers according to the context. Specifically, your tasks are as follows:

- First, determine whether the questions and answers are in documents provided in context.
- Then, you need to determine whether the problem is a multi-hop problem, using multi-hop logic.
- At the same time, you need to judge whether the question conforms to common-sense logic. Does the question conform to common sense in a normal context? Is the logic smooth?
- In addition, you need to rate the overall data quality from three aspects: logical rationality and fluency, question complexity, and answer clarity. All scores are between 0 and 10.
- Before giving an annotation, you need to give your rationale.

[[DOCUMENTS]]

{chunk}

[[QUESTION]]

{question}

[[ANSWER]]

{answer}

Finally, you should give me an overall quality mark in the format:  
““{ "in\_document": BOOL, "domain\_similarity": NUMBER, "quality": NUMBER }““

### C.4.3 Single-hop Question Generation Agent

The Single-hop Question Generation Agent is responsible for generating fundamental single-hop questions, which are characterized by their simplicity and directness.

In this framework, we employ Qwen2-72B-Instruct (Yang et al., 2024) as the foundational model, utilizing it to synthesize data through a question-answering paradigm. The process begins with the generation of prompts designed specifically for question creation, initiating a structured approach to the formulation of these queries.

The document content is as follows:

{chunk}

Extract the questions contained in the above document, and the extracted questions should meet the following conditions:

- No pictorial information should be included in the extracted questions;
- No referential information should be included in the extracted questions;
- Ensure the completeness of the extracted questions; if they are multiple-choice questions, provide corresponding option information, remove line breaks, and place the question body in a single question;
- If the document contains concepts such as numbers, time, people, or places, questions that involve this information must be extracted;
- The extracted questions should be presented in a parseable list format, such as ["xxx", "xxx"]. If there are no valuable questions, output an empty list [];
- Try to extract as many valuable questions as possible, but do not include duplicate questions;
- Extract no more than three questions;

Extracted questions:

Based on the questions extracted, the prompt for answer generation is as follows:

Generate answers to a given series of questions based on the content of the document, which must meet the following conditions:

- Respond based on the content in the document;
- If there is no corresponding answer to the question in the document, please reply based on your own knowledge;
- If the question is about factual issues such as numbers, time, people, places, etc., please provide the answer directly, and different question and answer pairs should be distinguished by line breaks;

The document content is as follows:

{chunk}

The problems are as follows:

{question}

The corresponding answers are as follows:

### C.4.4 Multiple Question Sampling

This strategy further enhances the generation of multi-hop instructions by selecting questions that address diverse elements within the document. This method facilitates the creation of comprehensive, multi-hop, long-text question-answer datasets that are meticulously customized to reflect the characteristics and requirements of specific domain data sources. The organization of the relevant documents begins by embedding them into vectors, where BGE-zh-1.5 and BGE-en-1.5 (Xiao et al., 2023) models are used to map the documents into 768-dimensional vectors. Following the methods inspired by Shi et al. (2024), the document vectors are embedded using Faiss to facilitate storage and efficient retrieval. This process relies on measuring vector distances to retrieve the 10 nearest documents for each document, creating a document graph.

Subsequently, a circular search strategy is employed to generate paths that consist of multiple documents, with the maximum path length constrained to 20. This process continues until all documents are sampled, with these paths serving as the initial sets of multiple related documents.

After conducting a sampling analysis, we observed the hop distribution in the constructed data,



as illustrated in Figure 9 (b). Additionally, the distribution corresponding to the sampling strategy is depicted in Figure 9 (c).

#### C.4.5 Multi-hop Question Merging Agent

Multi-hop questions are designed to require reasoning across multiple data points, either within a single domain or spanning different domains. This approach ensures that responses cannot be derived from isolated facts; rather, they necessitate a more profound comprehension and integration of the dataset’s overall content.

To achieve this, the Multi-hop Question Merging Agent consolidates single-hop questions into well-structured multi-hop queries. This process demands information synthesis from various sections of the document, promoting a deeper level of understanding and engagement. For the model architecture, we employ Qwen2-72B-Instruct (Yang et al., 2024) as the base model. The specific prompt for merging two QA pairs is as follows:

Based on the given two question-answer pairs, synthesize up to one question-answer pair that matches the real scenario. The synthesized question-answer pair should meet the following conditions:

- If both questions and answers are time-related, a comparative question can be synthesized to compare the order in which two events occur;
- If both questions and answers are related to the character, it can be synthesized to determine which character better fits the description of the composite question;
- The synthesized answer should provide the corresponding reasoning process, and the synthesized answer should make as much use of the content in the given two answers as possible;
- Do not arbitrarily change the original information of two questions and answers;
- The generated questions and answers are strictly output in JSON format using {"question": xxx, "answer": xxx}. Synthesized question-answer pairs should not have any line breaks;

The correct answers to two questions are as follows:

{qa1 }

{qa2 }

The synthesized question-answer pair is:

## D Highest Quality Strategy Details

### D.1 Different Generation Strategy

#### D.1.1 Generation Strategy Definition

**LongMIT+Best-Strategy** means that add all strategies that yield better performance but may incur higher costs. These include using GPT-4o as the backbone model, incorporating additional rationales, merging questions with corresponding documents, and other related techniques. More details are described in Sec. D.2.

**Self-Instruct strategy** involves prompting GPT-4o to autonomously generate questions and corresponding answers based on the provided document, leveraging its self-instruction capabilities for required outputs.

#### D.1.2 Generation Strategy Discussion

As shown in Figure 8, the quality of the data synthesized using Qwen2 within the MIMG framework significantly exceeds that generated by GPT-4 using the Self-Instruct strategy. This illustrates the effectiveness of our framework, even more effective than backbone replacement. Furthermore, as illustrated in Table 1, our approach consistently outperforms prior efforts utilizing GPT-3.5, GPT-4, and even human-created data, which demonstrates and further proves the effectiveness of our framework.

In addition, GPT4 is often impractical for training data synthesis. It is also important to note that the large-scale training data, distilled from only 4 billion tokens using GPT-4, incurs a cost exceeding \$10,000, rendering such an approach impractical for many applications.

### D.2 Implementation Details

To achieve the highest quality data, we deliberately prioritize the use of GPT-4o as the backbone for all processes, fully disregarding cost constraints. This decision is driven by the understanding that ensuring the best data quality is paramount for the success of our project. Furthermore, to maintain and enhance performance during the exploration

	InterLM-2.5-7B-Enhance					InterLM-2.5-7B-Enhance + LongMIT				
	4k	8k	16k	32k	128k	4k	8k	16k	32k	128k
S-NIAH Subtask-1	99.00	99.00	100.00	100.00	16.00	99.00	99.00	99.00	99.00	97.00
S-NIAH Subtask-2	100.00	99.00	100.00	99.00	97.00	100.00	100.00	100.00	100.00	100.00
S-NIAH Subtask-3	99.00	98.00	99.00	99.00	100.00	99.00	99.00	99.00	99.00	100.00
MK-NIAH Subtask-1	97.00	98.00	97.00	88.00	58.00	100.00	100.00	98.00	99.00	90.00
MK-NIAH Subtask-2	99.00	99.00	96.00	81.00	28.00	99.00	100.00	100.00	95.00	63.00
MK-NIAH Subtask-3	95.00	90.00	56.00	14.00	0.00	96.00	91.00	70.00	33.00	2.00
MV-NIAH	99.25	99.50	99.50	94.50	84.50	99.00	99.00	97.00	93.50	89.50
MQ-NIAH	98.00	98.75	97.50	94.00	86.00	100.00	100.00	100.00	99.25	94.25
VT	91.20	91.80	98.60	97.40	0.00	96.60	97.80	98.80	95.60	94.20
FWE	85.33	87.00	84.67	91.00	71.67	86.00	89.00	86.67	90.67	78.33
CWE	83.40	67.10	34.70	26.40	0.10	75.60	40.40	8.50	6.10	0.20
Single Hop QA	90.00	80.00	81.00	75.00	42.00	92.00	84.00	82.00	80.00	58.00
Multi Hop QA	70.00	67.00	64.00	53.00	35.00	73.00	72.00	68.00	63.00	47.00
Average	92.78	90.32	85.23	77.87	47.56	<b>93.48</b>	<b>90.09</b>	<b>85.15</b>	<b>81.01</b>	<b>70.27</b>

Table 4: The evaluation performance on Ruler (Hsieh et al., 2024) benchmark based on LongMIT extended to 128K.




Figure 10: Analysis of the impact of different training dataset sizes on the average accuracy score.




Figure 11: Analysis of the impact of hop on model performance, where 1-hop is the reproduced version of the Quest (Gao et al., 2024) dataset.

phase, we implement a range of strategies aimed at maximizing the data retention rate.

Specifically, for the Quality Verification Agent, we employ a multi-faceted approach that includes more-perspectives scoring mechanisms, the addition of rationales, the integration of multiple perspectives, and the application of detailed guidelines. For the Single-hop Question Generation Agent, we have adopted a question-then-answer strategy. This approach is complemented by the incorporation of

	High-quality	Diversity	Multi-hop
Self-Instruct	61.3	53.4	33.1
MIMG	94.8	88.2	94.8
w/o MQMA	85.0	85.7	45.7
w/o MQS	92.8	65.4	89.4
w/o SQGA	82.4	69.4	73.2
w/o QVA	78.3	88.1	91.0

Table 5: Results of the ablation study on MIMG, reporting the data quality score as the average of three independent manual evaluations.

rationales, which provide context and justification for each query generated. Additionally, we require LLMs to generate only one question per query, which is intended to reduce the logical burden on the model, thereby improving the coherence and relevance of the questions produced. In the case of Multiple Question Sampling, we utilize BGE embeddings for the retrieval of questions. This technique is applied both within individual documents (intra-document) and across multiple documents (inter-document). Finally, for the Multi-hop Question Merging Agent, we employ a strategy that involves merging questions and answers using document references. This method ensures that the merged questions and answers are contextually aligned and coherent. Notably, we have opted to remove the rationale for merging in this process, as we found that it adds unnecessary complexity without significantly improving the quality of the merged content.

### D.3 Ablation Analysis

We analyze the contributions of various agent components to the quality of human-annotated data.

Model	inst_level_loose_acc	inst_level_strict_acc	prompt_level_loose_acc	prompt_level_strict_acc
Llama-3-8B-ProLong	0.1259	0.1199	0.0610	0.0536
+ChatQA2	0.0156	0.0132	0.0074	0.0037
+LongMIT	0.1583	0.1463	0.0795	0.0684

Table 6: The instruction-following capabilities on IFEval for different instruction datasets.

	High-quality	Diversity	Multi-hop
InternLM2-Chat-20B (Cai et al., 2024)			
Self-Instruct	51.0	44.0	52.0
MIMG	86.0	66.0	80.0
Qwen-72B-Instruct (Yang et al., 2024)			
Self-Instruct	61.3	53.4	33.1
MIMG	94.8	88.2	94.8
GPT-4o (Achiam et al., 2023)			
Self-Instruct	86.0	81.0	66.0

Table 7: The human-annotated quality score for different LLMs and strategies.

As summarized in Table 5, each component positively impacts the overall quality. The Multi-hop Question Merger Agent notably improves the multi-hop quality, while the Multiple Question Sampling mechanism increases data diversity. The Single-hop Question Generation Agent is crucial for enhancing both quality and diversity. Lastly, the Quality Verification Agent acts as a safeguard, ensuring a lower bound of model quality and further improving data integrity.

#### D.4 More Backbone Exploration

To assess the data generation quality across additional backbone architectures, we manually labeled 100 samples for both open-source and closed-source models. As reported in Table 7, integrating MIMG with InternLM2 yields a marked improvement in output quality. Moreover, LongMIT surpasses the combination of GPT-4o and Self-Instruct on our quality metrics. We will elaborate on these findings in the next version.

#### D.5 Instruction Capabilities

As MIMG is designed to enhance models’ ability to follow instructions over extended contexts. Accordingly, it is essential both to improve instruction-following performance and to preserve this ability in long-context settings. Because there is currently no benchmark dedicated to long-context instruction following, we evaluated instruction adherence using two established short-context benchmarks: IFEval and ArenaHard.

ArenaHard	Score	95% Confidence Interval
Llama-3-8B-ProLong	7.2	(−1.2, 1.2)
+ChatQA2	4.4	(−0.9, 0.8)
+LongMIT	6.7	(−1.1, 1.0)

Table 8: The instruction-following capabilities on ArenaHard for different instruction datasets.

As shown in Table 6, remarkably, the model trained with LongMIT not only achieved substantial gains in handling long contexts but also outperformed the base model on the IFEval benchmark. In contrast, training on ChatQA2 led to a pronounced decline in instruction-following performance on IFEval.

In addition, as presented in Table 8, on the more demanding short-context instruction-following benchmark ArenaHard, our method continues to outperform ChatQA2 by a substantial margin. Moreover, the improvements to long-context processing do not compromise its effectiveness on challenging instruction-following tasks.

## E Instruction Tuning Experiments Details

### E.1 Training Details

All models were trained using 64 A800\*80G GPUs with the DeepSpeed+ZeRO-1 framework. The maximum sequence length was set from 4K to 128K, with any sequences exceeding this length truncated from the right. The training process utilized the Adam optimizer with a learning rate of  $3 \times 10^{-5}$ ,  $\beta_1 = 0.9$ , and  $\beta_2 = 0.95$ .

To enhance training efficiency, we employed a packing strategy that concatenates training samples to reach the maximum sequence length. Additionally, Flash Attention (Dao et al., 2022; Dao, 2024) is used to accelerate the computation of the attention mechanism. The global batch size consisted of 4 million tokens, and the entire dataset is trained over one epoch.

### E.2 Evaluation Details

Based on the methodology proposed by Bai et al. (2024a), evaluating Token F1 using a model optimized through Chain of Thought (CoT) (Wei et al.,

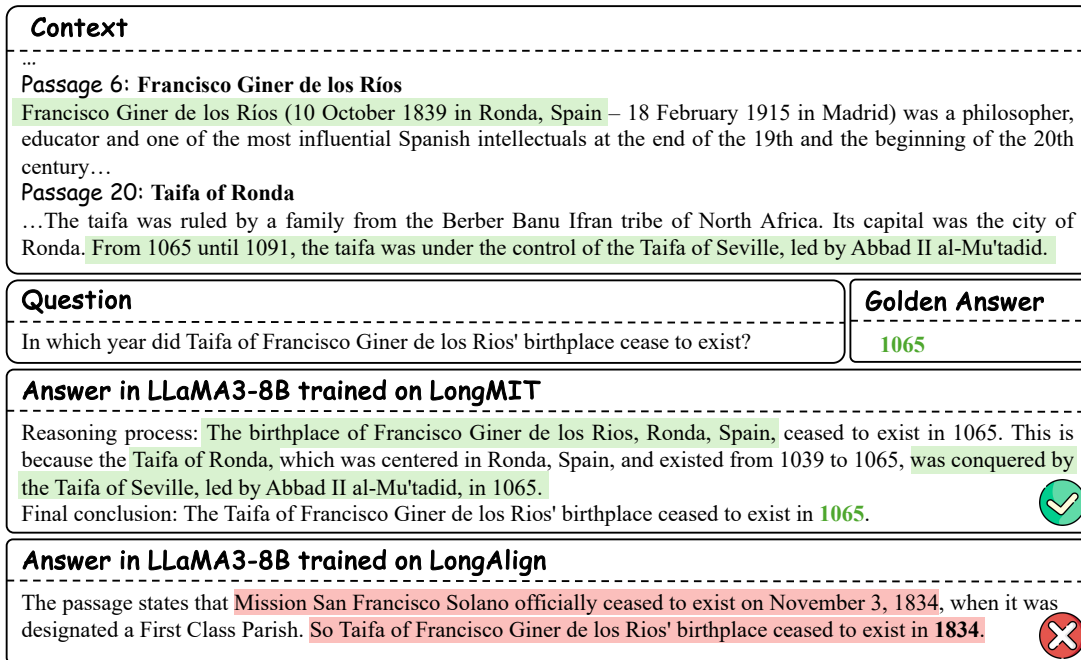


Figure 12: The case study of the effectiveness of LongMIT.

2022) reasoning proves to be challenging (Chen et al., 2025a). To address this limitation, we employ GPT-4 as a consistency evaluator. Our testing demonstrates that the error rate of GPT-4 in this role remains consistently low, with deviations falling within a 2% margin. The corresponding prompt used is outlined below:

```
Suppose you are a professional annotator.
Given the result predicted by a model, you
need to annotate whether the "[[PREDIC-
TION]]" is consistent with the given "[[REF-
ERENCE]]" based on the "[[QUESTION]]".
[[QUESTION]]
{question}
[[PREDICTION]]
{predictions}
[[REFERENCE]]
{answer}
Finally, you should give me an annotation in
the format:
"" { "short_pred_answer": "xxx", "pre-
dict_consistency": BOOL } ""
```

## F Case study

To gain a more nuanced and intuitive qualitative understanding of our model’s performance, we con-

ducted a detailed case study, resulting in two significant findings:

- **Impact of Instruction Quality:** As illustrated in Figure 13, models trained with high-quality multi-hop instruction data, specifically the LongMIT dataset, exhibit enhanced logical reasoning capabilities. These models are better equipped to process and analyze extensive textual information, enabling them to derive more accurate and reliable reasoning. In contrast, models trained using traditional, lower-quality instruction data, such as LongAlign (Bai et al., 2024a), demonstrate a reduced capacity for logical reasoning. This comparison underscores the importance of the quality of training data in developing models that can effectively handle complex reasoning tasks, especially when dealing with long and intricate texts.
- **Role of Rationale Incorporation in Training:** Furthermore, as depicted in Figure 12, our analysis reveals that the inclusion of additional rationales during the training process significantly enhances the model’s ability to focus on relevant information within long texts and make precise inferences. This finding is particularly evident when comparing models that underwent Chain-of-Thought (CoT) (Wei et al., 2022) training with those that did not. Specifically, models that lacked CoT training tend to falter during inference, often generating erroneous outputs, such



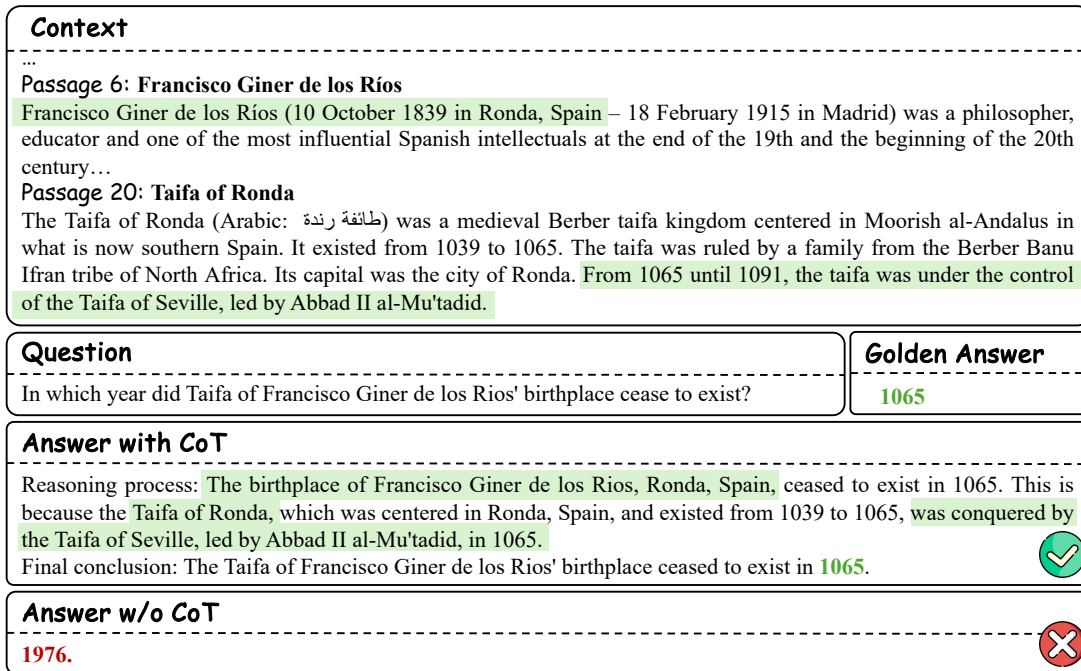


Figure 13: The case study of whether utilize reasoning process for instruction tuning.

as the completely incorrect answer "1976". On the other hand, models that were fine-tuned with CoT training not only demonstrate a coherent logical reasoning process but also consistently arrive at the correct answer, "1065". This result highlights the critical role of rationale-based training in improving the model's reasoning accuracy and its ability to tackle complex inferential challenges.

## G Discussion about Long Code Data Generation

In both coding and mathematical settings, it is essential to perform multi-step reasoning across source files and scholarly articles. To investigate this requirement, we evaluate our framework on a set of coding tasks. The resulting code-question-answer dataset exhibited a high level of quality, with approximately 80% of entries rated as satisfactory. Furthermore, by consolidating related questions, we are capable to prompt the advanced LLM (GPT-4o) to generate code-generation queries of competition caliber.

Below, we illustrate a simple example in which the context consists of code documents from the "torch.nn" module. First, the Single-Hop Question Generation Agent generates single-hop question-answer pairs as follows:

**Question 1:** What is the purpose of nn.LeakyReLU?

**Answer 1:** nn.LeakyReLU is similar to nn.ReLU but allows a small, non-zero slope for negative values, controlled by a "negative\_slope" parameter, making it less likely for neurons to become inactive.

**Question 2:** What is the nn.ReLU?

**Answer 2:** nn.ReLU (Rectified Linear Unit) applies the rectified linear unit function element-wise, where any negative input values are set to zero, leaving positive values unchanged.

After that, Multi-hop Question Merging Agent generates the merged question-answering pair as:

**Question:** How does the behavior of nn.ReLU differ from nn.LeakyReLU in terms of their activation functions when used in a neural network constructed with nn.Sequential?

**Answer:** In a neural network constructed with nn.Sequential, nn.ReLU applies the rectified linear unit function element-wise, where negative input values are set to zero, while positive values remain unchanged.

In contrast, `nn.LeakyReLU` behaves similarly but allows a small, non-zero slope for negative values through a `'negative_slope'` parameter, which helps prevent neurons from becoming inactive. This means `nn.LeakyReLU` can potentially maintain more active neurons compared to `nn.ReLU` in sequential layers.

Further, We try to propose a new prompt for Multi-hop Question Merging Agent to generate more difficult questions for competitive code based on merged question.

Using the following question and answer related knowledge and meanings as a reference, generate additional programming challenges and their solutions specifically designed for high-level coding competitions, such as Olympiad-style problems or advanced AI-related tasks (like those on Kaggle). The focus should be on tasks that require efficient code generation, algorithm design, or AI model implementation.

As shown below, large language models (LLMs) can efficiently produce sophisticated code-generation datasets. This capability represents a compelling and promising research direction. One of generated questions is presented as follows:

#### **Problem Statement:**

You are tasked with using a genetic algorithm (GA) to optimize the hyperparameters of a neural network for a given task (e.g., image classification using CIFAR-10 dataset). You need to optimize parameters such as the number of layers, the number of neurons in each layer, and the activation function used at each layer.

#### **Input:**

A dataset such as CIFAR-10 or MNIST (for image classification). An initial population of neural network configurations represented by a list of hyperparameters:

- Number of layers.
- Number of neurons per layer.

- Activation functions (ReLU, LeakyReLU, Sigmoid, etc.).

Further, the corresponding generated answer is shown below:

#### **Output:**

The optimized neural network configuration after a set number of generations.

#### **Solution:**

```
import random
from sklearn.datasets import
    fetch_openml
from sklearn.model_selection import
    train_test_split
import torch
import torch.nn as nn
import torch.optim as optim
import torch.nn.functional as F

# Simple feedforward neural network
class NeuralNet(nn.Module):
    def __init__(self, layers,
                 activation_fn):
        super(NeuralNet, self).__init__()
        self.layers = nn.ModuleList()
        for i in range(len(layers) - 1):
            self.layers.append(nn.Linear(
                layers[i], layers[i+1])
            )
        self.activation_fn =
            activation_fn

    def forward(self, x):
        for layer in self.layers:
            x = layer(x)
            x = self.activation_fn(x)
        return x

# Genetic algorithm for hyperparameter
# optimization

def genetic_algorithm(dataset,
                    generations=10, population_size=10):
    population = []
    for _ in range(population_size):
        layers = [random.randint(5, 100)
                 for _ in range(random.
                                randint(3, 6))] # Number of
            neurons in layers
        activation_fn = random.choice([
            nn.ReLU(), nn.LeakyReLU(),
            nn.Sigmoid()])
        population.append((layers,
                           activation_fn))

    best_score = float('inf')
    best_model = None

    for generation in range(generations):
        :
```

```

for layers, activation_fn in
population:
    # Convert layers to torch
    model

    model = NeuralNet(layers,
        activation_fn)
    optimizer = optim.Adam(model
        .parameters())
    criterion = nn.
        CrossEntropyLoss()

    # Training loop on dataset
    model.train()
    optimizer.zero_grad()
    # Assume a dummy data
    loading and training
    process here
    loss = criterion(model(
        dummy_inputs),
        dummy_labels)
    loss.backward()
    optimizer.step()

    if loss.item() < best_score:
        best_score = loss.item()
        best_model = model

    # Selection for next generation
    (randomly choose top models
    for simplicity)
    population = random.sample(
        population, k=
        population_size)

return best_model

# Example usage
best_model = genetic_algorithm(
    fetch_openml("CIFAR_10"))

```