# Mask-Enhanced Autoregressive Prediction: Pay Less Attention to Learn More

Xialie Zhuang<sup>\*12</sup> Zhikai Jia<sup>\*2</sup> Jianjin Li<sup>\*3</sup> Zhenyu Zhang<sup>4</sup> Li Shen<sup>5</sup> Zheng Cao<sup>2</sup> Shiwei Liu<sup>6</sup>

# Abstract

Large Language Models (LLMs) are discovered to suffer from accurately retrieving key information. To address this, we propose Mask-Enhanced Autoregressive Prediction (MEAP), a simple yet effective training paradigm that seamlessly integrates Masked Language Modeling (MLM) into Next-Token Prediction (NTP) to enhance the latter's in-context retrieval capabilities. Specifically, MEAP first randomly masks a small fraction of input tokens and then directly performs the standard next-token prediction autoregressive using a decoder-only Transformer. MEAP eliminates the need for bidirectional attention or encoderdecoder architectures for MLM, incurring no additional computational overhead during pre-training or inference. Intensive experiments demonstrate that MEAP substantially outperforms NTP on key information retrieval and long-context reasoning tasks, while performing on par or better on commonsense reasoning tasks. The benefits of MEAP also extend to supervised fine-tuning, where it shows remarkable advantages in lostin-the-middle scenarios, outperforming NTP by 11.77% percentage points. Our analysis indicates that MEAP's effectiveness arises from its ability to promote more distinguishable attention scores by concentrating on a reduced set of non-masked tokens. This mechanism improves the model's focus on task-relevant signals while mitigating the influence of peripheral context. These findings position MEAP as a promising training paradigm for large language models. Code is available at https://github.com/scitix/MEAP.

# 1. Introduction

Next-token prediction (NTP) (Radford, 2018) is the foundational training objective for many large language models (LLMs), including OpenAI's GPT series (Brown, 2020). NTP trains models to predict the next word (or token) in a sequence, given all preceding tokens. Its scaling efficiency and exceptional performance in text generation have established it as the dominant paradigm for state-of-the-art LLMs such as GPT-4 (Achiam et al., 2023), LLaMa3 (Dubey et al., 2024), Gemini 1.5 Pro (Team et al., 2024), and DeepSeek-V3 (Liu et al., 2024a). However, recent studies highlight the limitations of NTP-based LLMs in accurately retrieving key information from context (Liu et al., 2024b; Kamradt, 2023; Nelson et al., 2024).

In contrast, masked language modeling (MLM), used in BERT (Devlin, 2018), adopts a denoising objective that reconstructs masked inputs using bidirectional attention. This cloze-type nature makes MLM particularly effective for tasks requiring precise information retrieval and sentencelevel understanding. However, MLM's inherent focus on reconstructing masked tokens reduces its effectiveness in tasks requiring coherent and long-form text generation (Wang & Cho, 2019; Dong et al., 2019).

While intuitively appealing, combining NTP and MLM to leverage their respective strengths remains a non-trivial challenge. MLM typically operates best within two-stack encoder-decoder architectures, and performance degrades significantly when applied to decoder-only Transformers (Tay et al., 2022). Efforts to integrate the two often rely on unified pre-training pipelines where multiple objectives are alternated during the pretraining process (Dong et al., 2019; Tay et al., 2022). However, this multi-objective approach introduces substantial complexity to the training pipeline, making it cumbersome to scale, especially for models with billions or trillions of parameters.

To this end, we propose **Mask-Enhanced Autoregressive Prediction (MEAP)**, a simple yet effective LLM training paradigm that seamlessly integrates masked tokens into next-token prediction. Specifically, we first randomly mask a small fraction of the input tokens and then directly perform standard next-token prediction using a decoder-only Transformer in an autoregressive manner. This straightforward

<sup>&</sup>lt;sup>\*</sup>Equal contribution <sup>1</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences, China <sup>2</sup>SCITIX (SGP) TECH PTE. LTD., Singapore <sup>3</sup>South China Normal University, China <sup>4</sup>University of Texas at Austin, USA <sup>5</sup>Sun Yat-Sen University, China <sup>6</sup>University of Oxford, UK. Correspondence to: Zheng Cao <zcao@scitix.ai>, Shiwei Liu <shiwei.liu@maths.ox.ac.uk>.

Proceedings of the  $42^{nd}$  International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).



*Figure 1.* Overview of next token prediction, masked language modeling, and our MEAP.

modification eliminates the need for bidirectional attention or an expensive encoder-decoder architecture, thereby incurring no additional computational overhead during training. During inference, our resulting LLMs can work as simply as LLMs that are trained with NTP with no extra engineering effort. The simplicity of MEAP enables us to enhance LLMs' performance of key information retrieval and longcontext reasoning, while retaining the impressive scaling efficiency of decoder-only LLMs. Figure 1 shows the illustrations of different training paradigms.

As a general pre-training paradigm, MEAP works effectively for scenarios of pre-training and fine-tuning. For the pre-training setting, we conduct control experiments by pre-training 1.1B LLaMa-style LLMs (Zhang et al., 2024) with NTP and MEAP, where the training tokens scale from 40B to 200B. Our results demonstrate that MEAP substantially improves the performance in key information retrieval tasks such as Needle in a Haystack (Kamradt, 2023) by up to 33% average score and Multi-Document Question Answering (MDQA) (Liu et al., 2024b) by up to 27.2 percentage points, while preserving general knowledge learned during pre-training. It is noteworthy that MEAP achieves 85.8% accuracy with 60B training tokens on the Needle in a Haystack, while NTP requires 200B for similar performance, highlighting MEAP's superior data efficiency in key information retrieval. In addition, compared to the original NTP, MEAP also suffers less from hallucination.

In addition, the promise of MEAP also holds for LLM fine-tuning. Our MEAP framework demonstrates consistent improvements across multiple commonsense reasoning tasks, achieving an average gain of 1.12 scores over the NTP baseline. On Multi-Document Question Answering, MEAP achieves an average improvement of 11.77% across all positions.

Our analysis suggests that MEAP's effectiveness stems from its ability to enhance attention distinguishability by focusing on a reduced set of non-masked tokens. This mechanism sharpens the model's attention to task-relevant signals while reducing the impact of peripheral context. In essence, MEAP learns more by attending to fewer tokens.

The structure of this paper is as follows. Section 3 details the MEAP algorithm. The evaluation of MEAP on LLM pretraining and fine-tuning is presented in Sections 4.1 and 4.2, respectively. In Section 5, we further analyze the underlying reasons for MEAP's effectiveness. Section 6 provides an ablation study, and we conclude the paper in Section 7.

# 2. Related Work

Masked Language Modeling. Pre-training is one of the most important pillars of LLMs. BERT first trained a bidirectional, encoder-only Transformer with masked language modeling (MLM), where the model is trained to predict masked input tokens. XLNet (Yang, 2019) introduced the Permutation-based Language Modeling to account for dependencies between masked tokens during training. RoBERTa (Liu, 2019) further improves the pre-training of BERT by training the model longer, over more data, with longer sequences, etc. MLM was further advanced by T5 (Roberts et al., 2019). Specifically, T5 frames every text processing task as a 'text-to-text' problem, leveraging increased lengths of corrupted tokens to achieve improved performance on classification tasks, which has contributed to its growing popularity. However, these models have shown limited performance in open-text generation and in-context learning, limiting their usage in modern LLMs.

**Next Token Prediction.** In a parallel vein, Radford et al. (2019) proposed next-token prediction (NTP) where a decoder-only Transformer is trained to predict the next token from left to right using unidirectional attention ensured by casual mask. By predicting the next token based on previously generated tokens and the given input context, NTP maintains coherence and logical flow in the generated text, well-suited for text generation. Moreover, NTP eliminates the need for an encoder, significantly improving the scalability of language models. Due to the above advantages, NTP serves as the most popular pre-training objective of modern LLMs (Brown, 2020; Achiam et al., 2023; Touvron et al., 2023; Jiang et al., 2023; Yang et al., 2024; Liu et al., 2024a).

**Unified Training Paradigms.** There are works that propose unified training paradigms aiming to train one Transformer with multiple objective functions. For instance, UniLM (Dong et al., 2019) trains a bidirectional encoder on unidirectional language modeling (LM), bidirectional LM, and Sequence-to-Sequence LM. UL2 (Tay et al., 2022) proposes a unified pre-training paradigm with Mixture-of-Denoisers (MoD) to combine diverse pre-training paradigms together, improving the performance over T5 and GPT. While effective, the preference for encoder-decoder architectures and the complicated switches among different training objectives hinder their applications in practice.

In contrast, our approach seamlessly integrates masked tokens into NTP without incurring any additional pre-training



*Figure 2.* Training frameworks of MEAP: Left (Pre-training): A certain portion of input tokens is randomly masked, followed by standard next-token prediction (NTP). Right (Fine-tuning): Training samples are duplicated, and the random masking strategy is applied to the copied sequences. Standard NTP is then performed on the modified input for fine-tuning.

or inference costs, while preserving the ultra-efficiency of NTP. More importantly, MEAP is more suitable for modern LLMs, as our method does not alter the core mechanism of NTP, the resulting models remain fully compatible with existing pipelines, platforms, and hardware optimized for modern LLMs.

### 3. Mask-Enhanced Autoregressive Prediction

In this section, we introduce Mask-Enhanced Autoregressive Prediction (MEAP).

**LLM pre-training.** To enhance the performance of LLM in handling and understanding long texts, particularly in key tasks such as key information retrieval and long context, we designed and implemented a simple yet efficient random masking strategy. The core idea of this method is to selectively mask portions of the input during the pre-training phase. Specifically, we employed a fixed proportion masking mechanism, where tokens in the input sequence are randomly masked according to a predefined percentage *P*. In this way, the model is forced to learn in the absence of some contextual information, which helps improve its deep understanding and reasoning capabilities.

Formally, given a decoder-only Transformer  $\theta$  and a sequence of input  $X = (x_1, x_2, ..., x_{n-1}, x_n)$ , we first randomly mask a fraction of P tokens having  $X' = (x_1, [mask], ..., x_{t-1}, x_t)$ . We then perform the standard next-token prediction using the masked input in a left-to-right manner:

$$p_{\theta}(X') = \prod_{t=1}^{T} p_{\theta}(x_t \mid x_1, [mask], \dots, x_{t-1})$$

Same as NTP, when the model is tasked with predicting a masked token, it employs causal masked attention, using only the preceding tokens to predict the masked token. We carefully selected the masking ratio, P = 15% for pretraining, to ensure that the model receives an adequate level of training difficulty and learning signals, without excessively disrupting the pre-training process. The relatively moderate number of masked tokens allows this approach to be seamlessly integrated into existing NTP frameworks, without significantly increasing pre-training overhead or altering the original training procedure.

LLM fine-tuning. MEAP can also be extended to finetuning scenarios. In this scenario, we duplicate the training samples and apply the same random masking strategy to the copied sequences during fine-tuning. The original sequences and their masked counterparts are then combined into a single input sequence to feed into the model. The cross-entropy loss is computed only with the masked tokens  $(U_m)$  in the answer tokens  $(U_q)$ . This design addresses a critical concern: input sequences in supervised fine-tuning often contain key information essential for downstream tasks. Directly masking the original sequence risks removing crucial information, potentially compromising the model's performance on the target tasks. Masking the duplicated sequence incorporates MLM to NTP while avoiding this concern. We choose P = 10% for fine-tuning in our experiment. We only perform MEAP for the QA pair whose answer length exceeds 50, otherwise, we perform the standard NTP for the pair. Formally, the following objective of MEAP for fine-tuning is:

$$\mathcal{L}(\theta) = -\sum_{t \in U_q \cup U_m} \log p_{\theta}(x_t \mid x_1, \dots, x_{t-1}; \hat{x}_1, [\text{mask}], \dots, \hat{x}_{t-1})$$

where the sequence  $\{\hat{x}_i\}$  is a copy of the original sequence  $\{x_i\}$  (*i.e.*,  $\hat{x}_i = x_i$ ).

Notably, while MEAP doubles the sequence length during fine-tuning, Figure 5 shows that it achieves superior performance to NTP with only half the training time, essentially attaining stronger results with even fewer training tokens.

We believe that the effectiveness of MEAP stems from its ability to promote more distinguishable attention by focusing on fewer tokens during LLM training, as masked tokens typically receive negligible attention. This modification helps the model focus on task-relevant signals while reducing the impact of peripheral context, as verified in Section 5.

### **4. Experimental Results**

To evaluate the effectiveness of MEAP in training LLMs, we conduct controlled experiments comparing LLMs pretrained/fine-tuned by MEAP with those trained by NTP.

	ARC-c	ARC-e	BoolQ	PIQA	HellaSwag	WinoGrande	OBQA	Average
NTP	22.9	55.7	53.3	73.6	44.1	55.0	23.2	46.2
MEAP	25.4	56.4	59.5	72.3	43.4	55.3	22.6	47.8

*Table 1.* **Pre-training Evaluation.** Zero-shot performance of MEAP and NTP on various commonsense reasoning tasks. Results are measured directly after pre-training on 200B tokens with no fine-tuning.

#### 4.1. Pre-training Evaluation

**Setup.** We follow the Llama architecture (Touvron et al., 2023) as our base model. Specifically, we train 1.1B decoder-only Transformers (Vaswani, 2017) following the setting of Zhang et al. (2024). Our model has 24 layers with 32 attention heads, a hidden size of 2,048, an intermediate hidden size of 5,632, and a context length of 4096. We follow the common configurations of LLM components, e.g., Rotary Positional Embedding (RoPE) (Su et al., 2024), Pre-Norm (Ba, 2016) with RMSNorm (Zhang & Sennrich, 2019), SwiGLU (Shazeer, 2020), and Grouped-query Attention (Ainslie et al., 2023). To assess the scalability of MEAP, we increase the training token size from 40B to 60B, and further to 200B.

For all experiments, we implement a learning rate warm-up during the first 10% of the training steps, followed by a cosine annealing schedule, which decays the learning rate to 10% of its initial value. We use the AdamW optimizer with the following settings:  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ . The maximum learning rate is set to  $4 \times 10^{-4}$ , the minimum learning rate is  $4 \times 10^{-5}$ , and the weight decay is  $5 \times 10^{-2}$ .

#### 4.1.1. LANGUAGE MODELING EVALUATION

While the primary goal of MEAP is to enhance LLM performance in key information retrieval, it is essential to ensure that integrating MLM into NTP does not compromise the model's fundamental language modeling capability. To evaluate this, we employ the LM Eval Harness benchmark (Gao et al., 2024), assessing models in a zero-shot setting. The results, presented in Table 1, show that MEAP performs comparably to, or even outperforms, NTP, achieving a 1.6% improvement in the overall average score. This finding provides strong evidence that incorporating random masking into NTP does not degrade the model's language modeling capacity. In the following evaluations, we will examine whether MEAP further improves performance in key information retrieval and long-context modeling.

#### 4.1.2. NEEDLE-IN-A-HAYSTACK RETRIEVAL

For key information retrieval, we choose the wellestablished Needle-in-a-Haystack evaluation (Liu et al., 2024b), where the model is asked to retrieve the random fact or statement (the 'needle') in the middle of a long context Table 2. Single needle accuracy (%) with 32K context.

Token	40B	60B	200B
NTP	65.9	52.8	87.1
MEAP	80.2	85.8	98.2

window (the 'haystack'). This approach provides quantitative metrics for assessing precise information extraction from extended contexts, particularly relevant for document analysis applications.

As this evaluation involves long-context modeling capacity, we follow the setting of Ye et al. (2024) and conduct a length extension to 64K. In particular, we continue training our model for additional 4B tokens from SlimPajama (Soboleva et al.) using the approach proposed in (Fu et al., 2024). The implementation utilizes modified Rotary Position Embeddings with  $\theta_{\text{base}} = 640,000$ .

To demonstrate MEAP's scalability, we increase the training token size to 40B, 60B, and 200B, reporting the results of needle retrieval in Table 2. The results show that MEAP consistently outperforms NTP across different training scales. At 40B tokens, MEAP achieves 80.2% accuracy, significantly surpassing the baseline's 65.9%. The performance gap peaks at 60B tokens, with MEAP maintaining steady improvement and reaching 85.8% accuracy. At 200B tokens, MEAP approaches optimal performance, attaining 98.2% accuracy, while the NTP baseline still falls short of 90% accuracy. It is noteworthy that MEAP achieves 85.8% accuracy using just 60B training tokens, whereas NTP requires approximately three times as many (200B tokens) to reach a similar level. This demonstrates MEAP's superior data efficiency over NTP in key information retrieval.

We further illustrate the retrieval performance of our 200Btoken model with a 32K context length in Figure 3. The accuracy is reported across varying answer needle depths (y-axis) and context lengths (x-axis). The results show that MEAP generally maintains perfect accuracy across different context lengths and depths, with errors limited to only two grid cells. In contrast, NTP begins to exhibit accuracy degradation at a context length of 24K, affecting a wide range of depths from 50% to 100%.



*Figure 3.* Performance comparison between NTP and MEAP on Needle In A Haystack. Scores are computed using ROUGE-1, measuring unigram overlap between model responses and expected answers.

#### 4.1.3. MULTI-DOCUMENT QUESTION ANSWERING

*Table 3.* **Pre-training Evaluation.** Relative accuracy (%) improvement of MEAP over NTP on multi-document QA.

Answer Position	1	5	10	15	20
10 documents	+7.6	+7.0	+30.6	_	_
20 documents	+12.4	+4.0	+5.1	+3.7	+27.2

We evaluate the model's ability to retrieve information from long contexts using a multi-document QA task (Liu et al., 2024b) based on NaturalQuestions-Open (Kwiatkowski et al., 2019). The task requires identifying answers from a target document while ignoring k - 1 distractor documents, where k is the total number of documents in the context. We analyze performance across two context lengths with k = 10 and k = 20 documents and multiple answer positions. For each query, we construct an input context containing one target document with the annotated answer and k - 1 distractor documents that do not contain any of the answers. 64K-extended models are used for this evaluation.

We report the accuracy improvement of MEAP over NTP in Table 3. MEAP again consistently outperforms NTP by good margins across all configurations, with significant gains at later positions (+30.6% at position 3 in 10-doc, +27.2% at position 5 in 20-doc). These results indicate that MEAP enhances the model's ability to retrieve relevant information from long contexts, maintain performance across different context lengths and positions, and handle complex scenarios with multiple distractors. The improvements highlight the effectiveness of the masking strategy in enhancing the model's overall capability for long-context information retrieval tasks.



*Figure 4.* Long-context reasoning performance comparison between MEAP and NTP on the Multi-Needle Reasoning Task (M-RS) across different context lengths.

### 4.1.4. LONG-CONTEXT REASONING EVALUATION

We evaluate long-context reasoning capabilities using the Multi-Needle Reasoning Task (M-RS) (Li et al., 2024a), which requires models to retrieve and extract multiple pieces of information from long texts and using them to logically answer questions that demand an integrated understanding and reasoning of various text segments. This forces the model to distribute attention across contextually relevant tokens rather than focusing solely on local patterns.

We leverage the OpenCompass evaluation framework (Contributors, 2023) and report the results in Figure 4. MEAP consistently outperforms NTP across context lengths with **6.6** percentage point average improvement. demonstrates MEAP's enhanced capacity to maintain attention coherence over extended sequences. *Table 4.* Accuracy (i.e., free of hallucinations) on text summarization datasets evaluated by different LLM judges.

Task	XSum	MultiNews	WikiSum
NTP (Deepseek-V3)	0.09	0.17	0.24
MEAP (Deepseek-V3)	0.13	0.19	0.33
NTP (Qwen-Plus)	0.16	0.11	0.21
MEAP (Qwen-Plus)	0.19	0.14	0.27
NTP(GPT-40)	0.14	0.10	0.19
MEAP(GPT-4o)	0.16	0.13	0.24

### 4.1.5. CONTEXTUAL HALLUCINATION EVALUATION

Since MEAP improves more accurate key information retrieval, we expect it to suffer less from contextual hallucination. To verify, we evaluate MEAP in reducing contextual hallucinations on three summarization datasets: XSum (Narayan et al., 1808), WikiSum (Cohen et al., 2021), and MultiNews (Fabbri et al., 2019), following Ye et al. (2024). For this setting, we fine-tune the pre-trained models with Alpaca and evaluate them. We compare model-generated and reference summaries using (Deepseek-V3 (Liu et al., 2024a), Qwen-Plus (Yang et al., 2024) and GPT-40 (Hurst et al., 2024)) as the hallucination detector across 100 random samples per dataset. As shown in Table 4, our masking strategy achieves a consistent reduction in hallucination rates across all datasets.

### 4.2. Fine-tuning Evaluation

**Setup.** We fine-tune the Llama-3-8B (Dubey et al., 2024) on the Alpaca instruction dataset (Taori et al., 2023a). The training configuration uses a global batch size of 512. The model is optimized with AdamW ( $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ ), a learning rate of  $2 \times 10^{-5}$  (with 10% warmup and cosine decay), and weight decay set to 0.01. We retain key architectural components from Llama-3, such as RoPE embeddings (Su et al., 2024), RMSNorm (Zhang & Sennrich, 2019), and grouped-query attention (Ainslie et al., 2023).

During fine-tuning, we randomly mask a portion of tokens in the assistant's response, while keeping the source context intact. Only the masked tokens are predicted during finetuning. The training process uses bfloat16 precision with DeepSpeed Zero Stage 2 (Ren et al., 2021), and the Llama-3 tokenizer (Dubey et al., 2024) with a maximum sequence length of 4096 tokens.

### 4.2.1. LANGUAGE MODELING EVALUATION

Similar to the pre-training evaluation, we first assess MEAP's effectiveness in language modeling. Table 5 presents the evaluation results. Our MEAP framework demonstrates consistent improvements across multiple tasks, achieving an average gain of 1.12 scores over the NTP

baseline. The performance improvements are particularly notable on ARC-c and WinoGrande, indicating enhanced reasoning capabilities. The results highlight MEAP's effectiveness in fine-tuning complex reasoning tasks.

### 4.2.2. CROSS-MODEL GENERALIZABILITY

To verify that MEAP's effectiveness generalizes across different model architectures and scales, we conducted experiments on a diverse set of pre-trained LLMs. Table 6 presents results across various commonsense reasoning tasks, while Table 7 shows performance on multi-document QA tasks.

The results demonstrate that MEAP consistently matches or outperforms NTP across different model architectures and sizes. On the multi-document QA task, MEAP demonstrates substantial improvements across all model architectures and sizes. These results highlight MEAP's universal effectiveness in enhancing information retrieval capabilities.

#### 4.2.3. MULTI-DOCUMENT QUESTION ANSWERING

We evaluate MEAP's context-aware reasoning using the multi-document QA task with distractor suppression (Liu et al., 2024b). To ensure a fair comparison, we train MEAP for 2 epochs and NTP for 4 epochs, such that both approaches process a similar number of tokens. Table 8 quantifies the exact match (EM) improvements across critical document positions in the 20-document setting. MEAP consistently achieves notable gains across all positions, further demonstrating its superiority over NTP. Two key patterns emerge from the experimental results:

- *Consistent Improvement*: MEAP achieves substantial gains across all positions with an average improvement of 11.77%, showing robust performance throughout the document range.
- *Mid-Context Advantage*: The maximum improvement at position 20 (+15.22%) demonstrates enhanced long-range dependency modeling, crucial for connecting concepts across scientific documents.

These findings validate MEAP's effectiveness in preserving signal integrity across long contexts while highlighting opportunities for temporal reasoning enhancement and cross-document entity disambiguation.

### 4.3. Training Efficiency Analysis

MEAP introduces no additional overhead for pre-training or inference compared to standard NTP, as the only difference lies in the masking operation. During fine-tuning, MEAP requires duplicating the input sequence and training with a doubled sequence length, resulting in increased training overhead. This overhead, however, is effectively amortized

	ARC-c	ARC-e	BoolQ	PIQA	HellaSwag	WinoGrande	OBQA	Average
NTP	53.58	81.10	83.98	79.27	62.74	72.06	39.40	67.30
MEAP	55.12	83.21	83.82	81.01	63.31	74.27	38.20	68.42

*Table 5.* Fine-tuning Evaluation. Performance of MEAP and NTP on various commonsense reasoning tasks. Results are measured by fine-tuning with Llama-3-8B.

Table 6. Cross-Model Generalizability. Fine-tuning performance comparison of MEAP and NTP on commonsense reasoning tasks when applied to different architectures and model sizes.

Model	Method	ARC-c	ARC-e	BoolQ	PIQA	HellaSwag	WinoGrande	OBQA	Average
Llama-3.2-3B	NTP	47.95	69.07	75.54	76.50	72.43	64.33	44.40	64.32
Llama-3.2-3B	MEAP	<b>49.32</b>	<b>73.06</b>	71.80	<b>77.53</b>	<b>74.26</b>	<b>68.51</b>	<b>44.60</b>	65.58
Qwen2.5-14B	NTP	53.67	74.71	86.73	77.64	78.44	68.19	48.00	69.63
Qwen2.5-14B	MEAP	<b>56.83</b>	<b>79.38</b>	<b>87.37</b>	<b>79.33</b>	<b>79.37</b>	<b>72.69</b>	47.40	<b>71.77</b>
Mistral-7B-0.2	NTP	35.67	<b>60.10</b>	<b>75.81</b>	71.22	63.03	<b>61.40</b>	35.40	57.52
Mistral-7B-0.2	MEAP	<b>37.20</b>	59.18	72.63	<b>73.50</b>	<b>64.08</b>	61.17	<b>35.60</b>	<b>57.62</b>



*Figure 5.* Comparison of fine-tuning efficiency between MEAP and NTP. 'MEAP-n' refers to MEAP training for n epoch.

by MEAP's higher data utilization efficiency. Specifically, compared to NTP, MEAP requires only 50% of the epochs with a similar number of tokens being processed, while outperforming the latter significantly.

To verify, we report the results on the multi-document QA retrieval from 20 documents (Liu et al., 2024b), where retrieval performance is assessed by computing the average retrieval values across 5 positions. As shown in Figure 5, a single epoch of MEAP training significantly outperforms two epochs of NTP training by a large margin while also reducing total training time. This highlights MEAP's data efficiency, achieving similar or better results while reducing computational resources.

In summary, MEAP delivers significant training time reductions with improved or comparable performance on the retrieval task, highlighting its efficiency and effectiveness in large-scale training scenarios.

### 5. Why Does MEAP Work?

In this section, we attempt to interpret the underlying reasons for the effectiveness of MEAP. We conjecture that MEAP's effectiveness stems from its ability to promote more distinguishable attention by focusing on fewer tokens during LLM training, as masked tokens [MASK] are expected to receive marginal attention scores.

While effective, attention mechanisms in LLMs often struggle with long-context understanding, where redundant and non-informative attention is assigned to tokens (Liu et al., 2024b; Li et al., 2024b). A plausible explanation is that the attention module relies on the Softmax function to normalize attention scores within (0, 1), which tends to minimize differences among tokens, especially when training on sequences of thousands of tokens. This bears some resemblance to Martins et al. (2020)'s findings on sparse attention mechanisms which implicitly relate to their core mechanism (reducing attention to tokens).

Furthermore, LLMs exhibit a phenomenon known as attention sinks, where the initial few tokens receive disproportionately high attention scores compared to the rest (Xiao et al., 2023). Collectively, these factors lead to small and nearly indistinguishable attention scores across tokens, which is generally undesirable. When LLMs fail to properly differentiate between tokens, they are more likely to generate incorrect outputs.

By randomly replacing tokens with masks, MEAP implicitly

Model	Method	Position 1	Position 5	Position 10	Position 15	Position 20	Average
Llama-3.2-3B	NTP	13.60	12.09	12.54	12.69	14.35	13.05
Llama-3.2-3B	MEAP	<b>23.47</b>	<b>20.34</b>	<b>20.38</b>	<b>21.96</b>	<b>23.65</b>	21.96
Mistral-7B-0.2	NTP	36.96	30.55	27.82	27.55	38.79	32.33
Mistral-7B-0.2	MEAP	<b>37.91</b>	<b>32.98</b>	<b>31.46</b>	<b>32.22</b>	<b>43.45</b>	35.60
Qwen2.5-14B	NTP	60.00	51.98	56.01	56.05	63.39	57.49
Qwen2.5-14B	MEAP	<b>61.69</b>	<b>53.71</b>	<b>57.21</b>	<b>56.65</b>	<b>66.29</b>	<b>59.11</b>

Table 7. Cross-Model Generalizability. Accuracy (%) of MEAP and NTP on multi-document QA with 20 documents across different model architectures and sizes.

Table 8. Fine-tuning Evaluation. Accuracy (%) of MEAP and NTP on multi-document QA with 20 documents.

Position	1	5	10	15	20
NTP MEAP	24.29 33 22	22.82 34 16	24.11 36.01	25.46 <b>36 91</b>	31.11 46 33
$\frac{\Delta}{\Delta}$	+8.93	+11.34	+11.90	+11.45	+15.22

penalizes the attention scores at masked positions, thereby amplifying the attention differences among non-masked tokens. This masking mechanism encourages the model to generate more distinguishable attention scores, allowing it to focus on task-relevant texts while mitigating the influence of peripheral context. We validate this hypothesis through the following experiments.

#### 5.1. Masking Leads to More distinguishable Attention

To elucidate the mechanistic impact of our masking strategy on model behavior, we conducted a detailed analysis of attention distribution patterns. Our experimental protocol involved sampling 500 sequences. The original unmodified samples refer to the input sequence of NTP  $X_N$ , and their masked counterparts (same samples with 15% masks),  $X_M$ , designated as the input for MEAP. These sequence pairs were then processed through our 1.1B models pre-trained with NTP and MEAP, respectively. We compare two values: (1) Attention Score Decay: the percentage decrease in the averaged attention score at masked positions, computed as:

$$\frac{Att(X_N[mask = 1]) - Att(X_M[mask = 1])}{Att(X_N[mask = 1])}$$

(2) Attention Variance Increase: the attention variance increase at non-mask positions, computed as:

$$Var(Att(X_M[mask = 0])) - Var(Att(X_N[mask = 0]))$$

**Expectations.** We anticipate that the average attention score at masked positions will undergo a significant decline in the MEAP-trained model, indicating that masked tokens receive

minimal attention in MEAP. Consequently, this reduction is expected to increase the attention variance at non-masked positions, making the attention distribution in MEAP more distinguishable compared to NTP.

**Results.** Table 10 confirms our expectations. MEAP assigns 53.34% less attention to masked tokens, resulting in a 7.80% increase in attention variance. This finding validates that MEAP learns more distinguishable attention scores compared to NTP.

### 5.2. MEAP Focus More on Task-Relevant Tokens

To verify if MEAP learns more effective attention, we measure the average attention scores that the model assigns to different input segments. We structured our input sequences into distinct segments: context-before, answer, context-after, query, and EOS token. The complete input sequence was formed by concatenating these segments sequentially, followed by an EOS token. This structured format enabled precise tracking of attention allocation across different functional components.

**Expectation.** Our expectation is that MEAP tends to amplify attention to answer spans and meanwhile reduce the attention to less relevant tokens.

**Results.** The attention distributions during inference for both models are visualized in Figure 6. Notably, MEAP exhibits a substantial improvement in answer-relevant attention (34.5% vs. 9.4%) while reducing the dominance of context-before attention from 73.1% to 49.1%. Both models maintain similar attention levels for peripheral components, including context-after, query sections, and EOS tokens (all approximately 5%–6%). These results demonstrate that the MEAP framework enhances attention allocation during inference, prioritizing key information more effectively.

# 6. Ablation Study

We conduct ablation studies on the mask ratio for both pre-training and fine-tuning settings. Table 9 summarizes



*Figure 6.* Attention distribution comparison between NTP and MEAP during inference. The input sequence consists of: context-before ("In the heart of Paris, the Eiffel Tower stands tall, symbolizing both the city and the entire country."), answer ("Designed by Gustave Eiffel"), context-after (", it was completed in 1889 for the World's Fair. Originally criticized for its unusual design, it has since become one of the most recognizable landmarks in the world. Tourists from all over the globe visit it every year, making it one of the most photographed monuments."), and query ("question: Who designed the Eiffel Tower?"). MEAP allocates a much higher attention score to answer-relevant tokens (0.345) compared to NTP (0.094).

Table 9. Performance comparison of different mask ratios in pre-training and fine-tuning. The best results are highlighted in bold.

Pre-training						Fine-t	uning	
Mask Ratio	NTP	0.05	0.10	0.15	0.20   NTP	0.05	0.10	0.15
Accuracy	0.52	0.54	0.56	0.58	0.56   0.72	0.77	0.81	0.71

*Table 10.* Statistical analysis of attention score patterns between NTP and MEAP.

Length	Metric	Value	T-Stat/P-Value
1024	Score Decay	34.08%	-25.71/<1e-6
1024	Var. Increase	12.66%	12.26/<1e-6
4096	Score Decay	53.34%	-9.97/<1e-6
4096	Var. Increase	7.80%	5.22/<1e-6

the results. For pre-training, we evaluate our pre-trained model in Section 4.1 on the Multi-Document QA task using the nq-open-oracle dataset (Liu et al., 2024b). For fine-tuning, we train MEAP on the Alpaca dataset (Taori et al., 2023b) for 3 epochs with different mask ratios against standard NTP baselines with 6 epochs for a fair comparison. The results show that a mask ratio of **0.15** achieves the best performance in pre-training, while a mask ratio of **0.10** yields the highest accuracy in fine-tuning. MEAP consistently outperforms standard NTP in pre-training and fine-tuning, demonstrating its effectiveness in leveraging masked tokens for improved performance.

### 6.1. Effect of Different Masking Strategies

To further investigate the impact of masking patterns, we compare three distinct masking strategies: Random Masking

(our default approach), 5-Span Masking (consecutive spans of 5 tokens), and 50-Span Masking (longer spans of 50 consecutive tokens). We evaluate these strategies using a 0.3B parameter model pre-trained on 5B tokens, with results presented in Table 15.

# 7. Conclusion

This work addresses challenges in information processing through a straightforward approach that masks 10%–15% of input while maintaining traditional prediction methods. Our results show significant improvements in comprehension across longer contexts, achieved without additional computational costs. This approach demonstrates remarkable efficiency, matching performance metrics with just 60B training examples that typically require 200B examples with conventional methods. The results indicate that this strategy leads to more effective processing of key information through improved focus on relevant content. Since it requires no structural changes, this method can be readily integrated into existing systems without disrupting workflows.

# **Impact Statement**

This work proposes a modified pre-training paradigm that may influence how both industry and academia approach language model training. MEAP integrates seamlessly with existing LLM frameworks without requiring additional engineering effort or computational resources. While the improvement in information retrieval and reasoning capabilities could have broad implications for downstream applications, the method's computational efficiency and architectural compatibility mean it can be readily adopted within current training infrastructures. We anticipate this work will contribute to more efficient model development while maintaining established training pipelines and computational requirements.

# References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Ainslie, J., Lee-Thorp, J., de Jong, M., Zemlyanskiy, Y., Lebrón, F., and Sanghai, S. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. arXiv preprint arXiv:2305.13245, 2023.
- Ba, J. L. Layer normalization. arXiv preprint arXiv:1607.06450, 2016.
- Brown, T. B. Language models are few-shot learners. *arXiv* preprint arXiv:2005.14165, 2020.
- Cohen, N., Kalinsky, O., Ziser, Y., and Moschitti, A. Wikisum: Coherent summarization dataset for efficient human-evaluation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pp. 212– 219, 2021.
- Contributors, O. Opencompass: A universal evaluation platform for foundation models. https://github.com/open-compass/opencompass, 2023.
- Devlin, J. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dong, L., Yang, N., Wang, W., Wei, F., Liu, X., Wang, Y., Gao, J., Zhou, M., and Hon, H.-W. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32, 2019.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan,

A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

- Fabbri, A. R., Li, I., She, T., Li, S., and Radev, D. R. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. arXiv preprint arXiv:1906.01749, 2019.
- Fu, Y., Panda, R., Niu, X., Yue, X., Hajishirzi, H., Kim, Y., and Peng, H. Data engineering for scaling language models to 128k context. *arXiv preprint arXiv:2402.10171*, 2024.
- Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac'h, A., Li, H., McDonell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. A framework for few-shot language model evaluation, 07 2024. URL https://zenodo.org/records/ 12608602.
- Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al. Gpt-40 system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. Mistral 7b. arXiv preprint arXiv:2310.06825, 2023.
- Kamradt, G. Needle in a haystack-pressure testing llms. *Github Repository*, pp. 28, 2023.
- Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Li, M., Zhang, S., Liu, Y., and Chen, K. Needlebench: Can llms do retrieval and reasoning in 1 million context window?, 2024a. URL https://arxiv.org/abs/ 2407.11963.
- Li, T., Zhang, G., Do, Q. D., Yue, X., and Chen, W. Longcontext llms struggle with long in-context learning. arXiv preprint arXiv:2404.02060, 2024b.
- Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., et al. Deepseekv3 technical report. arXiv preprint arXiv:2412.19437, 2024a.
- Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., and Liang, P. Lost in the middle: How language models use long contexts. *Transactions of the*

Association for Computational Linguistics, 12:157–173, 2024b.

- Liu, Y. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692, 364, 2019.
- Martins, A., Farinhas, A., Treviso, M., Niculae, V., Aguiar, P., and Figueiredo, M. Sparse and continuous attention mechanisms. *Advances in Neural Information Processing Systems*, 33:20989–21001, 2020.
- Narayan, S., Cohen, S. B., and Lapata, M. Don't give me the details, just the summary. *Topic-Aware Convolutional Neural Networks for Extreme Summarization ArXiv, abs*, 1808.
- Nelson, E., Kollias, G., Das, P., Chaudhury, S., and Dan, S. Needle in the haystack for memory based large language models. arXiv preprint arXiv:2407.01437, 2024.
- Radford, A. Improving language understanding by generative pre-training. 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Ren, J., Rajbhandari, S., Aminabadi, R. Y., Ruwase, O., Yang, S., Zhang, M., Li, D., and He, Y. {Zero-offload}: Democratizing {billion-scale} model training. In 2021 USENIX Annual Technical Conference (USENIX ATC 21), pp. 551–564, 2021.
- Roberts, A., Raffel, C., Lee, K., Matena, M., Shazeer, N., Liu, P. J., Narang, S., Li, W., and Zhou, Y. Exploring the limits of transfer learning with a unified text-to-text transformer. *Google, Tech. Rep.*, 2019.
- Shazeer, N. Glu variants improve transformer. arXiv preprint arXiv:2002.05202, 2020.
- Soboleva, D., Al-Khateeb, F., Myers, R., Steeves, J. R., Hestness, J., and Dey, N. Slimpajama: A 627b token cleaned and deduplicated version of redpajama, 2023. URL https://huggingface. co/datasets/cerebras/SlimPajama-627B.
- Su, J., Lu, Y., Pan, S., Wen, B., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 2024.
- Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., and Hashimoto, T. B. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/ stanford\_alpaca, 2023a.

- Taori, R., Gulrajani, I., Zhang, T., et al. Alpaca: A strong, replicable instruction-following model. https://github.com/tatsu-lab/ stanford\_alpaca, 2023b.
- Tay, Y., Dehghani, M., Tran, V. Q., Garcia, X., Wei, J., Wang, X., Chung, H. W., Shakeri, S., Bahri, D., Schuster, T., et al. Ul2: Unifying language learning paradigms. *arXiv preprint arXiv:2205.05131*, 2022.
- Team, G., Georgiev, P., Lei, V. I., Burnell, R., Bai, L., Gulati, A., Tanzer, G., Vincent, D., Pan, Z., Wang, S., et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530, 2024.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971, 2023.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Wang, A. and Cho, K. Bert has a mouth, and it must speak: Bert as a markov random field language model. *arXiv* preprint arXiv:1902.04094, 2019.
- Xiao, G., Tian, Y., Chen, B., Han, S., and Lewis, M. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023.
- Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., et al. Qwen2.5 technical report. arXiv preprint arXiv:2412.15115, 2024.
- Yang, Z. Xlnet: Generalized autoregressive pretraining for language understanding. arXiv preprint arXiv:1906.08237, 2019.
- Ye, T., Dong, L., Xia, Y., Sun, Y., Zhu, Y., Huang, G., and Wei, F. Differential transformer. *arXiv preprint arXiv:2410.05258*, 2024.
- Zhang, B. and Sennrich, R. Root mean square layer normalization. Advances in Neural Information Processing Systems, 32, 2019.
- Zhang, P., Zeng, G., Wang, T., and Lu, W. Tinyllama: An open-source small language model. arXiv preprint arXiv:2401.02385, 2024.

# A. Experimental Details of Pre-training

### A.1. Architecture and Hyperparameters

This section outlines the pre-training hyperparameters of the MEAP model, designed to ensure efficient training and optimal performance. The sequence length is fixed at 4096 tokens, enabling the model to handle long-range dependencies while maintaining computational efficiency. The learning rate schedule includes an initial warm-up phase for the first 10% of training steps, followed by cosine decay to 10% of the initial value, allowing gradual and precise parameter adjustments. The AdamW optimizer is used with standard hyperparameters  $\beta_1 = 0.9$  and  $\beta_2 = 0.95$  to stabilize the optimization process. Learning rate bounds are set between  $4 \times 10^{-4}$  and  $4 \times 10^{-5}$  to ensure effective learning throughout training, while a weight decay of  $5 \times 10^{-2}$  helps prevent overfitting and promote generalization by penalizing excessively large weights. Complete training hyperparameters are documented in Table 11.

The model sizes and corresponding hyperparameters are shown in Table 12.

Name	Hyperparameter
optimizer	AdamW
lr schedule	cosine
clip	1.0
max learning rate	$4 \times 10^{-4}$
min learning rate	$4 \times 10^{-5}$
weight_decay	$5  imes 10^{-2}$
sequence length	4096
batch size	256
epoch	1

Table 11. Hyperparameters of training

Table 12. Hyperparameters of pretrained MEAP models. Data amount are specified in tokens.

Params	Hidden	Intermediate	Heads	Layers	Steps	Data amount
100M	768	2048	12	12	2K	$2 \mathrm{B}$
500M	1024	4864	16	24	10K	$10 \mathrm{B}$
$1.1 \mathrm{B}$	2048	5632	24	32	190K	$200 \mathrm{B}$

### A.2. Pre-training Loss of Difference Model Sizes

The loss curves of the MEAP model at various sizes, as shown in Figure 7, provide a detailed visualization of the model's performance across different scales.

#### A.3. Language Modeling Evaluation Of All Size Models for Pre-training

As shown in Table 13, we present the evaluation results of models of different scales implemented using our method. To comprehensively assess the language modeling performance of these models, we conducted a detailed analysis for each model, with particular focus on their performance at varying scales.

### A.4. Pretrained Model Evaluation Under Different Masking Rates

As shown in Table 14, we present the evaluation results of models implemented with our approach, where different mask rates are applied during training. A comprehensive and detailed analysis of the language modeling performance is conducted for each mask rate, with a focus on how varying levels of masking influence key performance metrics. This analysis elucidates the effects of mask rates on the model's ability to handle diverse linguistic tasks, highlighting any changes in accuracy as the mask rate is adjusted.



Figure 7. Overview of loss for all size pretrained-models

Benchmark	100M	500M	1.1B
ARC-Challenge	17.32	18.4	25.4
ARC-Easy	31.99	42.0	56.4
BoolQ	45.14	55.63	59.5
HellaSwag	26.82	30.77	43.4
OpenBookQA	11.41	16.40	22.6
PIQA	58.49	66.81	72.3
WinoGrande	52.09	49.57	55.3
Avg	34.75	39.94	47.85

Table 13. Results of all size MEAP pretrained models

#### A.5. Details Of Contextual Hallucination Evaluation

Here are the prompt for summarization generation, where "doc" is the original text to be summarized.

Summarize the following article: doc

We use the following prompts to let the Deepseek v3 model perform binary classification to determine whether there is hallucination in the model output compared to the human summary.

The "model output" is the output of the model, and the "predicted label" is the manually annotated label. Please compare the "model output" with the "predicted label". By comparing the two, check if the "model output" is similar. If it is similar, return 1; otherwise, return 0. An explanation of the output is required. Here is the output format I provide. Please follow it strictly!! Score: xx

Benchmark	Mask Ratio 0.15	Mask Ratio 0.05	Mask Ratio 0.1
ARC Challenge	25.4	26.11	24.3
ARC Easy	56.4	56.1	54.3
BoolQ	59.5	56.5	53.4
HellaSwag	43.4	43.69	43.85
OpenBookQA	22.6	22.0	21.8
PIQA	72.3	72.63	72.91
Winogrande	55.3	56.4	56.91
Avg	47.84	47.63	46.78

Table 14. Results of the 1.1b MEAP model under different masking rates

Table 15. Performance comparison of different masking strategies against NTP baseline.

Method	ARC-c	ARC-e	BoolQ	PIQA	HellaSwag	WinoGrande	OBQA	Average   MDQA
NTP-0.3B	18.00	37.75	58.44	62.62	28.56	50.67	13.60	40.09 0.187
Random Mask	21.84	35.44	61.25	61.04	29.50	51.46	27.40	41.13 0.218
5-Span Mask	21.42	35.61	60.40	62.08	29.81	51.07	27.60	41.14 0.168
50-Span Mask	23.46	36.20	59.54	62.84	30.43	50.99	28.00	41.64 0.189

### A.6. Details of Attention Distribution of MEAP and NTP

To validate the generality of attention changes, we conducted corresponding tests on additional examples and observed that the attention changes in these examples are consistent with the results presented in the main text. The specific changes are detailed in Table 16, Table 17, and Table 18.

Area	Content	MEAP attention	NTP attention
context before	The Great Wall of China, stretching over 13,000 miles, is one of the most impressive feats of ancient engineer- ing.	0.491	0.731
answer	Built to protect Chinese states from invasions	0.329	0.108
context after	the wall took several dynasties over 2,000 years to complete. Its im- mense length and historical signif- icance make it a popular tourist attraction today. The wall's con- struction involved countless work- ers, many of whom faced difficult conditions.	0.078	0.80
query	question:What was the purpose of the Great Wall of China?	0.067	0.070
eos	<s></s>	0.069	0.071

Table 16. Attention change of example 1

Table 17. Attention change of example 2

Area	Content	MEAP attention	NTP attention
context before	In the early 20th century, Albert Ein- stein introduced his theory of rela- tivity, which changed the way we understand space, time, and gravity.	0.435	0.694
answer	His famous equation, E=mc <sup>2</sup>	0.386	0.115
context after	shows the relationship between en- ergy and mass. Einstein's ideas rev- olutionized physics, and his work led to the development of technolo- gies like GPS and nuclear energy. Despite facing initial skepticism, his theories were eventually proven through experiments and observa- tions, earning him a Nobel Prize in Physics in 1921.	0.066	0.074
query	question:What famous equation did Albert Einstein create?	0.057	0.060
eos	<s></s>	0.055	0.057

Table 18. Attention change of example 3

Area	Content	MEAP attention	NTP attention
context before	At the center of Rome, the Colos- seum rises as a magnificent testa- ment to ancient Roman architecture, symbolizing the grandeur of the Ro- man Empire.	0.579	0.748
answer	Constructed between 70 and 80 AD under the emperors Vespasian and Titus,	0.219	0.065
context after	it was used for gladiatorial contests and public spectacles. Once a sym- bol of Roman power, the Colosseum has weathered centuries of change but remains one of the most iconic structures in the world. Tourists flock to see it every year, making it one of the most photographed mon- uments in history.	0.071	0.067
query	question: Who built the Colosseum?	0.063	0.050
eos	<s></s>	0.068	0.069

# **B.** Details of Fine-tuning Experiments

# **B.1.** Architecture and Hyperparameters

This section details the MEAP fine-tuning hyperparameters for the Llama3 model. The maximum sequence length is 4096 tokens, optimizing long-range dependencies and efficiency. The batch size is 512, and the learning rate schedule includes a warm-up for the first 10% of training steps. The AdamW optimizer is used with  $\beta_1 = 0.9$  and  $\beta_2 = 0.95$ , and the learning rate is set to  $2 \times 10^{-5}$ .

Table 19. MEAP fine-tuning hyperparameters of Llama3 model

Name	Hyperparameter	
optimizer	AdamW	
lr schedule	cosine	
clip	1.0	
learning rate	$2 \times 10^{-5}$	
weight_decay	$5 \times 10^{-2}$	
maximum sequence length	4096	
batch size	512	