

PEAR: POSITION-EMBEDDING-AGNOSTIC ATTENTION RE-WEIGHTING ENHANCES RETRIEVAL-AUGMENTED GENERATION WITH ZERO INFERENCE OVERHEAD

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) enhanced with retrieval-augmented generation (RAG) have introduced a new paradigm for web search. However, the limited context awareness of LLMs degrades their performance on RAG tasks. Existing methods to enhance context awareness are often inefficient, incurring time or memory overhead during inference, and many are tailored to specific position embeddings. In this paper, we propose **Position-Embedding-Agnostic attention Re-weighting (PEAR)**, which enhances the context awareness of LLMs with zero inference overhead. Specifically, on a proxy task focused on context copying, we first detect heads which suppress the models’ context awareness, thereby diminishing RAG performance. To weaken the impact of these heads, we re-weight their outputs with learnable coefficients. The LLM (with frozen parameters) is optimized by adjusting these coefficients to minimize loss on the proxy task. During inference, the optimized coefficients are fixed to re-weight these heads, regardless of the specific task at hand. Our proposed *PEAR* offers two major advantages over previous approaches: (1) It introduces zero additional inference overhead in terms of memory usage or inference time, while outperforming competitive baselines in accuracy and efficiency across various RAG tasks. (2) It is independent of position embedding algorithms, ensuring broader applicability.

1 INTRODUCTION

Retrieval-augmented generation (RAG, (Lewis et al., 2021)) is widely utilized to enhance large language models (LLMs) on tasks like question answering. Typically, an RAG framework retrieves documents related to users’ question from external knowledge bases or web pages, and then arranges them in the LLMs’ context as the references to form answers. This LLM-based question-answering paradigm has given rise to a promising web search paradigm (Microsoft, 2023; OpenAI, 2024).

Recent research demonstrated LLMs’ limitations on context awareness, especially when processing long context. These limitations in LLMs’ context awareness challenge the effectiveness and robustness of RAG frameworks. For instance, Liu et al. (2023) found that when performing in-context retrieval tasks, LLMs exhibit insensitivity to information located in the middle of the context, a phenomenon referred to as “lost in the middle.” Chen et al. (2024) identified a mathematical property in rotary position embedding (RoPE, (Su et al., 2023)) that results in LLMs assigning less attention to specific contextual positions, leading to varying context awareness throughout the entire context.

Existing approaches to enhancing LLMs’ context awareness are inefficient in terms of memory and time cost. Some works (Peysakhovich & Lerer, 2023) segment and re-arrange the input context, with the assumption that placing important information in positions the model attends well can improve RAG’s effectiveness. This method incur additional inference time costs, negatively affecting user experience, as it requires multiple forward passes to obtain attention weights for guiding segment rearrangement. Another group of studies modifies the model’s working mechanism, specifically by employing a set of position embeddings to adjust the attention preferences of attention heads. While these methods are input-agnostic, they also lack efficiency due to “parallelable” forward passes (Chen et al., 2024) although without increased time cost, disrupting the parallelism of multi-head attention (resulting in increased time cost) to achieve low memory cost (Lin et al.,

Table 1: An example input for the proxy task, where unique letters representing distinct tokens, with $n = 4$. For example, at position $n + i = 5$ (with $i = 1$), when an LLM receives “ABCD A” as input, it is likely to output “B.” This happens because the last occurrence of “A” in the preceding context is followed by “B.” If a head suppresses copying “B” from position $i + 1 = 2$ to position $n + i = 5$, it could negatively impact RAG performance.

Input Sequence	A	B	C	D	A	B	C	D
Position Index	1	2	3	4	5	6	7	8

2024), or still requiring “non-parallelable” multiple forward passes (Zhang et al., 2024) which is similar to (Peysakhovich & Lerer, 2023). Moreover, these studies are mainly designed for RoPE and face challenges in generalizing to other position embedding algorithms, limiting their broader applications.

In this paper, we introduce **Position-Embedding-Agnostic** attention head **Re-weighting** (PEAR), which unleashes the context awareness of LLMs, thereby improving their RAG performance. *PEAR* achieves zero additional overhead in memory usage and inference time. Our motivation relies on the following facts:

1. Prior research (McDougall et al., 2023; Lv et al., 2024a) detected some attention heads decrease the language model’s prediction confidence by suppressing the flow of contextual information to the final position within the context, where the output is to be generated.
2. This suppression negatively impacts LLMs’ context awareness, particularly abilities in in-context retrieval and context integration, which are crucial for effective RAG.

As a result, we contend that such suppression mechanism in LLMs can be safely¹ weakened in RAG scenarios, thereby improving the performance of RAG. Our proposed *PEAR* includes two stages:

In the first stage, we discover attention heads negatively affect performance on a proxy task. The proxy task involves feeding the model a random token sequence of length $2n$, which consists of a duplicated sub-sequence of length n . Table 1 illustrates an example input. At position $n + i$, the model typically predicts the token from position $i + 1$, as the natural continuation for a semantically meaningless context is to copy the existing in-context token pattern (Lv et al., 2024b). The negatively impactful attention heads are discovered using the path patching technique (Wang et al., 2023). Since this proxy task is free from semantic bias and requires both in-context retrieval and generation based on the context—fundamental capacities for RAG, we refer to discovered heads as RAG-suppression heads².

In the second stage, we weaken detected RAG-suppression heads by re-weighting their outputs using learnable coefficients. These coefficients are optimized by minimizing the LLM’s loss on the proxy task, with the objective of next-token prediction in a supervised fine-tuning process (loss is computed only for the second half of the random-token sequence). During the optimization, the original LLM parameters are frozen. Intuitively, most of the learned coefficients are optimized to values less than one, reducing the relative weight of these heads compared to others in the same layer when multi-head outputs are aggregated. Consequently, their influence during the forward pass are weakened. Once optimized, the coefficients remain fixed and are agnostic to downstream RAG tasks.

PEAR achieves three-fold contributions:

1. *PEAR* introduces zero inference overhead in terms of both memory usage and inference time—an advantage not achieved by competitive baselines. Across a wide range of RAG tasks, *PEAR* surpasses previous works in both efficiency and accuracy.
2. *PEAR* is independent of specific position embedding algorithms, making it broadly applicable. We demonstrate that *PEAR* enhances the RAG performance of various LLMs using

¹“Safely” means that the parametric knowledge and fundamental capabilities remain unaffected. Detailed experimental results are presented in Section 5.4.

²We do not imply that these heads hinder RAG through the same mechanisms (discussed in Section 4.1).

108 distinct position embeddings (e.g., learnable embeddings (Zhang et al., 2022), RoPE (Su
109 et al., 2023), and Alibi (Press et al., 2022)).
110

111 2 RELATED WORKS 112

113
114 In this section, we discuss two research areas closely related to this paper: enhancements to LLMs’
115 context awareness and studies on mechanistic interpretability.
116

117 2.1 CONTEXT AWARENESS ENHANCEMENT 118

119 Many studies highlighted limitations in LLMs’ context awareness. For example, Lu et al. (2022)
120 found that the order of in-context learning (ICL) demonstrations significantly affects ICL accuracy.
121 Liu et al. (2023) demonstrated that LLMs exhibit stronger awareness of content at the beginning
122 and end of context but weaker awareness in the middle, a phenomenon termed “lost in the middle.”
123 Chen et al. (2024) proposed that LLMs’ context awareness fluctuates across token positions due to
124 mathematical properties in position embeddings. These challenges impact applications like RAG
125 that rely on robust context awareness.

126 Several approaches have been proposed to tackle these issues. However, existing methods often
127 come at the cost of increased inference time or memory overhead. Attention Buckets (AB, (Chen
128 et al., 2024)) enhances context awareness by integrating positional information from a set of various
129 RoPE angles, but it incurs significant inference overhead due to multiple parallel forward passes,
130 leading to increased memory usage. Ms-PoE (Zhang et al., 2024) calculates distinct re-scaling fac-
131 tors for each attention head, requiring multiple non-parallel forward passes that introduce noticeable
132 time delays. MoICE (Lin et al., 2024) builds on Attention Buckets by employing a Mixture-of-
133 Experts (MoE, (Shazeer et al., 2017)) approach, treating each RoPE embedding as a unique in-
134 context expert, thereby limiting extra attention computations to within each layer rather than across
135 the entire forward pass.

136 Our proposed method, *PEAR*, enhances context awareness by weakening the RAG-suppression
137 heads during the forward pass. It introduces no additional modules or extra forward passes, result-
138 ing in *zero* additional overhead in memory usage and inference time. Additionally, *PEAR* operates
139 independently of position embedding algorithms, which is applicable to more LLMs compared to
140 existing approaches.

141 2.2 MECHANISTIC INTERPRETABILITY 142

143 Investigating the role of a specific head during a forward pass is one of key focuses in mechanistic
144 interpretability research. Wang et al. (2023) reported that in GPT-2 small (Radford et al., 2019),
145 the 7th head in the 10th layer, termed the “negative head,” significantly hinders answer copying
146 from context. McDougall et al. (2023) comprehensively studied this head, suggesting it functions
147 as a self-repair mechanism to prevent overconfident outputs. Lv et al. (2024a) found that negative
148 heads exist across various LLMs, employing different mechanisms to mitigate overconfidence, such
149 as generating counteracting vectors or introducing high-frequency tokens’ information. This paper
150 does not examine what specific mechanisms the heads employ to suppress RAG performance but
151 instead aims to discover and suppress heads negatively impactful across general RAG tasks.

152 Yu et al. (2023) detected two types of heads in Transformer-based language models during counter-
153 factual task execution (where counter-factual knowledge is provided in the context): memory heads,
154 which prefer to use stored knowledge, and in-context heads, which prefer to use facts in the context.
155 However, when re-weighting these heads, only reducing the weight of memory heads successfully
156 enhances the model’s preference of using contextual knowledge, while enhancing or mitigating in-
157 context heads does not bring much influence. Moreover, existing results only demonstrate effective-
158 ness of these heads in the “country-capital” task (i.e., prompting the model to answer the capital city
159 of the given country). Additionally, there is no evidence suggesting that re-weighting these heads
160 improves the comprehensive context awareness of LLMs. We owe these failures to their head detec-
161 tion method, and the same re-weighting value applied for all the heads of the same type. By contrast,
in our proposed *PEAR*, each individual head is re-weighted by a specific learnable coefficient, which
is optimized through a proxy task independent of downstream tasks.

Wu et al. (2024) also relates to head re-weighting. They reported that diminishing induction heads (Olsson et al., 2022), which is known for handling context copying, can impair LLMs’ long-context performance. However, they did not offer a solution to improve context awareness.

3 PRELIMINARIES: DISCOVERY OF INFLUENTIAL ATTENTION HEADS

For a particular task, research has shown that only a sparse sub-network is activated during the forward pass in Transformer language models (Wang et al., 2023; Merullo et al., 2024; Gong et al., 2024). Such a sub-network is referred to as a circuit (Olah et al., 2020). Discovering circuits provides interpretability into the working mechanisms of language models and offers insights for model enhancement.

The primary method for circuit discovery is based on causal mediation analysis. The core idea is to view the forward computation graph as a causal graph, where the output of one module serves as the input for the next. In such a case, if the output of a module is changed, the computation of subsequent modules in the causal graph is also affected, as their inputs change.

In this paper, we primarily focus on analyzing the working mechanism of attention heads in language models. We briefly introduce a paradigm from a series of works (Wang et al., 2023; Zhang & Nanda, 2024; Wang et al., 2024) that discovers which attention heads are crucial for processing an input sequence X of length n . Suppose the language model consists of L layers, with H attention heads per layer. Let $A^{(l,h)}$ denote the h -th attention head in the l -th layer, and let its outputs be denoted by $\mathbf{a}^{(l,h)} \in \mathbb{R}^{n \times d}$. We use $\mathbf{a}_i^{(l,h)} \in \mathbb{R}^d$, where $1 \leq i \leq n$, to represent the output at position i . The discovery paradigm typically includes three steps, as illustrated in Figure 1:

1. In the *normal run*, with an input sequence X (e.g., X = “The capital of France is”), $\mathbf{a}^{(l,h)}$ for every attention head are recorded.
2. In the *perturbed run*, the forward computation runs using the same input sequence X , but with some mediation. This mediation either changes the discrete input tokens within X by substituting specific keywords (e.g., replacing “France” with “England”), or corrupts the hidden states by adding noise. The modified $\tilde{\mathbf{a}}^{(l,h)}$ for each attention head are then recorded.
3. We conduct an intervention on a particular head $A^{(l,h)}$ at a specific position i (e.g., the country token position in above examples) in the normal run by substituting its outputs with $\tilde{\mathbf{a}}_i^{(l,h)}$. The subsequent activations in the computational graph are then recomputed (these recomputed activations are denoted as $\tilde{\mathbf{a}}$ in the figure). If the final output of the language model as the intervention expects (e.g., the predicted token changes from “Paris” to “London”), the head $A^{(l,h)}$ is considered to have a positive influence on the processing of sequence X .

This overview outlines a simplified discovery paradigm; detailed measurements of intervention impact are tailored to specific experimental needs.

4 METHODOLOGY

In this section, we provide a detailed introduction to our proposed method, *PEAR*, which is executed in two stages: (1) discovering RAG-suppression heads and (2) re-weighting coefficient learning. The first stage discovers attention heads that have a negative impact on general RAG tasks based

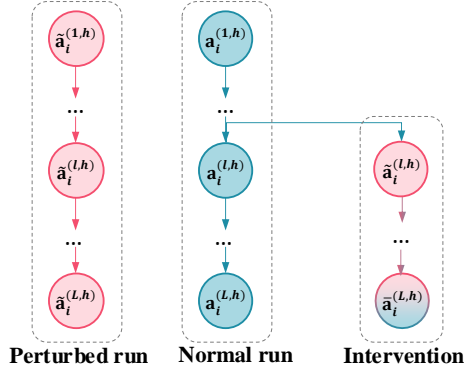


Figure 1: An illustration of causal mediation methods for circuit discovery.

on circuit discovery for a proxy task. In the second stage, we optimize learnable coefficients to re-weight the outputs of the discovered heads, aiming to mitigate their RAG-suppression effect. These coefficients remain fixed during inference, irrespective of the specific input. Figure 2 demonstrates the overview of *PEAR*.

4.1 DISCOVERY OF RAG-SUPPRESSION HEADS

We set up a proxy task and use this task as input for circuit discovery algorithms to discover influential attention heads that hamper LLMs’ performance on general RAG tasks.

Task Input For each input sample, we create a sequence of length n , denoted as $\{x_1, \dots, x_n\}$, where each x_i is a randomly sampled token from the vocabulary. This sequence is repeated to form an input sample $X = \{x_1, \dots, x_{2n}\}$, with $x_i = x_{i+n}$ for $i \in [1, n]$. Research has shown that, in semantically meaningless contexts, models tend to check if the last few tokens in the sequence appeared previously and copy the suffix of their last appearance as the output (Olsson et al., 2022; Lv et al., 2024b). We consider an arbitrary LLM to successfully perform the proxy task when, at position $n + i + 1$, the token with the highest output logits is x_i . Table 1 shows an example input.

This proxy task exhibits two key characteristics that facilitate the effective discovery of RAG-related heads:

1. Completing the proxy task requires LLM capabilities essential for a robust RAG framework, such as in-context retrieval and generation based on context, making it suitable for discovering RAG-related attention heads.
2. The random token composition in X ensures semantically meaningless input, minimizing knowledge bias and thereby enabling the discovered attention heads to have general RAG-related functions, independent of specific downstream tasks.

Head discovery We previously outlined the head discovery algorithm in Section 3. Here, we provide additional practical details for the first stage of *PEAR*.

1. During the *normal run*, the input sequences X are constructed as above described, with a length of $2n$.
2. In the *perturbed run*, we do not modify the input or hidden states; instead, we average the outputs of each attention head along the sequence dimension and record the resulting mean vectors.
3. We focus on detecting changes in logits at position $2n$, where the model is expected to copy the token from position $n - 1$. Consequently, we intervene at $a_{n-1}^{(l,h)}$ by replacing it with the saved mean vectors.
4. Our intervention measurements are based on the logits difference, defined as:

$$\Delta\pi^{(l,h)} = \frac{\tilde{\pi}_{2n}^{(l,h)}[x_{n-1}]}{\pi_{2n}[x_{n-1}]} - 1, \tag{1}$$

where π_{2n} represents the final logits at position $2n$ during the *normal run*, and $[x_{n-1}]$ denotes selecting the value of the token x_{n-1} from the logits. $\tilde{\pi}^{(l,h)}$ indicates the logits after intervention on $A^{(l,h)}$. We contend that a higher value of this metric suggests a stronger suppression effect from $A^{(l,h)}$.

5. For an arbitrary LLM, we repeat the proxy task multiple times with varying values of n to mitigate bias in context length. The final metric score for each head is the average of the results from these repeated experiments. The detailed setup is provided in Section 5.1.
6. Based on the final metric scores, we identify the heads with the top- K most negative influence on the proxy task as a set \mathcal{S} , defined as:

$$\mathcal{S} = \{A^{(l,h)} | A^{(l,h)} \text{ has one of the top-}K \text{ values of } \Delta\pi^{(l,h)}\}.$$

These heads are collectively referred to as RAG-suppression heads.

Notably, we do not suggest that the heads we discovered suppressing RAG tasks operate through the same mechanisms. In Section 5.6, we demonstrate that these heads may serve various functions, such as copy suppression (McDougall et al., 2023), incorporating high-frequency token information (Lv et al., 2024a), or influencing the behavior of other heads to indirectly affect outputs (Wang et al., 2023). Analyzing the specific mechanism for each head is not the focus of this paper and is left for future works.

4.2 RE-WEIGHTING COEFFICIENT LEARNING

Optimization In standard multi-head attention mechanisms, the outputs of all attention heads are aggregated with equal weighting. We propose that re-weighting these relative aggregation weights to values less than 1 can mitigate the RAG-suppression effect from our discovered heads. To implement this, we modify the forward computation by multiplying the output of each head, $A^{(l,h)}$, in the set \mathcal{S} by a learnable scalar, $\tau^{(l,h)}$, referred to as the re-weighting coefficient. The modified output for each head is:

$$\mathbf{a}^{(l,h)} = \tau^{(l,h)} * \mathbf{a}^{(l,h)}, \text{ for each } A^{(l,h)} \in \mathcal{S}. \tag{2}$$

To optimize these re-weighting coefficients for RAG-suppression heads, we freeze the original parameters of the LLM and train only the re-weighting coefficients to minimize the loss on a proxy task. Importantly, the loss is calculated only over the latter half of the sequence, optimizing the coefficients to enhance in-context retrieval capacities rather than predicting the next token. Formally, our adopted loss can be written as:

$$\mathcal{L} = - \sum_{i=n}^{2n-1} \log p(x_{i+1}|x_{1:i}) \tag{3}$$

Figure 2 illustrates a re-weighting process during optimization. Notably, the re-weighting process shown in this figure adds extra multiplication operations in a forward pass. In practice, when coefficient learning ends, we re-scale $W_O^{(l,h)}$ (the output projection matrix in head $A^{(l,h)}$) by $\tau^{(l,h)}$, which is equivalent to Eq. 2 and does not add any extra computation during inference.

Inference on Downstream Tasks We highlight several points regarding the inference process of our proposed *PEAR* on downstream RAG tasks:

1. In downstream RAG tasks, the re-weighting coefficients are task-independent and remain fixed.
2. RAG-suppression heads are optimized once for each LLM via the proxy task. For a new RAG task, head discovery and coefficient learning do not need to be repeated.

In theory, our approach, *PEAR*, introduces *zero* additional overhead during inference on downstream RAG tasks, as it does not incorporate extra computational modules; instead, it only adjusts the aggregation weights of specific heads. In practice, however, the re-weighting process involves an additional multiplication, which results in zero additional overhead in inference time and memory usage. Additionally, the learning of re-weighting coefficients is independent of the LLM architecture, thus making our method compatible with various position embedding algorithms.

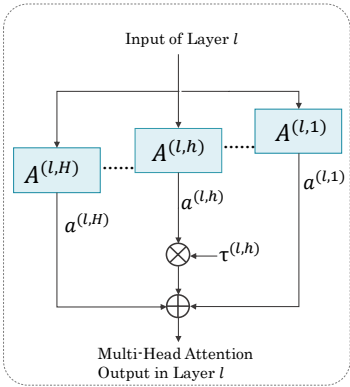


Figure 2: Suppose in layer l , $A^{(l,h)}$ is discovered as a RAG-suppression head. *PEAR* re-weights its output with a learnable scalar $\tau^{(l,h)}$.

Table 2: Discovered RAG-suppression heads in Llama2-Chat-7B-4k, OPT-6.7B-2k and Baichuan-13B-chat-4k, respectively.

Model Name	(l, h) for discovered $A^{(l,h)}$
Llama2-7B-chat-4k	(26, 28), (11, 6), (14, 15), (30, 9), (18, 9), (15, 10), (13, 9), (12, 10), (15, 14), (10, 18), (15, 25), (19, 15), (29, 15), (14, 0), (10, 2), (31, 17), (8, 22), (17, 0), (20, 26), (9, 13), (13, 14), (7, 9), (10, 1), (15, 12), (11, 9), (15, 7), (9, 16), (26, 9), (28, 22), (15, 2)
OPT-6.7B-2k	(29, 19), (0, 22), (0, 6), (26, 16), (26, 15), (30, 19), (0, 18), (23, 30), (0, 10), (31, 31), (28, 6), (30, 30), (21, 27), (0, 17), (31, 25), (12, 23), (22, 16), (0, 0), (23, 0), (0, 1), (24, 31), (23, 8)
Baichuan-13B-chat-4k	(26, 22), (33, 25), (28, 26), (32, 13), (23, 20), (25, 24), (19, 20), (38, 16), (22, 21), (21, 12), (3, 24), (39, 39), (20, 27), (37, 21), (0, 32), (24, 39), (39, 28), (39, 20), (27, 24), (2, 20), (36, 10)




Figure 3: Heatmaps of $\Delta\pi$ scores for each head across three LLMs ($n = 10$).

5 EXPERIMENTS

5.1 SETUP

In this section, we introduce the LLMs we used for experiments, the baseline methods for enhancing context awareness, the setups for the proxy tasks, hyperparameters for learning re-weighting coefficients.

Models and baselines We conducted experiments with three LLMs, each employing a different position embedding algorithm: Llama2-7B-chat-4k (Touvron et al., 2023) using RoPE (Su et al., 2023), OPT-6.7B-2k (Zhang et al., 2022) using learnable position embeddings, and Baichuan-13B-4k (Baichuan, 2023) using Alibi position embeddings (Press et al., 2022).

We also compared several competitive baseline methods for enhancing LLMs’ context awareness, including Attention Buckets (AB, (Chen et al., 2024)), Ms-PoE (Zhang et al., 2024), and MoICE (Lin et al., 2024). Details on these methods can be found in Section 2.1.

Detailed setups of proxy task For head discovery, we constructed 200 task samples. In the case of the Llama and OPT models, we repeat the discovery process four times with varying values of n : 10, 15, 25, 50. For the Baichuan model, the n values are 10, 20, 50, 80. We found that each model has a group of heads with significantly large $\Delta\pi$ values, leading us to select the K values based on the observed group sizes: 30, 22, and 21, respectively. Table 2 presents the discovered RAG-suppression heads for the LLMs under study. Figure 3 illustrates $\Delta\pi$ values when $n = 10$ for each model. Further detailed results can be found in Appendix A.

For the re-weighting coefficient learning stage, we constructed 500 task samples, setting n to 50 for all models.

Hyperparameters for re-weighting coefficient learning We employed the AdamW optimizer with a learning rate of 0.005 and parameters $(\beta_1, \beta_2) = (0.9, 0.999)$. τ are initialized at 1.0. Training was performed for a single epoch using BF16 precision on an A100-PCIE-40GB GPU.

Table 3: Performance comparison of Llama2-7B-chat-4k and its enhancements across three RAG tasks.

Method	2WikiMultiHopQA \uparrow	MuSiQue \uparrow	Qasper \uparrow	Avg. \uparrow
Llama2-7B-chat-4k	29.50	6.50	17.00	17.67
+ Ms-PoE (Zhang et al., 2024)	27.50	9.00	18.00	18.17
+ AB (Chen et al., 2024)	31.00	11.00	16.50	19.50
+ MoICE (Lin et al., 2024)	30.00	10.00	15.50	18.50
+ <i>PEAR</i> (Ours)	35.00	8.50	18.00	20.50

Table 4: Practical inference time (in seconds) and GPU memory cost (in GB) per test sample for different methods. For a fair comparison, Flash-Attention (Dao, 2024) was not applied. The experiments were conducted on a single H800-80G GPU.

Method	2WikiMultiHopQA \downarrow	MuSiQue \downarrow	Qasper \downarrow	Avg. \downarrow
Llama2-7B-chat-4k	0.63/31.33	0.70/31.33	1.23/31.33	0.88/31.33
+ Ms-PoE (Zhang et al., 2024)	0.95 (+0.32) / 39.21 (+7.88)	1.11(+0.41) / 39.21 (+7.88)	1.84 (+0.61) / 34.59 (+3.26)	1.30 (+0.42) / 37.67 (+6.34)
+ AB (Chen et al., 2024)	2.50(+1.87) / 66.19(+34.86)	2.70(+2.00) / 66.19 (+34.86)	5.67(+4.44) / 66.34(+35.01)	3.62(+2.74) / 66.24(+34.91)
+ MoICE (Lin et al., 2024)	2.91(+2.28) / 79.13(+47.80)	3.06(+2.36) / 79.12(+47.79)	5.86(+4.63) / 79.10(+47.77)	3.94(+3.06) / 79.12(+47.79)
+ <i>PEAR</i> (Ours)	0.63 (+0.00) / 31.33 (+0.00)	0.70 (+0.00) / 31.33 (+0.00)	1.23 (+0.00) / 31.33 (+0.00)	0.88 (+0.00) / 31.33 (+0.00)

5.2 COMPARISON WITH BASELINES ON RAG TASKS

We compare *PEAR* against various baselines on RAG tasks we constructed using three datasets: 2WikiMultiHopQA (Ho et al., 2020), MuSiQue (Trivedi et al., 2022), and Qasper (Dasigi et al., 2021). The first two datasets require the model to answer questions based on multiple documents, while the third focuses on questions related to NLP research papers, formulated and answered by NLP researchers. We truncate the context to 4,000 tokens for the first two datasets; the third dataset has an average context length of 3,619 tokens.

Our experiments are conducted with Llama2-7B-chat-4k, as the baselines are tailored specifically for RoPE. We evaluate the models’ performance using exact match scores, with results reported in Table 3. Notably, our method achieves the highest average improvement across all three tasks. Although *PEAR* does not achieve the top performance on the MuSiQue task, it outperforms the original model by a large margin.

Additionally, we present inference time and memory costs for these datasets in Table 4. *PEAR* does not increase GPU memory usage and inference time costs. This makes it significantly more efficient than other enhancement methods.

These experiments underscore the effectiveness and efficiency of *PEAR* in enhancing LLMs for RAG tasks.

5.3 APPLICABILITY TO LLMs USING VARIOUS POSITION EMBEDDINGS

In this section, we demonstrate the applicability of *PEAR* to LLMs utilizing different position embeddings. We conduct a multi-document question-answering (MDQA) experiment based on data from (Liu et al., 2023). Following Liu et al. (2023); Chen et al. (2024), we position the gold document (i.e., the document contains the ground truth answer) at various contextual positions to evaluate the robustness of a context-awareness enhancement method. In our experiments, we set the maximum document count to 10 and assess the question-answering accuracy when the gold document is placed as the 1st, 3rd, 5th, 7th, and 10th document, respectively. Since baseline methods are not compatible with the OPT and Baichuan models, we compare *PEAR* only with the original models. The results are presented in Table 5.

5.4 *PEAR* DOES NOT DIMINISH KNOWLEDGE CAPABILITIES IN LLMs

Previous research (Geva et al., 2023; Lv et al., 2024a) has shown that certain attention heads store or play a crucial role in eliciting parametric knowledge. This raises the question of whether *PEAR*

Table 5: Experimental results on the MDQA task show that *PEAR* achieves the highest accuracy in 24 out of 25 comparisons across three LLMs, demonstrating its broad applicability to various position embeddings and its robustness in enhancing awareness to different contextual positions.

Position Embedding	Method	Gold Document Position					Avg.
		1	3	5	7	10	
RoPE	Llama2-7B-chat-4k	64.14	65.95	64.97	62.67	67.53	65.05
	+ Ms-PoE (Zhang et al., 2024)	66.06	64.29	63.99	62.22	64.75	64.34
	+ AB (Chen et al., 2024)	66.36	66.14	65.25	63.20	64.93	65.18
	+ MolCE (Lin et al., 2024)	65.50	66.33	65.61	64.11	65.84	65.48
	+ <i>PEAR</i> (Ours)	62.71	67.01	68.32	66.44	69.57	66.81
Learnable Embeddings	OPT-6.7B-2k	19.07	15.45	17.03	16.54	22.61	18.14
	+ <i>PEAR</i> (Ours)	20.23	17.18	17.60	17.22	22.87	19.02
Alibi	Baichuan-13B-chat-4k	12.28	13.45	11.98	11.04	12.96	12.34
	+ <i>PEAR</i> (Ours)	14.16	14.84	13.67	12.77	13.94	13.88

enhances context awareness of LLMs at the expense of their ability to utilize this parametric knowledge. To investigate this, we evaluated a *PEAR*-enhanced Llama2-7B-chat model using the MMLU benchmark (Hendrycks et al., 2021), and the results are presented in Table 6. The performance of the enhanced Llama2-7B-chat and the original Llama2-7B-chat did not show a significant difference. Consequently, we argue that *PEAR*, through its precise head discovery and effective re-weighting learning approaches, does not compromise the knowledge capabilities of LLMs.

Table 6: Results on the MMLU benchmark showing that *PEAR* does not enhance LLMs’ context capacities at the expense of knowledge ability.

Model	Humanities	Social Science	STEM	Other	Avg.
Llama2-7B-chat-4k	42.55	52.29	37.14	52.47	45.81
+ <i>PEAR</i> (Ours)	42.06	52.03	36.61	52.19	45.41

5.5 ANALYSIS: THE EFFECT OF K

While we have demonstrated the effectiveness of *PEAR* from various angles, a key point for discussion is the role of K , representing the number of heads to re-weight. Using Llama2-7B-chat as a case study, we vary K and observe its impact on *PEAR*’s performance. Table 7 presents results on RAG tasks, while Table 8 details analysis for MDQA tasks. The findings indicate that *PEAR* performs optimally when K matches the inherent threshold of the model, i.e., the number of heads with a significantly higher $\Delta\pi$ than others. Re-weighting fewer heads fails to fully alleviate the suppression from RAG-suppression heads, while exceeding this optimal number can harm the performance of non-RAG-suppression heads, ultimately diminishing overall effectiveness.

Table 7: The experiment results on the question answering task with ablation settings, which show that our control over the number of suppression heads is effective.

Method	2WikiMultiHopQA	MuSiQue	Qasper	Avg.
Llama2-7B-chat-4k	29.50	6.50	17.00	17.67
+ <i>PEAR</i> ($K=10$)	33.00	8.00	16.00	19.00
+ <i>PEAR</i> ($K=20$)	33.50	8.50	16.50	19.50
+ <i>PEAR</i> ($K=30$)	35.00	8.50	18.00	20.50
+ <i>PEAR</i> ($K=40$)	32.50	8.00	17.00	19.17

Table 8: The experiment results on the MDQA task with ablation settings, which show that our control over the number of suppression heads is effective.

Method	Gold Document Position					Avg.
	1	3	5	7	10	
Llama2-7B-chat-4k	64.14	65.95	64.97	62.67	67.53	65.05
+ <i>PEAR</i> ($K=10$)	63.43	66.26	66.82	64.67	67.76	65.79
+ <i>PEAR</i> ($K=20$)	63.88	66.29	66.44	65.65	68.51	66.15
+ <i>PEAR</i> ($K=30$)	62.71	67.01	68.32	66.44	69.57	66.81
+ <i>PEAR</i> ($K=40$)	62.90	66.00	67.16	66.59	68.40	66.21

5.6 ANALYSIS: THE VALUE OF τ

Using Llama2-7B-chat as an example, we present the learned coefficients of *PEAR* in Figure 4, with heads ranked by their $\Delta\pi$ scores. Intuitively, most heads are optimized to values less than one, which reduces their relative weight compared to other heads within the same layer when multi-head outputs are aggregated. Due to BF16 training precision, many τ s are optimized to the same value. However, using FP32 precision for training did not significantly impact the results.

Notably, τ s for 9 heads, which have relatively low $\Delta\pi$, are greater than one. We do not attribute this to the precision of the discovery process, as constraining the re-weighting coefficients to be less than one led to suboptimal performance. Thus, a plausible explanation is that RAG suppression is a complex, cooperative effect involving multiple heads, each with distinct working mechanisms, as discussed in Section 4.1.

6 CONCLUSION

In this paper, we introduce *PEAR*, a position-embedding-agnostic method designed to enhance the performance of LLMs on RAG tasks with zero inference overhead. Our method not only outperforms competitive baselines in both effectiveness and efficiency but also demonstrates broad applicability across various LLMs. We also presented that *PEAR* improves context awareness in LLMs without compromising their inherent knowledge capabilities. These benefits make *PEAR* a promising approach for a wide range of applications that require robust context abilities, such as in-context learning and strict instruction following, which we leave for future research.




Figure 4: The learned coefficients of *PEAR* (on Llama2-7B-chat and $K = 40$).

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

REFERENCES

- 540
541
542 Baichuan. A 13b large language model developed by baichuan intelligent technology,
543 2023. URL [https://github.com/baichuan-inc/Baichuan-13B/blob/main/](https://github.com/baichuan-inc/Baichuan-13B/blob/main/README_EN.md)
544 [README_EN.md](https://github.com/baichuan-inc/Baichuan-13B/blob/main/README_EN.md).
- 545 Yuhan Chen, Ang Lv, Ting-En Lin, Changyu Chen, Yuchuan Wu, Fei Huang, Yongbin Li, and Rui
546 Yan. Fortify the shortest stave in attention: Enhancing context awareness of large language mod-
547 els for effective tool use. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings*
548 *of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*
549 *Papers)*, pp. 11160–11174, Bangkok, Thailand, August 2024. Association for Computational Lin-
550 guistics. URL <https://aclanthology.org/2024.acl-long.601>.
- 551 Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *Inter-*
552 *national Conference on Learning Representations (ICLR)*, 2024.
- 553 Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. A dataset
554 of information-seeking questions and answers anchored in research papers, 2021. URL [https://](https://arxiv.org/abs/2105.03011)
555 arxiv.org/abs/2105.03011.
- 556 Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual
557 associations in auto-regressive language models. In Houda Bouamor, Juan Pino, and Kalika Bali
558 (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Pro-*
559 *cessing*, pp. 12216–12235, Singapore, December 2023. Association for Computational Linguis-
560 tics. doi: 10.18653/v1/2023.emnlp-main.751. URL [https://aclanthology.org/2023.](https://aclanthology.org/2023.emnlp-main.751)
561 [emnlp-main.751](https://aclanthology.org/2023.emnlp-main.751).
- 562
563 Zhuocheng Gong, Ang Lv, Jian Guan, Junxi Yan, Wei Wu, Huishuai Zhang, Minlie Huang, Dongyan
564 Zhao, and Rui Yan. Mixture-of-modules: Reinventing transformers as dynamic assemblies of
565 modules, 2024. URL <https://arxiv.org/abs/2407.06677>.
- 566
567 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Ja-
568 cob Steinhardt. Measuring massive multitask language understanding, 2021. URL [https://](https://arxiv.org/abs/2009.03300)
569 arxiv.org/abs/2009.03300.
- 570 Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-
571 hop QA dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th*
572 *International Conference on Computational Linguistics*, pp. 6609–6625, Barcelona, Spain (On-
573 line), December 2020. International Committee on Computational Linguistics. URL [https://](https://www.aclweb.org/anthology/2020.coling-main.580)
574 www.aclweb.org/anthology/2020.coling-main.580.
- 575 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,
576 Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe
577 Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2021. URL [https://](https://arxiv.org/abs/2005.11401)
578 arxiv.org/abs/2005.11401.
- 579
580 Hongzhan Lin, Ang Lv, Yuhan Chen, Chen Zhu, Yang Song, Hengshu Zhu, and Rui Yan. Mixture of
581 in-context experts enhance llms’ long context awareness, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2406.19598)
582 [abs/2406.19598](https://arxiv.org/abs/2406.19598).
- 583 Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and
584 Percy Liang. Lost in the middle: How language models use long contexts, 2023.
- 585
586 Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered
587 prompts and where to find them: Overcoming few-shot prompt order sensitivity. In Smaranda
588 Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meet-*
589 *ing of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8086–8098,
590 Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.
591 [acl-long.556](https://aclanthology.org/2022.acl-long.556). URL <https://aclanthology.org/2022.acl-long.556>.
- 592 Ang Lv, Yuhan Chen, Kaiyi Zhang, Yulong Wang, Lifeng Liu, Ji-Rong Wen, Jian Xie, and Rui Yan.
593 Interpreting key mechanisms of factual recall in transformer-based language models, 2024a. URL
<https://arxiv.org/abs/2403.19521>.

- 594 Ang Lv, Ruobing Xie, Xingwu Sun, Zhanhui Kang, and Rui Yan. Language models "grok" to copy,
595 2024b. URL <https://arxiv.org/abs/2409.09281>.
596
- 597 Callum McDougall, Arthur Conmy, Cody Rushing, Thomas McGrath, and Neel Nanda. Copy sup-
598 pression: Comprehensively understanding an attention head, 2023. URL <https://arxiv.org/abs/2310.04625>.
599
- 600 Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Circuit component reuse across tasks in trans-
601 former language models. In *The Twelfth International Conference on Learning Representations*,
602 2024. URL <https://openreview.net/forum?id=fpoAYV6Wsk>.
603
- 604 Microsoft. Reinventing search with a new ai-powered microsoft bing and edge, your copilot for the
605 web, 2023.
- 606 Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter.
607 Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001.
608 <https://distill.pub/2020/circuits/zoom-in>.
609
- 610 Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan,
611 Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli,
612 Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane
613 Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish,
614 and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022.
615 <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.
- 616 OpenAI. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
617
- 618 Alexander Peysakhovich and Adam Lerer. Attention sorting combats recency bias in long context
619 language models, 2023.
- 620 Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases
621 enables input length extrapolation, 2022. URL <https://arxiv.org/abs/2108.12409>.
622
- 623 Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language
624 models are unsupervised multitask learners, 2019.
- 625 Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk, Andy Davis, Quoc V. Le, Geoffrey E. Hinton,
626 and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer.
627 *CoRR*, abs/1701.06538, 2017. URL <http://arxiv.org/abs/1701.06538>.
628
- 629 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: En-
630 hanced transformer with rotary position embedding, 2023. URL <https://arxiv.org/abs/2104.09864>.
631
- 632 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
633 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher,
634 Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy
635 Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,
636 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel
637 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,
638 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,
639 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,
640 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh
641 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen
642 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic,
643 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models,
644 2023. URL <https://arxiv.org/abs/2307.09288>.
- 645 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. MuSiQue: Mul-
646 ti-hop questions via single-hop question composition. *Transactions of the Association for*
647 *Computational Linguistics*, 10:539–554, 2022. doi: 10.1162/tacl_a.00475. URL <https://aclanthology.org/2022.tacl-1.31>.

- 648 Boshi Wang, Xiang Yue, Yu Su, and Huan Sun. Grokked transformers are implicit reasoners: A
 649 mechanistic journey to the edge of generalization, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2405.15071)
 650 2405.15071.
- 651 Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt.
 652 Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In
 653 *The Eleventh International Conference on Learning Representations*, 2023. URL [https://](https://openreview.net/forum?id=NpsVSN6o4u1)
 654 openreview.net/forum?id=NpsVSN6o4u1.
- 655 Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. Retrieval head mecha-
 656 nistically explains long-context factuality, 2024. URL [https://arxiv.org/abs/2404.](https://arxiv.org/abs/2404.15574)
 657 15574.
- 658 Qinan Yu, Jack Merullo, and Ellie Pavlick. Characterizing mechanisms for factual recall in language
 659 models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Confer-*
 660 *ence on Empirical Methods in Natural Language Processing*, pp. 9924–9959, Singapore, Decem-
 661 ber 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.615.
 662 URL <https://aclanthology.org/2023.emnlp-main.615>.
- 663 Fred Zhang and Neel Nanda. Towards best practices of activation patching in language models:
 664 Metrics and methods. In *The Twelfth International Conference on Learning Representations*,
 665 2024. URL <https://openreview.net/forum?id=Hf17y6u9BC>.
- 666 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christo-
 667 pher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer,
 668 Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettle-
 669 moyer. Opt: Open pre-trained transformer language models, 2022. URL [https://arxiv.](https://arxiv.org/abs/2205.01068)
 670 [org/abs/2205.01068](https://arxiv.org/abs/2205.01068).
- 671 Zhenyu Zhang, Runjin Chen, Shiwei Liu, Zhewei Yao, Olatunji Ruwase, Beidi Chen, Xiaoxia Wu,
 672 and Zhangyang Wang. Found in the middle: How language models use long contexts better via
 673 plug-and-play positional encoding, 2024. URL <https://arxiv.org/abs/2403.04797>.

674 A HEAD DISCOVERY RESULTS






Figure 5: Heatmaps of $\Delta\pi$ scores for each head of llama2-7B-chat ($n = 15, n = 25, n = 50$).



720 Figure 6: Heatmaps of $\Delta\pi$ scores for each head of OPT-6.7B ($n = 15, n = 25, n = 50$).



748 Figure 7: Heatmaps of $\Delta\pi$ scores for each head of Baichuan-13B-chat ($n = 20, n = 50, n = 80$).

749
750
751
752
753
754
755