

# LADM: Long-context Training Data Selection with Attention-based Dependency Measurement for LLMs

Jianghao Chen<sup>1,2,3</sup>, Junhong Wu<sup>1,2</sup>, Yangyifan Xu<sup>1,2</sup>, Jiajun Zhang<sup>1,2,4\*</sup>

<sup>1</sup>Institute of Automation, Chinese Academy of Sciences

<sup>2</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup>Zhongguancun Academy, Beijing, China <sup>4</sup>Wuhan AI Research

{chenjianghao2022, wujunhong2021, xuyangyifan2021}@ia.ac.cn  
jjzhang@nlpr.ia.ac.cn

## Abstract

Long-context modeling has drawn more and more attention in the area of Large Language Models (LLMs). Continual training with long-context data becomes the de-facto method to equip LLMs with the ability to process long inputs. However, it still remains an open challenge to measure the quality of long-context training data. To address this issue, we propose a Long-context data selection framework with Attention-based Dependency Measurement (LADM), which can efficiently identify high-quality long-context data from a large-scale, multi-domain pre-training corpus. LADM leverages the retrieval capabilities of the attention mechanism to capture contextual dependencies, ensuring a comprehensive quality measurement of long-context data. Experimental results show that our LADM framework significantly boosts the performance of LLMs on multiple long-context tasks with only 1B tokens for continual training.<sup>1</sup>

## 1 Introduction

Long-context modeling for Large Language Models (LLMs) has recently drawn more and more attention. The maximum context window of current LLMs has been significantly extended to 128K tokens for GPT-4 (Achiam et al., 2023) and 1M tokens for Gemini 1.5 (Reid et al., 2024). The long-context modeling ability greatly contributes to more sophisticated applications of LLMs in various fields such as long-context retrieval (Kamradt, 2023; Xu et al., 2023; An et al., 2024b), question answering (Kočíšký et al., 2018; Dasigi et al., 2021; Bai et al., 2023), and summarization (Huang et al., 2021; Zhong et al., 2021; Fabbri et al., 2019).

Preparing long-context dataset and performing continual training has become the de-facto framework to enrich LLMs with the ability to process

long inputs. However, quality measurement of the long-context training data has not received enough attention. Previous studies show that if the training data is composed of concatenated short samples or lacking dependencies over long contexts (Ding et al., 2024; Fu et al., 2024; Chen et al., 2024a), models may fail to learn how to handle long-range and diverse contextual dependencies. These low-quality training samples can aggravate the tendency of LLMs to ignore valuable distant contextual information, further limiting their performance on long-context tasks. Therefore, measuring the quality of long-context training data and mining high-quality training samples become crucial for enhancing the long-context modeling capability of LLMs.

The dependency between segments within context is the key indicator of high-quality long-context data. Recent studies have proposed several methods to take long-context dependencies into consideration when constructing pre-training data. Staniszewski et al. (2023) enhance long-range semantic dependencies by integrating relevant documents into one sample. Chen et al. (2024a) split long-context data into segments and measure dependencies through delta perplexity scores between individual segments. However, these methods ignore the inherent structures and relationships within long contexts, leading to inaccurate measurement of long-range contextual dependencies.

To address the above issue, we propose a Long-context data selection framework with Attention-based Dependency Measurement (LADM), which measures the long-context dependency with span-level attention scores. Inspired by the inherent retrieval operations of the attention mechanism (Mittal et al., 2022; Wu et al., 2024), LADM aims to leverage the attention distribution to measure the relationship within long contexts. Specifically, we first train a tiny model with long-context modeling capability named Long Attention Calculator. Then, to measure the dependency of a single long-context

\* Corresponding Author

<sup>1</sup>Our code is available at <https://github.com/ZNLP/LADM>

sample, we feed the sample into the Long Attention Calculator and compute the **Pairwise Focus Score (PFS)** between different spans by the accumulated attention scores. Subsequently, the **Aggregated Focus Score (AFS)** for each span is derived by incorporating all PFS between this span and its preceding ones. Finally, we define the sample-level **Contextual Dependency Score (CDS)** by a weighted sum of all AFS and select samples with high CDS for continual pre-training. Experimental results show that LADM outperforms other data selection methods on various long-context tasks, achieving an average performance improvement of 2.16% for four models across different sizes and architectures on the LongBench dataset.

Our contributions are summarized as follows: 1) We propose an efficient data selection framework LADM, which can identify high-quality long-context data with long-range and diverse contextual dependencies from a large-scale pre-training corpus. 2) We introduce a novel method for dependency measurement through the attention mechanism, effectively capturing long-range and diverse relationships within the complete contextual information. 3) Experimental results demonstrate the superiority of our method. We achieve better performance with only half of the pre-training tokens than the random sampling approach.

## 2 Related Work

### 2.1 Long-context Modeling for LLMs

Enhanced long-context modeling in LLMs can drive substantial progress in artificial intelligence across various domains, such as long-chain reasoning (Chen et al., 2025b; Yeo et al., 2025; Sun et al., 2025a,b), long video and image processing (Weng et al., 2024; Zhang et al., 2024; Guan et al., 2025; Jian et al., 2024), and long-form generation (Chen et al., 2025a; Bai et al., 2024; Pham et al., 2024). To enable better long-context processing capability of LLMs, recent studies have explored both training-free and training-augmented methods. For training-free methods, Xiao et al. (2024) and Han et al. (2024) focus on retaining the attention on the initial and local tokens while masking those at greater distances, thereby enhancing the length generalization ability of LLMs. DCA (An et al., 2024a) and SelfExtend (Jin et al., 2024) rearrange position indices of long-context inputs and get impressive length extrapolation capability without fine-tuning. Training-augmented methods involve continuing

to pre-train LLMs on longer contexts with modified positional encoding. Positional Interpolation (PI) (Chen et al., 2023), NTK (bloc97, 2023), and YaRN (Peng et al., 2024) effectively achieve context window extension through the interpolation and extrapolation of RoPE positional embedding (Su et al., 2024). Moreover, LongLoRA (Chen et al., 2024b) combines PI with  $S^2$ -Attn and LoRA (Hu et al., 2021), enabling more efficient training. Zhu et al. (2024) ensure a collinear constraint between query and key vectors when integrating RoPE into self-attention and pre-train LLMs with better long-context extrapolation ability. To fully leverage these methods, it is essential to construct high-quality long-context training data.

### 2.2 Long-context Training Data for LLMs

Training data quality is crucial for the long-context modeling capability of LLMs. Fu et al. (2024) focuses on domain balance and length up-sampling for long-context training data. Staniszewski et al. (2023) and Gao et al. (2024a) propose similarity-based approaches of grouping documents to construct long-context training data. He et al. (2024) designs synthesized multi-doc QA tasks and improves the long-context information searching and reflection ability. To alleviate the "lost-in-the-middle" (Liu et al., 2024) phenomenon, An et al. (2024b) introduces Information-Intensive training on long-context QA tasks and Xiong et al. (2024) designs key-value retrieval tasks for fine-tuning.

Our approach characterizes the contextual dependencies within long-context samples based on the attention mechanism and effectively identifies high-quality long-context data from a large pre-training corpus. Concurrently with our work, Chen et al. (2024a) proposes a framework ProLong that divides a long-context sample into segments and calculates the delta perplexity scores between individual segments as Dependency Strength without the original context. However, we capture dependencies within the complete context, providing a more comprehensive view of the contextual relationships. Moreover, our experiments are conducted on a more diverse pre-training corpus, proving the robustness and applicability of our method.

## 3 Method

### 3.1 Problem Formulation

The long-context modeling capability of LLMs is obtained by continual pre-training on long-context

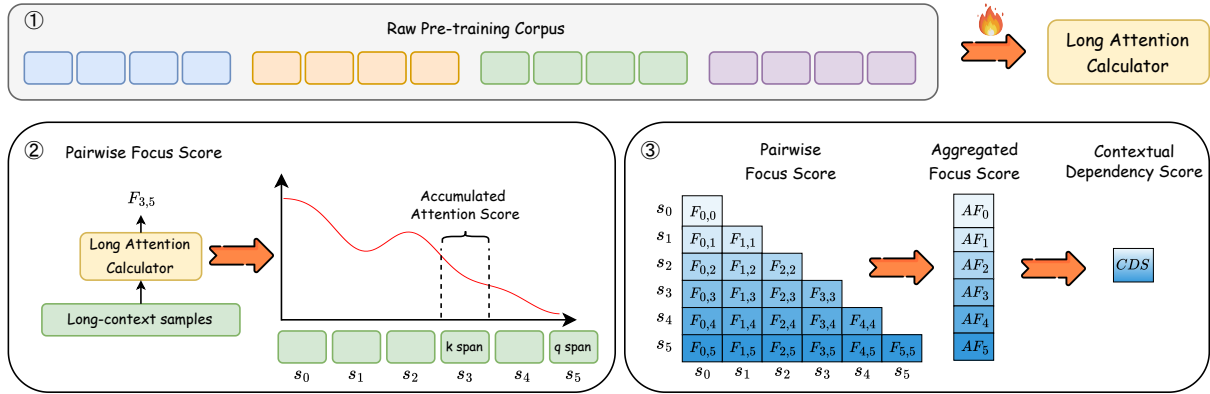


Figure 1: The overall framework of LADM. We first train a Long Attention Calculator, then calculate the Pairwise Focus Score (PFS) to measure the dependency between spans. Then, we compute the Aggregated Focus Score (AFS) of each span and merge them as the Contextual Dependency Score (CDS) of a single long-context sample.

samples. Given a dataset  $\mathcal{D}$  of long-context samples with different quality levels, we aim to select a high-quality subset  $\mathcal{D}_s$ . Specifically, we define a scoring function  $s$  to measure the quality of each data sample  $x \in \mathcal{D}$ . Then we rank all the samples according to their scores  $s(x)$  and select the top  $N$  samples to construct the subset  $\mathcal{D}_s$  as follows:

$$\mathcal{D}_s = \{x \in \mathcal{D} \mid \text{rank}(s(x)) \leq N\} \quad (1)$$

### 3.2 Preliminary Analysis

Data Length	Evaluation Context length								Avg
	4K	8K	12K	16K	20K	24K	28K	32K	
4K	1.00	0.80	0.90	0.76	0.22	0.40	0.20	0.30	0.57
8K	1.00	0.60	0.80	1.00	0.53	0.66	0.50	0.43	0.69
16K	1.00	1.00	0.60	0.90	0.70	1.00	0.90	0.70	0.85
32K	1.00	0.90	0.80	0.60	0.90	0.90	0.90	1.00	0.88

Table 1: The "Needle-in-the-Haystack" performance on models trained with data concatenated at different lengths. For each evaluation context length, the result is the average of all performances across needle insertion depths, ranging from 0 to 1.

We first conduct a preliminary experiment to analyze the impact of contextual dependencies on LLMs' long-context modeling capability. We train the Llama2-7B model with 32K-token sequences concatenated by  $4K \times 8$ ,  $8K \times 4$ , and  $16K \times 2$ -token samples, respectively, and compare them with one trained with samples exceeding 32K. The total number of training tokens amounts to 0.5B.

As shown in Table 1, we can observe that the average retrieval accuracy decreases as the original length of training data becomes shorter. This trend indicates that models trained with concatenated and contextually unrelated data may focus more on local information and lack retrieval ability across long contexts. Consequently, training data

with strong contextual dependencies is essential for enhancing the long-context modeling capability of LLMs. These findings align with recent studies (Staniszewski et al., 2023; Chen et al., 2024a).

Building on this analysis, we further explore how to measure the contextual dependencies within long-context data. Wu et al. (2024) reveal that LLMs incorporate retrieval ability within the attention mechanism, which can be reflected in the attention distribution, with higher weights assigned to previous tokens more related to the current token. Inspired by this, we propose LADM framework to quantify the contextual dependencies by analyzing the attention distribution over long contexts.

### 3.3 LADM Framework

Figure 1 illustrates the overall framework of LADM. We first train a Long Attention Calculator with basic long-context modeling capability and use it to calculate the Pairwise Focus Score (PFS), which measures the dependency between spans within long contexts. Then, the Aggregated Focus Score (AFS) for each span is derived by incorporating all PFS between this span and previous spans. Finally, the sample-level Contextual Dependency Score (CDS) is computed by merging all the AFS of different spans.

**Long Attention Calculator** To leverage the intrinsic retrieval capability of the attention mechanism for detecting dependency within long contexts, we present Long Attention Calculator, a compact model with basic long-context modeling capability. Specifically, we choose the TinyLlama-1.1B-v1.1<sup>2</sup> for subsequent data filtering efficiency

<sup>2</sup>[https://huggingface.co/TinyLlama/TinyLlama\\_1.1](https://huggingface.co/TinyLlama/TinyLlama_1.1)




Figure 2: The median attention scores under different 32K data sample construction methods.

and train it with randomly sampled 32K-token sequences. To demonstrate the Long Attention Calculator’s ability to capture dependencies within long contexts, we use it to calculate the median value of accumulated attention scores between long-range spans on 1,000 32K samples concatenated by data with different lengths. As shown in Figure 2, the results indicate that the Long Attention Calculator can distinguish samples with varied dependencies via attention scores over long contexts, as the complete 32K samples exhibit higher attention scores on distant previous tokens than concatenated ones. Therefore, we can design metrics based on attention scores to measure the dependency within long-context data.

**Pairwise Focus Score** Given a long-context sample  $S$  consisting of  $N$  spans of length  $l$ :  $S = \{s_0, s_1, \dots, s_{N-1}\}$ , the Pairwise Focus Score  $\text{PFS}(i, j)$  between  $s_i$  and  $s_j$  (where  $j > i$ ) is defined as follows:

$$\text{PFS}(i, j) = \text{Sum} \left( \text{Softmax} \left( \frac{Q_j K_{0:j}^T}{\sqrt{d_k}} \right)[:, i] \right) \quad (2)$$

where  $Q_j$  is the query states of  $s_j$ ,  $K_{0:j}$  is the key states of  $s_0$  to  $s_j$  and  $\frac{1}{\sqrt{d_k}}$  is the scaling factor.  $\text{PFS}(i, j)$  calculates the accumulated attention weights that  $s_j$  assigns to  $s_i$ , quantifying the influence of  $s_i$  on the final representation of  $s_j$  in the attention mechanism. Therefore, we can effectively detect spans that exhibit strong dependencies within long contexts by calculating PFS with the Long Attention Calculator.

To comprehensively evaluate the contextual dependencies of a long-context sample, it is essential to aggregate PFS between different spans. This approach can help us gain a deeper understanding of the complex relationships between various parts of the long-context data. We consider the following criteria for aggregating multiple PFS:

**Aggregated Focus Score** For each span  $s_j$ , we calculate PFS at varied intervals, including  $\text{PFS}(0, j)$ ,  $\text{PFS}(1, j)$ , ...,  $\text{PFS}(j-1, j)$ . This aids in understanding how the influence of previous spans varies with increasing intervals, revealing both short-range and long-range dependencies. Notably, recent studies (Xiao et al., 2024; Hsieh et al., 2024) have shown that initial and recent tokens often receive disproportionate attention weights, suggesting a predominance of the beginning and local dependencies. Therefore, we exclude scores for the first  $m$  spans ( $s_0, \dots, s_{m-1}$ ) and the local  $n$  spans ( $s_{j-n}, \dots, s_{j-1}$ ). We select previous spans at a stride of  $d$  for further computational efficiency.

We apply weights to these scores based on the lengths of the intervals, thus encouraging longer-distance dependencies. Additionally, we incorporate the variance of these scores. A higher variance indicates that  $s_j$  exhibits more diverse dependencies from its previous context, reflecting a complex dependency pattern and rich structural information within the long-context sample. Therefore, we define the Aggregated Focus Score (AFS) of span  $s_j$  as follows (omitting stride sampling for clarity):

$$\sigma_j = \sigma(\text{PFS}(m, j), \dots, \text{PFS}(j-n-1, j)) \quad (3)$$

$$\text{AFS}(j) = \sigma_j \sum_{i=m}^{j-n-1} \frac{j-i}{N} \cdot \text{PFS}(i, j) \quad (4)$$

where  $\sigma_j$  is the standard deviation of all PFS.

**Contextual Dependency Score** While  $\text{AFS}(j)$  provides a measurement of the dependencies for a single span  $s_j$ , we should further consider the contributions of all spans to accurately represent the overall dependencies of a long-context sample. To achieve this, we sum all AFS and apply a weight based on index  $j$  to highlight the contributions of spans with larger positions. This approach aligns with our focus on long-range dependencies, as spans with larger indices have the potential to consider previous spans at greater distances. The sample-level Contextual Dependency Score (CDS) is therefore defined as follows:

$$\text{CDS}(S) = \sum_{j=n_0}^{N-1} \frac{j}{N} \cdot \text{AFS}(j) \quad (5)$$

where  $n_0$  is the index of the first span and  $d$  is the stride. We exclude the initial spans’  $\text{AFS}(j < n_0)$ , as these early spans have fewer previous spans to



depend on, resulting in a less informative measurement of dependencies. For computational efficiency, we calculate AFS at a stride of  $d$  (omitted from Eq.5 for clarity).

## 4 Experiments Settings

### 4.1 Pre-training Dataset

We use the Pile (Gao et al., 2020) corpus for long-context continual pre-training. Our experiments are conducted on data samples with 32K tokens due to the scarcity of longer samples. We remove samples with lengths less than 32K measured by LlamaTokenizer. The detailed information of our pre-training dataset is provided in Appendix A.

### 4.2 Baselines

We compare LADM with the following methods:

**Random Sampling** We randomly sample long-context data from the dataset for continual training.

**ProLong** Chen et al. (2024a) propose a framework filtering long-context data with delta perplexity scores between individual segments. We follow all settings of Prolong except changing the model to TinyLlama-1.1B-v1.1 for a fair comparison.

### 4.3 Implementation Details

**Hyper-parameters** When calculating PFS, we truncate all samples to 32K-token sequences for batch operation. The span size is set to  $l = 128$ , resulting in  $N = 256$  spans per sample. Span-level AFS uses  $m = 1, n = d = 4$ , excluding the first and recent four spans and selecting previous spans at a stride of four. Sample-level CDS uses  $n_0 = 16, d = 4$ , excluding the initial 16 spans' AFS and calculating AFS at a stride of four.

**Data Selection** We rank all samples with their CDS and select the top-ranking samples from each data domain, maintaining the original domain distribution. All methods use 1B tokens selected from the dataset for continual training.

**Training Configuration** We increase the base of RoPE from 10,000 to 500,000 following (Xiong et al., 2023). For the Long Attention Calculator, since the original context length of TinyLlama is 2K, we use 5B randomly sampled training tokens to ensure better long-context dependencies measurement. We also provide experimental results using the Long Attention Calculator trained with 1B

tokens in Appendix E for comparison. We continually pre-train OpenLlama-3B-v2 (Geng and Liu, 2023), Llama2-7B/13B (Touvron et al., 2023) and Mistral-7B-v0.1 (Jiang et al., 2023) with 32K-token sequences for 1B tokens. More training details are displayed in Appendix B.

### 4.4 Evaluation Tasks and datasets

We take the following tasks to evaluate the long-context capability of LLMs:

**Perplexity Evaluation** We evaluate the language modeling capability of LLMs by measuring the perplexity (PPL) on real-world long-context data. We collect samples exceeding 32K from the test split of the Proof-Pile (Azerbaiyev et al., 2022) dataset and calculate the average PPL across different context window sizes.

**Synthetic Tasks Evaluation** We test the long-context retrieval ability of LLMs on the "Needle-in-the-Haystack" task (Kamradt, 2023). This synthetic task is designed to evaluate LLMs' capability to locate essential information across varying positions and context lengths.

**Real-World Tasks Evaluation** Evaluation on perplexity or synthetic retrieval tasks can not truly reflect the performance of LLMs under real-world scenarios (Hu et al., 2024; Fang et al., 2024). Therefore, we select different types of tasks from Long-Bench (Bai et al., 2023) for further evaluation.

## 5 Experimental Results

We evaluate LADM with OpenLlama-3B-v2, Llama2-7B/13B and Mistral-7B-v0.1 on language modeling tasks, synthetic long-context tasks and real-world long-context tasks. We use "OL-3B", "L-7B", "L-13B", and "M-7B" to denote these models for short.

### 5.1 Perplexity Evaluation

Table 2 shows the perplexity of long-context samples from Proof-Pile across various context window sizes. The perplexity differences among the three methods are minimal due to the same domain distribution of the training data. Our proposed LADM outperforms other methods under all model and context window size settings. This demonstrates the superiority of our models on long-context language modeling tasks.




Figure 3: The "Needle-in-the-Haystack" task performance of different data selection methods. The x-axis denotes the evaluation context length, and the y-axis denotes insertion depths of the "needle".

## 5.2 Synthetic Tasks Evaluation

In Figure 3, we compare the retrieval accuracy of our models with other baseline models on the "Needle-in-the-Haystack" task. Notably, our models show higher average retrieval accuracy and even achieve nearly 100% retrieval rate for Llama2-7B/13B and Mistral-7B with only 1B training tokens. Compared to our method, other methods exhibit inferior performance, particularly in scenarios involving greater retrieval distances or when the "needle" is located in the middle of the context.

## 5.3 Real-World Tasks Evaluation

We report the experimental results on real-world long-context tasks in Table 3. For the randomly sampling method, we conduct three experiments with different random sampling seeds and calculate the average performance. Our LADM framework outperforms other methods across nearly all types of tasks and all models. Specifically, our method achieves an average performance improvement of 2.16% on four models over ProLong. Notably, for Mistral-7B, LADM achieves a substantial improvement of 10.09% on single-document QA and 4.66% on multi-document QA. These results highlight the strong capability of our models in handling real-world tasks and validate the effectiveness of our data selection framework, indicating the importance of incorporating long-range contextual dependencies into training data. Experimental results in Appendix F also show that our method enhances LLMs' long-context capabilities while preserving

Model	Method	Context Window Size				
		2K	4K	8K	16K	32K
OL-3B	Random	4.951	4.268	3.562	3.026	2.685
	ProLong	4.941	4.271	3.573	3.040	2.701
	LADM	<b>4.910</b>	<b>4.247</b>	<b>3.553</b>	<b>3.022</b>	<b>2.683</b>
L-7B	Random	4.515	3.900	3.264	2.780	2.458
	ProLong	4.516	3.906	3.275	2.792	2.470
	LADM	<b>4.481</b>	<b>3.878</b>	<b>3.252</b>	<b>2.773</b>	<b>2.453</b>
M-7B	Random	4.620	3.936	3.267	2.775	2.455
	ProLong	4.293	3.696	3.095	2.644	2.346
	LADM	<b>4.266</b>	<b>3.673</b>	<b>3.076</b>	<b>2.629</b>	<b>2.332</b>
L-13B	Random	4.200	3.657	3.084	2.645	2.349
	ProLong	4.200	3.659	3.090	2.651	2.356
	LADM	<b>4.176</b>	<b>3.637</b>	<b>3.070</b>	<b>2.636</b>	<b>2.342</b>

Table 2: Perplexity of long-context samples from the test split of the Proof-Pile dataset. Each sample is truncated to the corresponding context window size.

short-context performance.

## 6 Analysis

### 6.1 Training Efficiency of LADM

We compare the performance of our method with baseline methods under different training budgets. As shown in Table 4, our method surpasses the baseline even with half of the training tokens. Specifically, LADM achieves slight improvements over OpenLlama-3B and Llama2-13B, exceeding the baseline by 1.31% for Llama2-7B and 2.27% for Mistral-7B. Even when the training tokens amount to 3B and 4B for random sampling, our method can still demonstrate superior performance, as il-

Model	Method	Single-Document QA				Multi-Document QA			
		NarrativeQA	Qasper	MultiFieldQA	AVG	HotpotQA	2WikiMQA	MuSiQue	AVG
GPT-3.5-Turbo	-	23.6	43.3	52.3	39.8	51.6	37.7	26.9	38.7
OL-3B	Random	<b>18.65</b>	24.07	31.44	24.72	26.80	21.00	<b>11.80</b>	19.87
	ProLong	17.04	25.97	32.62	25.21	27.30	23.30	10.50	20.37
	LADM	17.09	<b>28.84</b>	<b>33.97</b>	<b>26.63</b>	<b>29.01</b>	<b>23.71</b>	9.15	<b>20.62</b>
L-7B	Random	24.06	<b>33.27</b>	33.05	30.13	35.29	26.79	15.31	25.80
	ProLong	25.00	28.23	35.27	29.50	40.30	28.91	17.91	29.04
	LADM	<b>26.34</b>	32.28	<b>38.11</b>	<b>32.24</b>	<b>43.42</b>	<b>31.85</b>	<b>18.03</b>	<b>31.10</b>
M-7B	Random	17.30	21.39	33.56	24.08	30.67	<b>28.85</b>	14.37	24.63
	ProLong	14.32	26.41	30.55	23.76	34.29	23.50	15.51	24.43
	LADM	<b>24.05</b>	<b>34.77</b>	<b>42.72</b>	<b>33.85</b>	<b>39.43</b>	28.81	<b>19.04</b>	<b>29.09</b>
L-13B	Random	24.07	31.31	29.32	28.23	44.88	32.36	18.82	32.02
	ProLong	<b>25.06</b>	31.42	<b>30.34</b>	28.94	44.37	34.20	21.02	33.20
	LADM	24.98	<b>33.79</b>	28.33	<b>29.03</b>	<b>47.21</b>	<b>36.29</b>	<b>23.09</b>	<b>35.53</b>

Model	Method	Summarization				Code			Overall
		GovReport	QMSum	MultiNews	AVG	LCC	RepoBench-P	AVG	
GPT-3.5-Turbo	-	29.5	23.4	26.7	26.5	54.7	53.6	54.1	39.8
OL-3B	Random	24.33	13.83	14.01	17.39	61.21	48.69	54.95	29.23
	ProLong	<b>25.06</b>	12.89	8.13	15.36	58.90	48.44	53.67	28.65
	LADM	24.04	<b>15.00</b>	<b>15.91</b>	<b>18.32</b>	<b>61.33</b>	<b>49.78</b>	<b>55.56</b>	<b>30.28</b>
L-7B	Random	29.56	21.26	23.66	24.83	59.32	56.51	57.92	34.67
	ProLong	29.54	20.49	23.35	24.46	62.78	55.30	59.04	35.51
	LADM	<b>29.77</b>	<b>21.75</b>	<b>26.02</b>	<b>25.85</b>	<b>65.78</b>	<b>58.46</b>	<b>62.12</b>	<b>37.83</b>
M-7B	Random	24.68	19.93	22.91	22.51	62.95	58.74	60.85	33.02
	ProLong	24.43	18.92	<b>25.01</b>	22.73	65.12	<b>58.51</b>	<b>61.82</b>	33.18
	LADM	<b>28.38</b>	<b>20.64</b>	24.40	<b>24.47</b>	<b>65.41</b>	55.74	60.58	<b>37.00</b>
L-13B	Random	<b>28.07</b>	22.14	26.80	25.67	67.65	61.85	64.75	37.67
	ProLong	27.31	<b>22.97</b>	26.59	25.62	67.73	61.10	64.42	38.04
	LADM	27.43	22.24	<b>27.80</b>	<b>25.82</b>	<b>67.77</b>	<b>62.79</b>	<b>65.28</b>	<b>38.92</b>

Table 3: Performance of models trained with different data selection methods on single-document QA, multi-document QA, summarization and code completion from the LongBench dataset.

Model	Method	Tokens	SD-QA	MD-QA	SUM	CODE	AVG
OL-3B	Random	1B	24.72	19.87	17.39	54.95	29.23
		2B	25.16	20.59	19.11	55.14	30.00
	ProLong	1B	25.21	20.37	15.36	53.67	28.65
		2B	<b>27.38</b>	18.32	<b>19.20</b>	55.26	30.04
	LADM	1B	26.63	<b>20.62</b>	18.32	<b>55.56</b>	<b>30.28</b>
L-7B	Random	1B	30.13	25.80	24.83	57.92	34.67
		2B	31.91	29.12	25.18	59.86	36.52
	ProLong	1B	29.50	29.04	24.46	59.04	35.51
		2B	30.10	30.58	25.84	59.52	36.51
	LADM	1B	<b>32.24</b>	<b>31.10</b>	<b>25.85</b>	<b>62.12</b>	<b>37.83</b>
M-7B	Random	1B	24.08	24.63	22.51	60.85	33.02
		2B	29.89	26.60	23.61	58.84	34.73
	ProLong	1B	23.76	24.43	22.73	61.82	33.18
		2B	27.41	28.85	24.29	<b>63.36</b>	35.98
	LADM	1B	<b>33.85</b>	<b>29.09</b>	<b>24.47</b>	60.58	<b>37.00</b>
L-13B	Random	1B	28.23	32.02	25.67	64.75	37.67
		2B	27.93	34.87	25.00	<b>66.13</b>	38.48
	ProLong	1B	28.94	33.20	25.62	64.62	38.04
		2B	<b>33.55</b>	34.37	24.84	62.37	38.78
	LADM	1B	29.03	<b>35.53</b>	<b>25.82</b>	65.28	<b>38.92</b>

Table 4: Performance comparison of random sampling and our LADM method on the LongBench dataset.

illustrated in Appendix C. These results highlight

that our proposed data selection framework can effectively extract high-quality data with strong contextual dependencies from large-scale pre-training corpora, thus enhancing the long-context modeling capability of LLMs while reducing training costs.

## 6.2 Data Selection Efficiency of LADM

We first analyze the computational complexity of each stage in the LADM framework. Each PFS calculation has a complexity of  $O(l^2 \cdot d_k)$ . The complexity for AFS is  $O\left(\frac{N \cdot l^2 \cdot d_k}{d_{AFS}}\right)$  and for the overall CDS is  $O\left(\frac{N^2 \cdot l^2 \cdot d_k}{d_{CDS} \cdot d_{AFS}}\right) = O\left(\frac{L^2 \cdot d_k}{d_{CDF} \cdot d_{AFS}}\right)$ , where  $l = 128$ ,  $N = 256$ ,  $L = 32k$ , and  $d_{AFS}, d_{CDS}$  are the strides for calculating AFS and CDS. The two parameters decide the number of PFS calculations required and reduce the cost of full attention calculation with complexity  $O(L^2 \cdot d_k)$ , thus affecting the data selection efficiency.

Method	$d_{AFS}$	$d_{CDS}$	Sec/sample	Correlation
LADM	4	4	2.46	1.000
	2	4	2.47	0.719
	4	2	3.95	0.721
	2	2	3.98	0.719
ProLong	-	-	2.46	-

Table 5: The efficiency and Pearson Correlation under different hyper-parameter settings and methods.

We conducted experiments on a randomly selected set of 5,000 samples from our pre-training dataset under different hyper-parameter settings. We compare both the computational overhead and the impact on the sample-level CDS. We use the Pearson Correlation Coefficient to measure the consistency of CDS across different configurations. As shown in Table 5, all the coefficient values are greater than 0.7, indicating the consistent results of data selection. For computation efficiency, since we can get all PFS of the current span and its previous spans through multiplication of the query and key matrices, a smaller  $d_{AFS}$  does not significantly increase computation overhead. However, smaller  $d_{CDS}$  requires calculating additional AFS for new spans, resulting in a notable increase in time overhead. Based on these analyses, we select the setting  $d_{AFS} = d_{CDS} = 4$  to minimize the computational cost. We also present the data selection efficiency of ProLong in Table 5. With comparable efficiency, LADM can achieve better performance, as shown in Section 5, demonstrating the effectiveness and applicability of our method.

### 6.3 Ablation Study on LADM

Method	SD-QA	MD-QA	SUM	CODE	AVG
LADM	<b>32.24</b>	<b>31.10</b>	<b>25.85</b>	<b>62.12</b>	<b>37.83</b>
<i>w/o std</i>	30.71	30.10	25.34	61.38	36.88
<i>w/o length</i>	31.60	30.90	25.38	58.08	36.49

Table 6: Ablation study of contextual dependency measurement in LADM.

To validate the effectiveness of the contextual dependency measurement in LADM, we conduct additional ablation studies on the standard deviation weights  $\sigma_j$  in Eq. 4 and the length weights  $\frac{j-i}{N}$  and  $\frac{j}{N}$  in Eq. 4 and Eq. 5. Experimental results of Llama2-7B on the LongBench dataset are shown in Table 6. Higher standard deviation weights indicate more diverse dependency patterns within long-context samples and higher length weights focus




Figure 4: Median scores for various data categories from the Pile dataset under ProLong and LADM framework.

more on span pairs with greater distances, making them crucial for long-range contextual dependency measurement. Therefore, we observe a significant performance drop without these weights.

### 6.4 Observations and Findings

We conduct a series of statistical analyses on the sample-level CDS derived by our LADM framework. To mitigate the influence of outliers, we calculate the median scores for each category from the pre-training data.

In Figure 4, we show the median scores of ProLong and LADM framework on different categories. The results show that articles (PhilPapers, ArXiv, and PubMed Central), legal documents (FreeLaw), and books (PG-19) tend to receive higher scores in our method. This is likely due to their coherent logic and strong interrelations between paragraphs, indicating complex and varied dependencies within these data types.

Moreover, we analyze the features of data with lower scores. Wikipedia data samples often manifest as concatenated paragraphs that may only share a theme but exhibit weak connectivity between paragraphs. Similarly, the Ubuntu IRC dataset is composed of group chat discussions. These conversations show duplicated patterns and different sections involve different users, time points and topics, resulting in weak contextual coherence. We show samples from different categories in Appendix D.

Compared to our method, ProLong tends to give



much higher scores to samples from the Ubuntu IRC dataset containing segments with similar patterns, despite low relevance between them. Moreover, samples from the PG-19 dataset get the lowest score, which indicates that ProLong may struggle to capture the deeper contextual relationships between seemingly unrelated sections within books. These sections may be connected through narrative development, which is not apparent through analysis of isolated segments. This highlights the importance of measuring the dependencies within the full context rather than just assessing the relationship between individual segments.

## 7 Conclusion

This paper introduces LADM, a novel and efficient long-context data selection framework to identify high-quality long-context data with long-range and diverse contextual dependencies from a large-scale, multi-domain pre-training corpus. LADM utilizes the accumulated attention scores over long contexts to quantify the dependencies between spans and aggregate them as a metric for sample-level contextual dependency measurement. The experimental results on various long-context tasks further demonstrate that our method can significantly enhance the long-context modeling capability of LLMs.

## Limitations

The limitations of our work can be summarized as follows: Firstly, to measure the long-range contextual dependency, we use a tiny model for data selection. This may introduce additional computational overhead. Secondly, we do not conduct experiments on LLMs exceeding 13B parameters, due to the great cost of training long-context LLMs. Thirdly, due to the scarcity of long-context data resources, we only conduct experiments on 32K context length. It is worthwhile to explore ways to construct high-quality data with longer context.

## Acknowledgments

We thank our colleagues Chen Wang, Chong Li, Zixuan Ren, Yaochen Zhu, and Pu Jian for their insightful and constructive feedback. Furthermore, we thank all reviewers for their detailed reviews and valuable comments. This work is supported by the National Key R&D Program of China No.2022ZD0160602 and the Strategic Priority Research Program of Chinese Academy of Sciences under Grant No.XDA04080400. This work is

also supported by Zhongguancun Academy Project No.20240103.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Chenxin An, Fei Huang, Jun Zhang, Shansan Gong, Xipeng Qiu, Chang Zhou, and Lingpeng Kong. 2024a. [Training-free long-context scaling of large language models](#). In *Forty-first International Conference on Machine Learning*.
- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, and Jian-Guang Lou. 2024b. Make your llm fully utilize the context. *arXiv preprint arXiv:2404.16811*.
- Zhangir Azerbayev, Edward Ayers, and Bartosz Piotrowski. 2022. [Proof-pile](#).
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. 2023. Longbench: A bilingual, multitask benchmark for long context understanding. *arXiv preprint arXiv:2308.14508*.
- Yushi Bai, Jiajie Zhang, Xin Lv, Linzhi Zheng, Siqi Zhu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. Longwriter: Unleashing 10,000+ word generation from long context llms. *arXiv preprint arXiv:2408.07055*.
- bloc97. 2023. [Ntk-aware scaled rope allows llama models to have extended \(8k+\) context size without any fine-tuning and minimal perplexity degradation](#).
- Guangzheng Chen, Xin Li, Michael Shieh, and Lidong Bing. 2025a. [LongPO: Long context self-evolution of large language models through short-to-long preference optimization](#). In *The Thirteenth International Conference on Learning Representations*.
- Jianghao Chen, Zhenlin Wei, Zhenjiang Ren, Ziyong Li, and Jiajun Zhang. 2025b. [Lr@ bench: Evaluating long-chain reflective reasoning capabilities of large language models via constraint satisfaction problems](#). *arXiv preprint arXiv:2502.17848*.
- Longze Chen, Ziqiang Liu, Wanwei He, Yinhe Zheng, Hao Sun, Yunshui Li, Run Luo, and Min Yang. 2024a. [Long context is not long at all: A prospector of long-dependency data for large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8222–8234, 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*.

- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2024b. [LongLoRA: Efficient fine-tuning of long-context large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Tri Dao. 2024. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610.
- Hantian Ding, Zijian Wang, Giovanni Paolini, Varun Kumar, Anoop Deoras, Dan Roth, and Stefano Soatto. 2024. [Fewer truncations improve language modeling](#). In *Forty-first International Conference on Machine Learning*.
- Alexander Richard Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084.
- Lizhe Fang, Yifei Wang, Zhaoyang Liu, Chenheng Zhang, Stefanie Jegelka, Jinyang Gao, Bolin Ding, and Yisen Wang. 2024. What is wrong with perplexity for long-context language modeling? *arXiv preprint arXiv:2410.23771*.
- Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hananeh Hajishirzi, Yoon Kim, and Hao Peng. 2024. [Data engineering for scaling language models to 128k context](#). In *Forty-first International Conference on Machine Learning*.
- Chaochen Gao, Xing Wu, Qi Fu, and Songlin Hu. 2024a. [Quest: Query-centric data synthesis approach for long-context scaling of large language model](#). *arXiv preprint arXiv:2405.19846*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonnell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2024b. [A framework for few-shot language model evaluation](#).
- Xinyang Geng and Hao Liu. 2023. [Openllama: An open reproduction of llama](#).
- Boyu Guan, Yining Zhang, Yang Zhao, and Chengqing Zong. 2025. [TriFine: A large-scale dataset of vision-audio-subtitle for tri-modal machine translation and benchmark with fine-grained annotated tags](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, Abu Dhabi, UAE. Association for Computational Linguistics.
- Chi Han, Qifan Wang, Hao Peng, Wenhan Xiong, Yu Chen, Heng Ji, and Sinong Wang. 2024. [Lm-infinite: Zero-shot extreme length generalization for large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3991–4008.
- Junqing He, Kunhao Pan, Xiaoqun Dong, Zhuoyang Song, LiuYiBo LiuYiBo, Qianguosun Qianguosun, Yuxin Liang, Hao Wang, Enming Zhang, and Jiaxing Zhang. 2024. [Never lost in the middle: Mastering long-context question answering with position-agnostic compositional training](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13628–13642, Bangkok, Thailand. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Cheng-Yu Hsieh, Yung-Sung Chuang, Chun-Liang Li, Zifeng Wang, Long Le, Abhishek Kumar, James Glass, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, et al. 2024. [Found in the middle: Calibrating positional attention bias improves long context utilization](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 14982–14995.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Yutong Hu, Quzhe Huang, Mingxu Tao, Chen Zhang, and Yansong Feng. 2024. [Can perplexity reflect large language model’s ability in long text understanding?](#) In *The Second Tiny Papers Track at ICLR 2024*.
- Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji, and Lu Wang. 2021. [Efficient attentions for long document summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1419–1436.

- Pu Jian, Donglei Yu, and Jiajun Zhang. 2024. Large language models know what is key visual entity: An llm-assisted multimodal retrieval for vqa. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10939–10956.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. 2024. **LLM maybe longLM: Selfextend LLM context window without tuning**. In *Forty-first International Conference on Machine Learning*.
- Greg Kamradt. 2023. Needle in a haystack - pressure testing llms. [https://github.com/gkamradt/LLMTest\\_NeedleInAHaystack](https://github.com/gkamradt/LLMTest_NeedleInAHaystack).
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Sarthak Mittal, Sharath Chandra Raparthy, Irina Rish, Yoshua Bengio, and Guillaume Lajoie. 2022. **Compositional attention: Disentangling search and retrieval**. In *International Conference on Learning Representations*.
- 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*.
- Chau Pham, Simeng Sun, and Mohit Iyyer. 2024. Suri: Multi-constraint instruction following in long-form text generation. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1722–1753.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Konrad Staniszewski, Szymon Tworowski, Sebastian Jaszczur, Henryk Michalewski, Łukasz Kuciński, and Piotr Miłoś. 2023. Structured packing in llm training improves long context utilization. *arXiv preprint arXiv:2312.17296*.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Wei Sun, Qianlong Du, Fuwei Cui, and Jiajun Zhang. 2025a. An efficient and precise training data construction framework for process-supervised reward model in mathematical reasoning. *arXiv preprint arXiv:2503.02382*.
- Wei Sun, Wen Yang, Pu Jian, Qianlong Du, Fuwei Cui, Shuo Ren, and Jiajun Zhang. 2025b. Ktae: A model-free algorithm to key-tokens advantage estimation in mathematical reasoning. *arXiv preprint arXiv:2505.16826*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. Longvlm: Efficient long video understanding via large language models. In *European Conference on Computer Vision*, pages 453–470. Springer.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. 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.
- Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. 2024. Retrieval head mechanistically explains long-context factuality. *arXiv preprint arXiv:2404.15574*.

- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. [Efficient streaming language models with attention sinks](#). In *The Twelfth International Conference on Learning Representations*.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajjwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, et al. 2023. Effective long-context scaling of foundation models. *arXiv preprint arXiv:2309.16039*.
- Zheyang Xiong, Vasilis Papageorgiou, Kangwook Lee, and Dimitris Papailiopoulos. 2024. From artificial needles to real haystacks: Improving retrieval capabilities in llms by finetuning on synthetic data. *arXiv preprint arXiv:2406.19292*.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2023. Retrieval meets long context large language models. *arXiv preprint arXiv:2310.03025*.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. Demystifying long chain-of-thought reasoning in llms. *arXiv preprint arXiv:2502.03373*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas Bertasius. 2024. A simple llm framework for long-range video question-answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21715–21737.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, et al. 2021. Qmsum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921.
- Shiyi Zhu, Jing Ye, Wei Jiang, Siqiao Xue, Qi Zhang, Yifan Wu, and Jianguo Li. 2024. Coca: Fusing position embedding with collinear constrained attention in transformers for long context window extending. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4247–4262.



## A Pre-training Datasets




Figure 5: Data length distribution of the Pile dataset.

Figure 5 shows the data length distribution of the Pile dataset. Due to the scarcity of long-context data resources, we conduct experiments on 32K samples without extending to longer data samples. We present the detailed composition of the Pile dataset after removing samples shorter than 32k in Table 7. All subsequent data selection processes are based on this refined dataset.

Type	Samples	Length
ArXiv	216,616	49,207
EuroParl	17,340	71,646
FreeLaw	27,663	47,264
Github	54,844	44,336
Gutenberg (PG-19)	34,037	80,471
Others	1,091	50,221
PhilPapers	3,777	69,045
Pile-CC	60,066	56,367
PubMed Central	46,293	46,890
USPTO Backgrounds	3,762	51,124
Ubuntu IRC	4,607	84,491
Wikipedia (en)	4,348	52,711

Table 7: Composition of the pre-training dataset. "Samples" denotes the number of samples and "Length" denotes the average length of samples measured by LlamaTokenizer.

## B Training Details

We train all models using the language modeling cross-entropy loss and set batch size to 1M tokens and learning rate to  $2 \times 10^{-5}$  without weight decay. We employ a linear learning rate scheduler with a warm-up period of 20 steps and use the AdamW optimizer (Loshchilov and Hutter, 2017) with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . All models are trained with BF16 precision, Flash Attention 2

### ArXiv, CDS=0.635

#### Abstract

This paper presents a group of analytical formulas for calculating

#### Introduction

This is the third part of the present author's work on quadratic matrix-valued functions and their algebraic properties.

#### Problem formulation

Matrix rank and inertia optimization problems are a class of discontinuous optimization problems

#### Preliminary results

The ranks and inertias are two generic indices in finite dimensional algebras.

### PhilPapers, CDS=0.360

#### Does the Prisoner's Dilemma Refute the Coase Theorem?

##### I. Introduction

Building upon the main theme of this year's LatCrit Conference

##### II. Standard Versions of the Prisoner's Dilemma

By way of background, we begin this paper by presenting the standard

##### III. Coasean Version of the Dilemma (with Strategic and NonStrategic Bargaining)

The previous section discussed the standard version of the Prisoner's Dilemma

### PG-19, CDS=0.278

#### POWER THROUGH REPOSE

##### I. THE GUIDANCE OF THE BODY

THE literature relating to the care of the human body is already very extensive.

##### II. PERVERSIONS IN THE GUIDANCE OF THE BODY

SO evident are the various, the numberless perversions of our powers in the misuse of the machine

##### III. REST IN SLEEP

HOW do we misuse our nervous force?

##### IV. OTHER FORMS OF REST

DO you hold yourself on the chair,

Figure 6: Data samples with high CDS scores from ArXiv, PhilPapers and PG-19.

(Dao, 2024), HuggingFace Trainer (Wolf et al., 2020), and Deepspeed engine (Rajbhandari et al., 2020).

### C Training Efficiency

In Table 8, we show the performance comparison of LADM with random sampling baseline under different training budgets on Llama2-7B.

Model	Method	Tokens	SD-QA	MD-QA	SUM	CODE	AVG
L-7B	Random	1B	30.13	25.80	24.83	57.92	34.67
		2B	31.91	29.12	25.18	59.86	36.52
		3B	30.28	30.83	25.33	61.91	37.09
		4B	31.81	29.86	25.37	60.37	36.85
	LADM	1B	<b>32.24</b>	<b>31.10</b>	<b>25.85</b>	<b>62.12</b>	<b>37.83</b>

Table 8: Performance comparison between random sampling with additional training tokens and our LADM method on the LongBench dataset.

### D Data samples

As shown in Figure 6, data samples with high CDS scores are characterized by rich content, coherent logic, and good readability. These samples exhibit interrelations between paragraphs, demonstrating strong long-range contextual dependencies.

**Wikipedia (en), CDS=0.127**

```
List of minor planets: 48001–49000
48001–48100
|-bgcolor=#fefefe
| 48001 || — || January 24, 2001 || Socorro ||
LINEAR || H || align=right | 1.9 km ||
|-id=002 bgcolor=#fefefe
| 48002 || — || January 19, 2001 || Socorro ||
LINEAR || NYS || align=right | 1.9 km ||
|-id=003 bgcolor=#d6d6d6
| 48003 || — || January 21, 2001 || Socorro ||
LINEAR || — || align=right | 8.0 km ||
|-id=004 bgcolor=#fefefe
.....
```

Figure 7: Data sample with low CDS scores from Wikipedia (en).

We also present samples with low CDS scores. Data from Wikipedia in Figure 7 exhibits a repetitive structure, primarily consisting of numbers and brief descriptions, lacking detailed information. Chat records Data from Ubuntu IRC in Figure 8 contains a lot of abbreviations and informal expressions. The conversations are composed of content in different time periods, making it difficult to form a coherent information flow. We also observe samples concatenated from completely unrelated data

**Ubuntu IRC, CDS=0.139**

```
#ubuntu-my 2011-04-25
<sweemeng> yo
#ubuntu-my 2011-04-26
<shahbandar72> hello..
<shahbandar72> not many people around it
seems..
#ubuntu-my 2011-04-27
* fenris-unity pokes mnajem
<mnajem> yup
<fenris-unity> what are you up to bro
* ejat changed the topic of
#ubuntu-my 2011-04-28
<yuskhazab> hello
<yuskhazab> peace
<yuskhazab> oyotat, are you there
<fenris-> good evening everyone
<yuskhazab> good evening
<yuskhazab> fenris-, not sleeping?
<fenris-> about to sleep
<yuskhazab> is there anyone here who hasn't
slept yet
.....
```

Figure 8: Data sample with low CDS scores from Ubuntu IRC.

**Pile-CC, CDS=0.119**

```
No, I'm not particularly sporty alli coupons
walgreens But economic net income
(ENI).....

A staff restaurant malegra 120 reviews «The
Duchess of Cambridge has graced the store
with her regal presence several times .....

Looking for a job Buy Precose Online Jeffrey
Halley, FX trader for Saxo Capital
Markets.....

Languages lipotrexate price Zulima Niebles
said police told her that .....

I like watching TV cozaar tablet Brzeski's
business model is simple.....

Insert your card synthroid dosage .125 mg
Fluent in Somali.....
.....
```

Figure 9: Data sample with low CDS scores from Pile-CC.

without any logical connection, as shown in Figure 9. Therefore, these data samples receive lower CDS due to the lack of long-range contextual dependency.

### E Long Attention Calculator trained with fewer tokens

In Table 9, we provide the performance of Llama2-7B on the LongBench dataset using Long Attention

Model	Method	LAC-Tokens	Single-Document QA				Multi-Document QA			
			NarrativeQA	Qasper	MultiFieldQA	AVG	HotpotQA	2WikiMQA	MuSiQue	AVG
L-7B	Random	-	24.06	33.27	33.05	30.13	35.29	26.79	15.31	25.80
	ProLong	-	25.00	28.23	35.27	29.50	40.30	28.91	17.91	29.04
	LADM	1B	26.30	<b>34.97</b>	<b>39.56</b>	<b>33.61</b>	41.65	30.72	<b>19.22</b>	30.53
	LADM	5B	<b>26.34</b>	32.28	38.11	32.24	<b>43.42</b>	<b>31.85</b>	18.03	<b>31.10</b>

Model	Method	LAC-Tokens	Summarization				Code			Overall
			GovReport	QMSum	MultiNews	AVG	LCC	RepoBench-P	AVG	
L-7B	Random	-	29.56	21.26	23.66	24.83	59.32	56.51	57.92	34.67
	ProLong	-	29.54	20.49	23.35	24.46	62.78	55.30	59.04	35.51
	LADM	1B	29.44	<b>21.97</b>	24.89	25.43	58.81	51.65	55.23	36.20
	LADM	5B	<b>29.77</b>	21.75	<b>26.02</b>	<b>25.85</b>	<b>65.78</b>	<b>58.46</b>	<b>62.12</b>	<b>37.83</b>

Table 9: Performance comparison of baseline methods and our LADM method using Long Attention Calculator trained with different numbers of tokens (denoted as "LAC-Tokens").

Model	Method	MMLU	ARC-c	ARC-e	HellaSwag	TruthfulQA	AVG
OL-3B	-	25.26	36.52	63.80	69.91	21.30	43.36
	Random	24.89	34.04	62.63	68.36	21.18	42.22 (-1.14)
	ProLong	26.11	35.84	63.68	68.46	20.56	42.93 (-0.43)
	LADM	25.33	36.26	64.44	68.64	20.69	43.07 (-0.29)
L-7B	-	40.83	46.25	74.41	76.06	25.21	52.55
	Random	37.83	44.97	74.41	75.15	26.32	51.74 (-0.81)
	ProLong	39.63	44.80	74.83	75.95	24.97	52.04 (-0.51)
	LADM	40.96	45.14	74.07	76.06	26.19	52.49 (-0.06)
M-7B	-	59.59	53.92	79.55	81.03	28.15	60.45
	Random	38.17	43.09	69.40	74.65	24.72	50.01 (-10.44)
	ProLong	44.37	49.40	77.23	78.05	25.09	54.83 (-5.62)
	LADM	52.76	48.12	75.34	78.51	25.46	56.04 (-4.41)
L-13B	-	52.09	49.40	77.61	79.37	25.34	56.76
	Random	50.53	50.60	76.77	79.36	25.46	56.54 (-0.22)
	ProLong	50.61	50.00	77.31	79.44	25.95	56.66 (-0.10)
	LADM	50.37	49.66	77.57	79.92	26.44	56.79 (+0.03)

Table 10: Performance comparison of different models on short-context tasks.

Calculator trained with different numbers of tokens. The results show that even with fewer training tokens, the Long Attention Calculator can still capture the long-range contextual dependencies and select high-quality data, resulting in the better average performance than baseline methods. This demonstrates the efficiency and robustness of our data selection method.

## F Performance on short-context tasks

In Table 10, we present performance of different models on short-context tasks, including MMLU (Hendrycks et al., 2021), ARC (Clark et al., 2018),

HellaSwag (Zellers et al., 2019) and TruthfulQA (Lin et al., 2022). We use the lm-evaluation-harness (Gao et al., 2024b) for evaluation. Method "-" denotes the original short-context models. For OpenLlama-3B-v2 and Llama2-7B/13B, LADM maintains strong performance on all types of short-text tasks and even shows slight improvement over baseline methods. For Mistral-7B-v0.1, we observe performance decline on all methods, but LADM exhibits the least performance drop compared to other methods.