Exploring Context Window of Large Language Models via Decomposed Positional Vectors

Zican Dong¹*, Junyi Li³*, Xin Men⁴, Wayne Xin Zhao¹†, Bingning Wang⁴ Zhen Tian¹, Weipeng Chen⁴, Ji-Rong Wen^{1,2}

¹ Gaoling School of Artificial Intelligence, Renmin University of China ² School of Information, Renmin University of China

³ Department of Computer Science, National University of Singapore ⁴ Baichuan Inc.

dongzican@ruc.edu.cn, junyi_cs@nus.edu.sg
batmanfly@gmail.com, daniel@baichuan-inc.com

Abstract

Transformer-based large language models (LLMs) typically have a limited context window, resulting in significant performance degradation when processing text beyond the length of the context window. Extensive studies have been proposed to extend the context window and achieve length extrapolation of LLMs, but there is still a lack of in-depth interpretation of these approaches. In this study, we explore the positional information within and beyond the context window for deciphering the underlying mechanism of LLMs. By using a mean-based decomposition method, we disentangle positional vectors from hidden states of LLMs and analyze their formation and effect on attention. Furthermore, when texts exceed the context window, we analyze the change of positional vectors in two settings, *i.e.*, direct extrapolation and context window extension. Based on our findings, we design two training-free context window extension methods, **positional vector replacement** and **attention window extension**. Experimental results show that our methods can effectively extend the context window length.

1 Introduction

Recently, Transformer-based large language models (LLMs) have demonstrated excellent capabilities on downstream tasks [1–3], in which positional encodings (e.g., absolute or relative) are widely used in Transformers to better capture positional information within input sequences [4, 5]. However, LLMs typically suffer from a limited input length (called *context window*), which is constrained by the maximum length of training data. Beyond the context window, the positional encodings at larger position indices are out-of-distribution (OOD), not encountered during the training phase. Therefore, when the input sequence exceeds the context window length, there would often be a significant degradation in model performances, as evidenced by a surge in perplexity (PPL) score [6].

Prior work has primarily focused on extending the context window of existing LLMs by manipulating positional encodings. Owing to its excellent performance and long-term decay nature, RoPE [7] has been widely used to learn positional encodings for existing LLMs [8, 9]. To circumvent the OOD positional encodings in RoPE, various methods have been proposed to modify the base [10–12] or positional indices [13–16]. In addition, special relative positional encodings that apply larger negative biases to attention based on the relative distance have achieved promising length extrapolation, which

10320

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

^{*}Equal Contribution.

[†]Corresponding author.

can effectively stabilize the model performance beyond the context window [6, 17, 18]. Furthermore, decoder-only Transformers without positional encodings (NoPE) have been found to be capable of learning implicit positional information [19], and their context window size can be extended via the adjustment of temperature hyper-parameters [20]. However, the above extension methods solely focus on adapting positional encodings or attention scores, lacking a detailed analysis of the underlying mechanisms of hidden states in LLMs.

In this work, we aim to investigate the inner working mechanism of LLMs within and beyond the context window to interpret these context window extension approaches. As the basis, our work is developed by analyzing the positional information implicitly encoded in the hidden states of LLMs across various layers and positions, both within and outside the context window. Inspired by previous work [21], we use a mean-based decomposition approach to disentangle **positional vectors** from the hidden states, which captures the information independent of semantics but related to positions.

Specifically, we first investigate how positional information is formed and examine its impact on the attention mechanism within the context window. Second, for inputs beyond the context window, we analyze the change of positional vectors in two settings, *i.e.*, direct extrapolation and context window extension. Our key findings include: (1) After the first layer, initial tokens can form distinct positional vectors, serving as anchors for shaping positional vectors in subsequent tokens; (2) Positional vectors play a critical role in modulating the long-term decay and establishing attention sinks; (3) When exceeding the context window, OOD positional vector is the major factor contributing to performance degradation, while length extrapolation can effectively keep the consistency of positional vectors both within and beyond the context window; (4) Context window extension methods enable interpolation of positional vectors by adjusting the information flow from initial tokens to subsequent tokens.

Based on the empirical findings, we further propose two training-free context window extension methods from the perspective of interpolating positional vectors: **positional vector replacement** and **attention window extension**. For LLMs with NoPE, the former method replaces the positional vectors in critical layers with interpolated ones; while for LLMs with window attention and NoPE, the latter method directly scales the window size and adjusts the temperature hyper-parameter. We evaluate the length generalization capacities of the proposed methods on PG-19 [22]. Experimental results demonstrate that our methods can effectively generalize to longer texts without fine-tuning, achieving comparable performance to previous methods.

Our main contributions are summarized as follows:

- We explicitly delineate the formation process and the effect of positional vectors, highlighting
 the anchoring role of initial tokens in shaping different positional vectors across tokens and their
 importance in achieving long-term decay and attention sinks.
- We are the first to unify length extrapolation and context window extension from the perspective
 of positional vectors, identifying that preventing OOD positional vectors is crucial for avoiding
 performance degradation.
- We propose two training-free context window extension methods via the lens of adjusting positional vectors, *i.e.*, positional vector replacement and attention window extension. Experimental results show that our methods can effectively generalize to longer texts without fine-tuning.

2 Background

Transformer Decoder-only Transformer [4] has become the foundational architecture for LLMs [4, 8, 1]. For a Transformer with L layers and a context window size C, given an input sequence s of T tokens, i.e., $\{x_1,\ldots,x_T\}$, we denote the output of the l-th layer l as $\mathbf{H}_l^s = \{\mathbf{h}_{l,1}^s,\ldots,\mathbf{h}_{l,T}^s\}$. At each layer, the output \mathbf{H}_l^s is obtained through multi-head attention (MHA) and feed-forward network (FFN) with residual connections applied to both components as follows:

$$\widetilde{\mathbf{H}}_{l}^{s} = \mathrm{MHA}(\mathbf{H}_{l-1}^{s}) + \mathbf{H}_{l-1}^{s}, \quad \mathbf{H}_{l}^{s} = \mathrm{FFN}(\widetilde{\mathbf{H}}_{l}^{s}) + \widetilde{\mathbf{H}}_{l}^{s}. \tag{1}$$

Finally, the output of the last layer \mathbf{H}_L^s is then projected into the logits, which will be used to generate the prediction probability for each token in the vocabulary.

Positional Vector Previous work has found that positional information can be learned and encoded in the hidden states of Transformers [19]. Drawing inspiration from prior work [21], we hypothesize

that each hidden state (*e.g.*, query, key, value, output of each layer) within Transformer can be decomposed into two parts, *i.e.*, a *positional vector* that captures positional information and a *semantic vector* that captures the contextual information. Taking the output $\mathbf{h}_{l,t}^s$ of the *l*-th layer at *t*-th position as an example, it can be decomposed into a positional vector $\mathbf{p}_{l,t}$ and a semantic vector $\mathbf{c}_{l,t}^s$:

$$\mathbf{h}_{l,t}^s = \mathbf{p}_{l,t} + \mathbf{c}_{l,t}^s. \tag{2}$$

Such a decomposition can disentangle two primary factors, namely positional and semantic vectors, for interpreting the internal mechanism of LLMs. Notably, since positional vectors are globally shared across different inputs, there is no superscript s for $\mathbf{p}_{l,t}$. Further, the positional vector $\mathbf{p}_{l,t}$ can be decomposed into a *mean vector* \mathbf{u}_l and a *positional basis* $\mathbf{m}_{l,t}$:

$$\mathbf{p}_{l,t} = \mathbf{u}_l + \mathbf{m}_{l,t},\tag{3}$$

where the mean vector \mathbf{u}_l denotes the mean of the distribution of positional vectors and the positional basis $\mathbf{m}_{l,t}$ denotes the offset of t-th position from the mean vector within the context window size C. Following previous work [21], we adopt a mean-based decomposition method to obtain the above three vectors based on N samples from the training corpus as follows:

$$\mathbf{p}_{l,t} = \frac{1}{N} \sum_{s=1}^{N} \mathbf{h}_{l,t}^{s}, \quad \mathbf{m}_{l,t} = \mathbf{p}_{l,t} - \frac{1}{C} \sum_{t'=1}^{C} \mathbf{p}_{l,t'}, \quad \mathbf{c}_{l,t}^{s} = \mathbf{h}_{l,t}^{s} - \mathbf{p}_{l,t}.$$
(4)

With this decomposition, it offers an explicit way to analyze and explore the positional information encoded in the hidden states of Transformer models. For example, we can use similarity measurements to compare the positional vectors of different positions and also can visualize them in low-dimensional embedding space. In the following sections, we will mainly focus on studying the formation and impact of the positional vector $\mathbf{p}_{l,t}$, and conduct the analysis experiments.

3 Empirical Analysis

3.1 Experimental Settings

To better analyze positional information, we consider model variants with different positional encodings (PE) and attention mechanisms: variants without positional encodings (NoPE) [19] as well as variants with two different positional encodings: RoPE [7] and ALiBi [6]. We continually pre-train the TinyLlama-1.1B checkpoint [23] on 50B tokens from RedPajama [24] with a context window C=2048, resulting in a set of comparison models with different positional encodings and attention mechanisms, as shown in the Table 1. *Full attention* means that each token can attend to all previous tokens, while *window attention* restricts each token to attend only to previous tokens within a window size W. The training details are described in Appendix A. We also evaluate common LLMs (e.g., Llama-3-8B) and LLMs without positional encodings trained from scratch, and the evaluated results are listed in Appendix F.

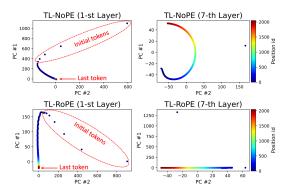
Specifically, we subsample 32K samples with the same number of tokens from RedPajama. We perform the inference on these data to obtain hidden states of LLMs. By using the mean-based decomposition method (Eq. 4), we can obtain the positional vectors $\mathbf{p}_{l,t}$ of tokens in these sample texts.

Table 1: The compared model variants. Full attention is denoted as *Full* and window attention with a window size of W tokens is denoted as *Window* (W). We abbreviate TinyLLaMA as TL.

Model	TL-NoPE	TL-RoPE	TL-ALiBi	TL-Window	TL-Window-80	TL-Window-RoPE
PE	NoPE	RoPE	ALiBi	NoPE	NoPE	RoPE
Attention	Full	Full	Full	Window (512)	Window (80)	Window (512)

3.2 Formation and Effect of Positional Vectors within Context Window

In existing LLMs, the bottom (first) layer typically takes as input token embeddings that lack inherent positional information; while interestingly, the hidden states from top layers can implicitly capture positional information, even without explicit positional encodings [19, 21, 14]. In order to have a deep understanding of implicit positional information, we next investigate the formation and effect of positional vectors in Transformers, with both full attention and window attention.



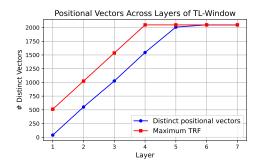


Figure 1: PCA visualization of positional vectors from the 1-st and 7-th layers.

Figure 2: Comparison of distinct positional vectors and theoretical receptive field.

3.2.1 Formation of Positional Vectors with Full Attention

Positional Vector After the First Layer To study how positional information is distributed over different positions, we first visualize the positional vectors $\mathbf{p}_{1,t}$ (Eq. 4) decomposed from the outputs of the first layer using principal component analysis (PCA). As shown in Figure 1 (left column), initial tokens (e.g., ≤ 4 tokens) exhibit significantly distinct positional vectors, while the positional vectors of subsequent tokens are similar to each other. As a comparison, we also present the PCA results of all positional vectors at the 7-th layer. Interestingly, position vectors are evenly distributed across all the positions in Figure 1 (right column). Such a finding indicates that position vectors have captured the corresponding positional information since these vectors are distinct from each other across positions. In other words, being distinct can be considered as a kind of positional evidence. By comparing the left and right columns of Figure 1, it seems that only initial tokens are different from the rest tokens after the first layer, which might suggest that after the first layer, initial tokens have already formed distinct positional information but subsequent tokens have not yet established such information. To investigate the reasons behind this phenomenon, we select the first attention head in the first layer (similar to other heads) to analyze attention scores, as detailed in Appendix B. We can prove that the positional vector $\mathbf{p}_{1,1}$ for the first token is different from the following tokens and the attention scores affect the formation of positional information. Thus, through several layers, the tokens after the first token will gradually form distinct positional vectors (Figure 1 right column).

Positional Information Flow From Initial Tokens By applying positional vectors at top layers (PCA visualized in Appendix G), we find that after forwarding several layers, tokens at all positions can also exhibit distinct positional vectors, and similar findings are also found in previous work [19]. To trace back to the source of positional information, a reasonable speculation is that initial tokens play a key role in the formation of positional information for the rest tokens since only initial tokens capture positional information after the first layer. To validate this, we select two groups of reference tokens: initial tokens (1 \sim 4) and secondary tokens (4 \sim 256), and further analyze what information of these tokens is critical for the formation of positional information in subsequent tokens (>256). Thus, based on a top-down strategy, we conduct an ablation study for each group by respectively deleting the value $\mathbf{v}_{l,t}^s$ of attention module (w/o value), the semantic vector $\mathbf{c}_{l,t}^s$ (w/o semantic vector), positional vector $\mathbf{p}_{l,t}$ (w/o positional vector), and positional basis $\mathbf{m}_{l,t}$ (w/o positional basis). Then, for each variant, we average the outputs of all layers and decompose new positional vectors based on Eq. 4. Finally, we compute the average cosine similarity between the original and new positional vectors for those subsequent tokens (>256) and also report PPL on samples in RedPajama. From Table 2, we can see that removing the positional vector and basis of $1\sim4$ tokens largely affect the positional vectors at later positions (low similarity). Conversely, removing the semantic vector or altering secondary tokens has slight effects on both similarity and PPL. From these findings, we conclude that the positional vectors of initial tokens seem to serve as the role of anchors, largely contributing to the formation of positional information in subsequent tokens.

Table 2: Results of removing different components in attention. *Sim* denotes the cosine similarity between original and new positional vectors, and *PPL* denotes perplexity on RedPajama.

		original	w/o v	alue	w/o posit	ional vector	w/o posit	ional basis	w/o sema	antic vector
		-	1~4	$4 \sim 256$	1~4	$4 \sim 256$	1~4	$4 \sim 256$	1~4	$4 \sim 256$
	Sim	1	-0.1558	0.9797	-0.1810	0.9086	-0.1817	0.9046	0.9985	0.9514
	PPL	7.56	>1000	8.97	>1000	13.36	>1000	10.23	8.20	10.55
TL-RoPE	Sim PPL	1	0.8394	0.9902	0.8505	0.9874	0.1711	0.9944	0.9970	0.9596
	PPL	6.06	11.98	6.44	12.28	6.24	>1000	6.11	6.63	6.85

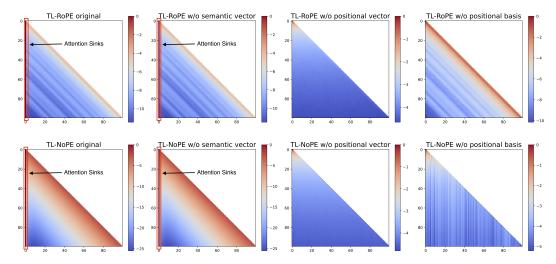


Figure 3: Logarithmic attention maps of TL-RoPE, and TL-NoPE.

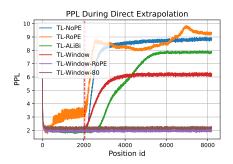
3.2.2 Formation of Positional Vectors with Window Attention

Unlike full attention, LLMs employing window attention restrict each token to attend only to tokens within a window size. Previous work has shown that the maximum theoretical receptive field (TRF) in window attention is equal to the product of the window size W and the layer index l [18].

To analyze how positional vectors change across layers, we compute *the number of distinct positional vectors* within the maximum TRF. Notably, those tokens beyond the maximum TRF share the same positional vectors due to translation invariance [18]. Specifically, we first randomly select a positional vector outside the maximum TRF and then compute the cosine similarity between positional vectors within the maximum TRF and the selected vector. We consider the positional vector with a similarity score lower than a threshold (*i.e.*, 0.99) as *distinct*. Figure 2 presents the number of distinct positional vectors and TRF at each layer. We can see that after the first layer, only initial tokens show distinct positional vectors, further verifying the findings in Section 3.2.1. As the layer increases, more tokens display different positional information and the number of distinct positional vectors increases by 512 (a window size W) with each additional layer. The reason is that due to the constraint of window attention, each token at the preceding layer can only influence tokens within the window size at the next layer. As *being distinct* indicates the formation of position information, **similar positional information flow from initial tokens to subsequent tokens also occurs for window attention, but gradually propagating across both windows and layers.**

3.2.3 Effect of Positional Vectors on Attention

After discussing the formation of positional vectors, we explore their impact on the attention module, mainly focusing on the attention scores. We first extract queries and keys from each head in all layers, and then compute the average attention scores in the following four settings, including (1) *original:* the original model, (2) *w/o semantic vector:* removing the semantic vectors of keys and queries, (3) *w/o positional vector:* removing the positional vectors of keys and queries, (4) *w/o positional basis:* removing the positional basis of keys and queries. Figure 3 presents the logarithmic attention scores for the first 100 tokens in the first head and fifth layer (similar results in many other heads and layers).



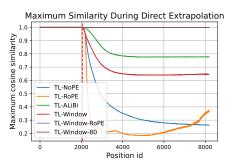


Figure 4: **Left**: The average PPL across positions during direct extrapolation. **Right**: The maximum cosine similarity between positional vectors within and beyond context window during extrapolation.

Effect of Positional Vectors on Attention Sinks Previous work has found that the initial tokens will be assigned high attention scores, called "*attention sinks*" [15], which can be clearly observed in Figure 3. However, once the positional vector or positional basis is removed from the keys and queries, the attention scores between initial tokens and other tokens drop significantly for TL-NoPE and TL-RoPE. This finding suggests that **the presence of attention sinks is likely attributed to the inherent positional information in the positional vectors of initial tokens.**

Effect of Positional Vectors on Long-term Decay For long texts, the attention scores of LLMs often exhibit a long-term decay pattern, which means that the score decreases as the relative distance between tokens increases [7, 6]. However, as shown in Figure 3, when removing the positional vector or positional basis, TL-NoPE fails to exhibit long-term decay. Even with explicit relative positional encoding, the distribution of attention scores in TL-RoPE tends to be smooth after removing decomposed positional vectors. Therefore, positional vectors also play a crucial role in the long-term decay property of attention scores.

3.3 Effect of Positional Vectors beyond Context Window

Typically, when dealing with texts that exceed the context window, there are two lines of research, *i.e.*, direct extrapolation and context window extension. In this section, we aim to investigate the change of positional vectors in these two methods for revealing their effectiveness.

3.3.1 Direct Extrapolation

Relationship Between Positional Vectors and Length Extrapolation Ability To examine the impact of positional vectors in direct extrapolation, we reuse the trained model variants in Table 1 to perform inference on samples consisting of 8192 tokens. Further, we analyze the change in PPL score and the maximum cosine similarity between positional vectors within and beyond the context window. As shown in Figure 4 Left, only TL-Window-RoPE and TL-Window-80 demonstrate the length extrapolation ability, maintaining stable PPL across longer texts. These models can preserve the consistency of positional vectors both within and beyond the context window (high similarity in Figure 4 Right). Conversely, the rest models, including those with extrapolated positional encodings or window attention (*e.g.*, TL-ALiBi), struggle to generalize to longer contexts. Notably, these models exhibit rapid changes in positional vectors (beyond 2048), diverging from the distributions observed within the context window. Thus, our findings underscore the critical role of the stability of positional vectors in enhancing the capability for length extrapolation.

Effect of OOD Positional Vectors Beyond the context window, position vectors are not encountered during training and are out-of-distribution from those vectors within the context window. To explore whether OOD positional vector is a key factor in performance degradation, we select TL-NoPE for evaluation, which does not use explicit positional encodings. First, we compare the attention distribution within and beyond the context window. Figure 5 shows the attention map and scores between initial and rest tokens by averaging all heads of the 5-th layer (similar results in other layers). Once exceeding the context window (T=2048), the attention distribution in these positions changes sharply, losing the characteristics of attention sinks and long-term decay. Since these properties

highly depend on the positional vectors within the context window, we speculate that **OOD positional vectors disrupt the original attention distribution**. Besides, we feed the positional vectors of the last layer into the linear projection of the softmax layer to get the logits at different positions. Figure 5 (Right) presents that the logits within the context window are tightly similar while others show different distributions. Thus, **the OOD positional vectors can damage the token prediction probability distribution**, thereby leading to performance degradation.

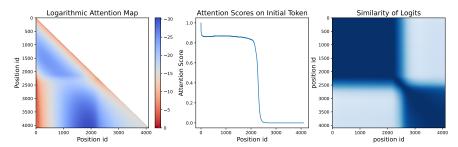


Figure 5: **Left**: Attention map of TL-NoPE. **Middle**: Attention Scores between initial token and others in TL-NoPE. **Right**: Similarity of logits of positional vectors across positions in TL-NoPE.

Table 3: The interpolation results of positional vectors, where Factor (= Target Length/C) is the expansion factor of the context window, Ratio is the effective interpolation ratio of positional vectors (detailed in Appendix C), and Similarity is the average cosine similarity between the scaled positional vector and the original most similar positional vector by averaging all layers.

Model	Method	Target Length	Factor	Ratio	Similarity	$PPL/\Delta PPL$
	Attention Scaling ($\lambda = 1.2$)		2	2.56	0.98	8.95/+1.42
TL-NoPE	Attention Scaling ($\lambda = 1.3$)	8192	4	4.30	0.94	17.87/+10.34
	Initial Scaling ($\lambda = 1.2$)	4096	2	2.38	0.97	9.82/+2.29
	Initial Scaling ($\lambda = 1.3$)	8192	4	4.10	0.91	32.78/+25.25
TL-RoPE	Dynamic NTK	4096	2	2.05	0.99	6.00/-0.02
IL-ROPE	Dynamic NTK	8192	4	3.75	0.96	6.78/+0.76

3.3.2 Context Window Extension

Change of Positional Vectors When Extending Context Windows To investigate why context window extension can prevent performance degradation, we analyze the change of positional vectors in two training-free context window extension methods, including dynamic-NTK [11] for TL-RoPE and attention scaling $(\mathbf{q}_i \mathbf{k}_j)$ multiplied by a scaling factor λ [20] for TL-NoPE. From Figure 6, we can see that after context window extension, positional vectors have undergone interpolation compared to the original ones. Comparing the *Factor* and *Ratio* metrics in Table 3, we conclude that the effective interpolation ratio is close to the expansion factor (e.g., 2 vs 2.56). Besides, as the expansion factor increases, there is a decrease in *Similarity* and an increase in *PPL*. Therefore, we suspect that imperfect interpolation may be a major reason for the decline in model performance.

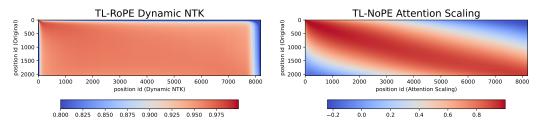


Figure 6: The average cosine similarity between the scaled and original positional vectors.

Effect of Initial Tokens on Context Window Extension Since the initial tokens serve as the anchor for the formation of subsequent positional vectors, we evaluate whether changing the information flow from the initial tokens to the rest tokens can achieve the interpolation effect. To avoid the effect of OOD positional encodings, we follow the attention scaling method on TL-NoPE but only scale the attention logits between the initial tokens and others, denoted as *Initial Scaling*. As shown in Table 3, it can achieve comparable performance and interpolation ratios closer than scaling all attention logits in Attention Scaling (*e.g.*, 2.38 vs 2.56), further underscoring that the interpolation of positional vectors is mainly achieved by adjusting the information flow of anchor tokens.

4 Extending Context Window via Positional Vectors

Inspired by our analysis of the formation of positional vectors and the interpolation of positional vectors when extending the context window, we propose two training-free context window extension methods, *i.e.*, **positional vector replacement** and **attention window extension**. The pseudocode of these methods can presented in Appendix E.

4.1 Positional Vector Replacement

In Section 3.2.3, we show that when exceeding the context window, the OOD positional vectors tend to cause the collapse of attention distribution. Further, we observe that context window extension methods can achieve length interpolation of positional vectors and the effective interpolation ratio is close to the expansion factor of the context window. Thus, we propose to replace all the implicitly learned positional vectors with the interpolated ones, called *positional vector replacement*, to avoid the OOD issue in LLMs without positional encodings (NoPE).

Specifically, we linearly interpolate the positional vectors within the context window with an interpolation ratio r and multiply the interpolated ones with a times $\alpha \ (\geq 1)$. In practice, we find that properly increasing the interpolation ratio r and times α can achieve better effectiveness of interpolation (details are discussed in Appendix D). Owing to the critical role of initial tokes, the positional vectors of the first four tokens remain unchanged, while those of subsequent tokens are replaced with the interpolated vectors. The replaced output $\hat{\mathbf{h}}_{l,t}$ for each layer can be formulated as:

$$\hat{\mathbf{h}}_{l,t} = \mathbf{h}_{l,t} - \mathbf{p}_{l,t} + \alpha \hat{\mathbf{p}}_{l,t}, \tag{5}$$

$$\{\hat{\mathbf{p}}_{l,5},\dots,\hat{\mathbf{p}}_{l,r(C-4)+5}\}$$
 = Interpolation $(\{\mathbf{p}_{l,5},\dots,\mathbf{p}_{l,C}\}),$ (6)

where C, l, and s represent the original context window size, replaced layer, and interpolation ratio. Since replacing positional vectors for all layers requires heavy recalculation efforts and the positional information is passed across layers, we only apply the replacement strategy to a single early layer. We find that the 4-th layer is the optimal layer for replacement in TL-NoPE, as shown in Figure 8.

4.2 Attention Window Extension

As discussed in Section 3.2.2, the positional vectors are shaped across layers and windows by the distinct positional information of initial tokens. Inspired by these observations, we propose *attention window extension*, the first training-free length interpolation method for window attention-based LLMs without positional encodings. The core idea is to extend the attention window size to control the formation of positional vectors. When scaling the context window by a ratio, the window size also needs to be extended by the same interpolation ratio r. However, for positions in the extended first window $\{W+1,\ldots,rW\}$, their position vectors are OOD. To avoid this, we follow the attention scaling method [20] and scale the attention logits with a scaling factor λ , achieving better interpolation of positional vectors. We define the attention score a_{ij} between query $\mathbf{q}_i \in \mathbb{R}^{D_H}$ and key $\mathbf{k}_j \in \mathbb{R}^{D_H}$ for any heads and layers as:

$$a_{ij} = \frac{\exp(\lambda \mathbf{q}_i \mathbf{k}_j / \sqrt{D_H})}{\sum_{z=i-r_W}^i \exp(\lambda \mathbf{q}_i \mathbf{k}_z / \sqrt{D_H})}.$$
 (7)

4.3 Results on Language Modeling

To assess the effectiveness of our proposed methods, we evaluate language modeling performance on the test set of PG-19 [22]. In line with previous work [6], we measure PPL across various input

Table 4: Results of language modeling in PG-19. The context window size C is 2048.

Model	Interpolation Method	Factor	2K	4K	6K	8K
TL-RoPE	-	-	10.17	$> 10^{3}$	$> 10^{3}$	$> 10^{3}$
	Dynamic NTK	-	10.17	10.45	11.28	28.58
TL-NoPE	-	-	11.92	$> 10^{3}$	$> 10^{3}$	$> 10^{3}$
	Attention Scaling	$\lambda = 1.2$ $\lambda = 1.3$	17.03 32.07	17.05 43.84	54.26 51.50	$> 10^3$ 46.59
	Positional Vector Replacement (ours)	$ r = 2, \alpha = 1.1 $ $ r = 5, \alpha = 1.3 $	13.54 28.15	15.58 47.65	- 49.79	73.79
	-	-	12.86	713.51	660.30	660.51
TL-Window	Attention Window Extension (ours)	$\left \begin{array}{l} r=2, \lambda=1.1 \\ r=4, \lambda=1.2 \end{array} \right $	13.70 17.23	14.10 31.66	- 29.27	29.30

lengths (from 2K to 8K) using a sliding window approach. We apply positional vector replacement to TL-NoPE and attention window extension to TL-Window. All the hyper-parameters are selected according to the PPL and the change of positional vectors across layers. For compared baselines, we select Dynamic-NTK [11] for TL-RoPE and Attention Scaling [20] for TL-NoPE.

The results are shown in Table 4. First, without interpolation, the PPL increases extremely after beyond the context window (e.g., $>10^3$). When using the positional vector replacement or attention window extension methods, we observe that PPL decreases substantially, showing the effectiveness of our proposed methods. Compared to attention scaling, our attention window extension method successfully extends the context window to 8K tokens with lower PPL. Moreover, our positional vector replacement method achieves similar performance to attention scaling within 6K tokens but shows increased PPL at 8K. We attribute this phenomenon to the decreasing effective interpolation ratio across layers, as shown in Figure 9. Additionally, an increase in PPL with the rising interpolation ratio r is also observed in both our methods, likely due to imperfect interpolation of positional vectors.

5 Related Work

Position Information in Transformers Positional information was crucial in Transformer-based LLMs, to enhance the sequence modeling abilities. The vanilla Transformer introduced absolute positional encodings, using a unique embedding to each position and adding it to the corresponding input embedding [4]. In contrast, relative positional encodings introduced biases based on the relative distance between tokens within attention modules [25–27, 6, 7]. Besides explicit positional encodings, some work investigated the implicit positional information within hidden states of Transformers. Even without positional encodings, positional information was found in hidden states of Transformer decoders [19, 28, 29]. Besides, prior work decoupled positional basis from hidden states in Transformers and analyzed geometric properties [21]. Our work mainly explores positional information embedded in the hidden states of LLMs, examining the formation and impact of positional vectors, and using it to analyze the mechanism of context window for LLMs.

Extending Context Window LLMs were often constrained by pre-defined context windows. When processing inputs that exceed these windows, models typically encountered OOD issues, leading to significant performance degradation. To meet the growing demands of long context tasks [30, 31], various methods were proposed to address this limitation and model longer texts, which can be roughly categorized into length extrapolation and context window extension [32]. Length extrapolation techniques aimed to maintain stable PPL regardless of text length by designing specialized positional encodings or window attention mechanisms [6, 17, 18, 15, 14]. Conversely, context window extension methods focused on extending the context window of existing models by adapting positional encodings or temperature hyper-parameters, thereby enlarging the context window with minimal performance loss [13, 12, 16, 20, 11, 10]. This paper bridges the concepts of length extrapolation and context window extension through the lens of positional vectors, enhancing the interpretability of context windows in LLMs.

6 Conclusion

In this work, we explored the inner working mechanism of LLMs within and beyond the context window via decomposed positional vectors. We found that the initial tokens initially present different positional information and serve as anchors for shaping the positional vectors of subsequent tokens. Besides, after exceeding the context window, length extrapolation methods maintain the stability of positional vectors, while context window extension methods achieve the interpolation of positional vectors. Based on our observations, we proposed two methods: positional vector replacement and attention window extension, which achieve training-free context window extension for specific LLMs. We believe that positional vectors will serve as an effective tool for analyzing the context window of LLMs and promote the design of better algorithms for extending the context windows of LLMs.

7 Limitation

Our work provides an extensive discussion and analysis of the context window through the lens of positional vectors. However, our study is mainly constrained by the use of small-scale LLMs that we trained ourselves, due to the unavailability of existing LLMs with the specific positional encodings and attention patterns required for our experiments. Though some mainstream LLMs are evaluated, these models are all based on RoPE. Furthermore, we have demonstrated the effectiveness of our proposed methods solely on our own models, again limited by the absence of suitable external models. In future work, we aim to seek a broader range of models to validate our findings more comprehensively.

Acknowledgement

This work was partially supported by National Natural Science Foundation of China under Grant No. 62222215, Beijing Municipal Science and Technology Project under Grant No. Z231100010323009, and Beijing Natural Science Foundation under Grant No. L233008. Xin Zhao is the corresponding author.

References

- [1] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020.
- [2] OpenAI. GPT-4 technical report. CoRR, abs/2303.08774, 2023.
- [3] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. A survey of large language models. *CoRR*, abs/2303.18223, 2023.
- [4] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [5] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. In ACL, 2019.
- [6] Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net, 2022.

- [7] Jianlin Su, Murtadha H. M. Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024
- [8] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023.
- [9] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023.
- [10] bloc97. NTK-Aware Scaled RoPE allows LLaMA models to have extended (8k+) context size without any fine-tuning and minimal perplexity degradation., 2023.
- [11] emozilla. Dynamically Scaled RoPE further increases performance of long context LLaMA with zero fine-tuning, 2023.
- [12] Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models. *CoRR*, abs/2309.00071, 2023.
- [13] Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. Extending context window of large language models via positional interpolation. *CoRR*, abs/2306.15595, 2023.
- [14] Chi Han, Qifan Wang, Wenhan Xiong, Yu Chen, Heng Ji, and Sinong Wang. Lm-infinite: Simple on-the-fly length generalization for large language models. *CoRR*, abs/2308.16137, 2023.
- [15] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *CoRR*, abs/2309.17453, 2023.
- [16] Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. LLM maybe longlm: Self-extend LLM context window without tuning. *CoRR*, abs/2401.01325, 2024.
- [17] Ta-Chung Chi, Ting-Han Fan, Peter J. Ramadge, and Alexander Rudnicky. KERPLE: kernelized relative positional embedding for length extrapolation. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 December 9, 2022, 2022.*
- [18] Ta-Chung Chi, Ting-Han Fan, Alexander Rudnicky, and Peter J. Ramadge. Dissecting transformer length extrapolation via the lens of receptive field analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13522–13537. Association for Computational Linguistics, 2023.
- [19] Adi Haviv, Ori Ram, Ofir Press, Peter Izsak, and Omer Levy. Transformer language models without positional encodings still learn positional information. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, *Abu Dhabi, United Arab Emirates, December 7-11*, 2022, pages 1382–1390. Association for Computational Linguistics, 2022.

- [20] Jie Wang, Tao Ji, Yuanbin Wu, Hang Yan, Tao Gui, Qi Zhang, Xuanjing Huang, and Xiaoling Wang. Length generalization of causal transformers without position encoding. CoRR, abs/2404.12224, 2024.
- [21] Jiajun Song and Yiqiao Zhong. Uncovering hidden geometry in transformers via disentangling position and context. *CoRR*, abs/2310.04861, 2023.
- [22] Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. Compressive transformers for long-range sequence modelling. In 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net, 2020.
- [23] Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. Tinyllama: An open-source small language model, 2024.
- [24] Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, 2023.
- [25] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 464–468. Association for Computational Linguistics, 2018.
- [26] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhut-dinov. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, pages 2978–2988. Association for Computational Linguistics, 2019.
- [27] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [28] Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Payel Das, and Siva Reddy. The impact of positional encoding on length generalization in transformers. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023*, 2023.
- [29] Ta-Chung Chi, Ting-Han Fan, Li-Wei Chen, Alexander Rudnicky, and Peter J. Ramadge. Latent positional information is in the self-attention variance of transformer language models without positional embeddings. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1183–1193. Association for Computational Linguistics, 2023.
- [30] Zican Dong, Tianyi Tang, Junyi Li, Wayne Xin Zhao, and Ji-Rong Wen. BAMBOO: A comprehensive benchmark for evaluating long text modeling capacities of large language models. In Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy, pages 2086–2099. ELRA and ICCL, 2024.
- [31] Yuhao Wang, Ruiyang Ren, Junyi Li, Wayne Xin Zhao, Jing Liu, and Ji-Rong Wen. Rear: A relevance-aware retrieval-augmented framework for open-domain question answering. *CoRR*, abs/2402.17497, 2024.
- [32] Saurav Pawar, S. M. Towhidul Islam Tonmoy, S. M. Mehedi Zaman, Vinija Jain, Aman Chadha, and Amitava Das. The what, why, and how of context length extension techniques in large language models A detailed survey. *CoRR*, abs/2401.07872, 2024.
- [33] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *CoRR*, abs/2307.08691, 2023.

- [34] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. The llama 3 herd of models. CoRR, abs/2407.21783, 2024.
- [35] Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. Yi: Open foundation models by 01. ai. *CoRR*, abs/2403.04652, 2024.
- [36] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *CoRR*, abs/2309.16609, 2023.

A Training and Experimental Details

We continue pre-training all our models from the TinyLlama³ [23] checkpoint on the RedPajama [24] dataset. All models undergo the same training process, with differences only in their positional encoding and attention patterns. Each model is trained on 16 A800 GPUs over two days. Detailed training parameters are provided in Table 5.

Table 5: Training Details of Models.

Training Data	RedPajama [24]
Tokens	50B
Parameters	1.3B
Context Window Size	2048
Decay style	cosine
Learning Rate	2e-5
Min Learning Rate	1e-6
Optimizer	AdamW(0.95,0.9)
Warmup Steps	3000
Batch size	48
Gradient clipping	1.0

In addition, all the subsequent experiments are computed in 8 A800 GPUs.

B The Positional Vectors After the First Layer

Though previous work [29, 28] have proven that implicit positional information can be encoded in hidden states after one attention module, they only set the attention logits are equal regardless of queries and keys, which does not hold in actual Transformers. In this section, we demonstrate the preference in attention scores promotes the formation of different positional information in the initial tokens.

B.1 Details

For the s-th sample in the corpus, we denote $\mathbf{v}_{1,i}^s$, $\hat{\mathbf{h}}_{1,i}^s$ as the value and output of the first attention head and the first layer at position i, and $a_{i,j}^s$ as the attention score between position i and position j. We also denote $\hat{\mathbf{p}}_{1,i}$ as the positional vector of the attention output, and \mathbf{u}_v as the mean vector of values. The positional vector can be represented as the formula:

$$\hat{\mathbf{p}}_{1,1} = \frac{1}{N} \sum_{s=1}^{N} \hat{\mathbf{h}}_{1,1}^{s} = \frac{1}{N} \sum_{s=1}^{N} \mathbf{v}_{1,1}^{s} = \mathbf{u}_{v}$$
 (8)

$$\hat{\mathbf{p}}_{1,i} = \frac{1}{N} \sum_{s=1}^{N} \hat{\mathbf{h}}_{1,i}^{s} = \frac{1}{N} \sum_{s=1}^{N} \sum_{i=1}^{i} a_{i,j}^{s} \mathbf{v}_{1,j}^{s}$$
(9)

For the first token, the output is equal to the value of itself, so the positional vector is equal to the mean of values. In addition, When $a_{i,j}^s=1/i$, positional information of all positions is equal as follows:

$$\hat{\mathbf{p}}_{1,i} = \frac{1}{N} \sum_{s=1}^{N} \sum_{j=1}^{i} \frac{1}{i} \mathbf{v}_{1,j}^{s} = \frac{1}{i} \sum_{j=1}^{i} \frac{1}{N} \sum_{s=1}^{N} \mathbf{v}_{1,j}^{s} = \frac{1}{i} \sum_{j=1}^{i} \mathbf{u}_{v} = \mathbf{u}_{v}.$$
 (10)

However, due to the preferences in attention scores, values that can be decomposed into more vector $\hat{\mathbf{p}}_{1,i} - \mathbf{u}_v$ will be assigned larger weights, making the positional information of the following tokens differ from the beginning token. In summary, the preferences of attention scores force the formation of different positional vectors of the following tokens with the initial ones.

³https://huggingface.co/TinyLlama/TinyLlama-1.1B-intermediate-step-1431k-3T

B.2 Proof of Preference in Attention Scores

Following previous work [28], we only focus on a single attention head in the first layer with specific weights. In our parameterization, we only consider the first two dimensions. We demonstrate the capacity of causal Transformers to learn this capacity with either NoPE or RoPE.

In our settings, the Transformer has H attention heads in each layer and the word embedding matrix is $\mathbf{W}_E \in \mathbb{R}^{D \times V}$, where D is the dimension of the hidden states and V is the number of vocabulary. We set that the first dimension in word embedding conforms to normal distribution $\mathcal{N}(0,1)$, (i.e., $e_{1,t} \sim \mathcal{N}(0,1), \forall t \in \{1,\ldots,V\}$), while the second dimension is 1. Other dimensions are arbitrary values. Thus, the word embedding matrix \mathbf{W}_E is:

$$\mathbf{W_{E}} = \begin{bmatrix} e_{1,1} & e_{1,2} & e_{1,3} & \dots & e_{1,V} \\ 1 & 1 & 1 & \dots & 1 \\ e_{3,1} & e_{3,2} & e_{3,3} & \dots & e_{3,V} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ e_{D,1} & e_{D,2} & e_{D,3} & \dots & e_{D,V} \end{bmatrix}_{D \times V}$$
(11)

Next, we set the projection matrix of query, key, and value as \mathbf{W}_Q , \mathbf{W}_K , \mathbf{W}_V of the first layer and first head of Transformers.

$$\mathbf{W}_{Q} = \begin{bmatrix} 0 & 1 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & \dots & 0 \end{bmatrix}_{\frac{D}{H} \times D}, \mathbf{W}_{K} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \dots & 0 \end{bmatrix}_{\frac{D}{H} \times D}, \mathbf{W}_{V} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \dots & 0 \end{bmatrix}_{\frac{D}{H} \times D}.$$

$$(12)$$

 \mathbf{W}_K make sure that only the first dimension will be considered in attention scores. \mathbf{W}_Q transfers the second dimension in the input to the first dimension. For Transformers without positional vectors (NoPE), the attention logits for each key at position i is $a_{i,j}^s = e_{1,s_i}$. Thus, the vectors with large first dimensions will be assigned with large attention logits, which proves that the model without positional vectors can learn to preference some values regardless of the query.

For Transformer with RoPE, we assume that the first two dimensions correspond to the basis of the rope with a value of 1/10000 and the maximum context window size is 2048. Thus, for the largest relative distance i-j, the rotation angle is smaller than $\frac{\pi}{2}$. Thus, the attention score can be represent as $a_{i,j}^s = e_{1,s_i} \cos{((i-j)/10000)}$. Thus, the attention logits will be larger than zero only if the first dimension of the key is larger than zero. In addition, for the same relative distance, keys with larger first dimensions have large attention logits. Thus, keys with positive values in the first dimension will be assigned greater attention weights.

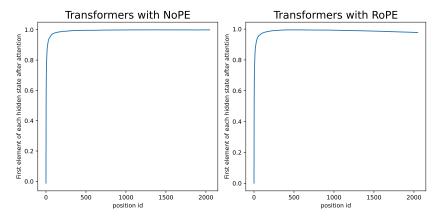
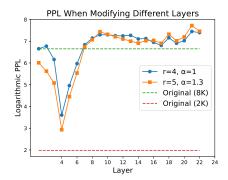


Figure 7: The values of first elements of the output of single head attention due to attention preferences in Transformers with NoPE and RoPE.

We also experimentally examine the first element of the output at different positions. With the above settings, we generate 10000 sequences with a length of 2048 from this distribution. Then,



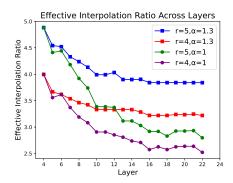


Figure 8: Changes of PPL with replacement at different layers.

Figure 9: Effective scaling factors of positional vectors at each layer.

we compute the averaged first element of each hidden state after attention in the first layer and head, which represents the positional vectors at that position for both NoPE and RoPE. As shown in Figure 7, the first element of the output increases fast at initial positions. Here, we can observe that the attention preferences make the first element of the output increase fast at the initial positions and tend to stabilize at a later position, further proving the different positional vectors of following tokens from the initial ones.

In summary, Transformers with either NoPE or RoPE can learn preference in attention score. In addition, the preference in attention scores forces the different positional vectors of the subsequent tokens with the initial tokens.

C Defination of Effective Interpolation Ratio

Given a Transformer with a context window size C. For samples with length C'(C' > C), we disentangle positional vectors $\{\mathbf{p}_{l,1},\ldots,\mathbf{p}_{l,C'}\}$. Subsequently, we employ a context window extension method and re-compute the position vectors $\{\mathbf{p}'_{l,1},\ldots,\mathbf{p}'_{l,C'}\}$. We first define the corresponding position indices $f(\mathbf{p}'_{l,t})$ for positional vector after extension $\mathbf{p}'_{l,t}$) as the indices of the most similar positional vectors before extension:

$$f(\mathbf{p'}_{l,t}) = \arg\max_{1 \le i \le C'} \frac{\mathbf{p}_{l,i}^T \mathbf{p'}_{l,t}}{\|\mathbf{p}_{l,i}\| \|\mathbf{p'}_{l,t}\|}.$$
(13)

Then, the effective interpolation ratio r' can be represented as the ratio between the maximum indices of position vectors after extension whose corresponding position indices are the context window size C, as shown in Equation 14.

$$r = \frac{\arg\max_{1 \le t \le C'} (f(\mathbf{p'}_{l,t}) = C)}{C}.$$
(14)

D Analysis of Positional Vector Replacement

D.1 Optimal Replacement Layer

Replacing positional vectors for all layers requires heavy recalculation efforts, so we only select one critical layer to apply the replacement strategy. We evaluate the performance (logarithmic PPL score) of our replacement strategy at each layer in TL-NoPE, using different interpolation ratios and expansion factors, i.e., $(r=4,\alpha=1)$ and $(r=5,\alpha=1.2)$, on samples with 8K tokens from RedPajama. As shown in Figure 8, the PPL is the lowest when replacing the 4-th layer of TL-NoPE. Thus, we choose the 4-th layer for replacement in TL-NoPE in other experiments.

D.2 Effective Interpolation Ratio

We further examine the effective interpolation ratio of positional vectors with different settings. In our setting, we only replace the positional vectors in the 4-th layer of TL-NoPE with interpolated

positional vectors in four settings on samples with 8K tokens from RedPajama: (1) $r=4, \alpha=1,$ (2) $r=5, \alpha=1,$ (3) $r=4, \alpha=1.3,$ (4), $r=5, \alpha=1.3.$

Figure 9 presents the effective interpolation ratio in each layer. As the layer increases, the effective interpolation ratio decreases as the layer increases. When the interpolation ratio is equal to the expansion factor of the context window, i.e., r=4, the effective interpolation ratio is unavoidably smaller than the expansion factor, leading to degraded performance. In addition, multiplying the interpolated positional vectors with a larger times e.g., $\alpha=1.3$, alleviates the decrease of effective interpolation ratio across layers. Thus, we suggest to properly increase the interpolation ratio r and times α .

E Pseudo Code

E.1 Positional Vector Replacement

For Positional Vector Replacement, we give the implementation with Pytorch code in Algorithm 1, which can be inserted after the output of the selected layer.

Algorithm 1 PyTorch-style Pseudocode of Positional Vector Replacement

h, p # hidden states, positional vectors

T, layer, s, alpha # context window size, interpolation ratio, selected layer, scaling factor of positional vectors.

h[:,4:] -= p[layer, 4:h.shape[1]].unsqueeze(0) # removing original positional vectors.

interpolated = torch.nn.functional.interpolate(p[layer, 4:T].transpose(0,1).unsqueeze(0), size = int(T*s), mode = 'linear', align_corners=True).transpose(1,2) # Linear interpolation of positional vectors.

h[:,4:] += alpha*interpolated[:,:h.shape[1]-4] # Replacing with new positional vectors.

E.2 Attention Window Extension

For attention window extension, we give the implementation with Flash-Attention-2 [33] in Algorithm 2.

Algorithm 2 PyTorch-style Pseudocode of Attention Window Extension

query, key, value, attn_output # queries, keys, values, and output in the attention model W, lambda, s # original window size, scaling factor of attention logits, window extension factor flash_attn_varlen_func # attention function in flash attention 2.

new_window = W*s # extended window size

attn_output = flash_attn_varlen_func(query, key*lambda, value, ..., window_size = (new_window, new_window))

F Results of Additional LLMs

To ensure a fair comparison of positional vectors across various attention mechanisms and positional encodings, we conducted continual pre-training using TinyLlama under consistent settings. However, TinyLlama is a relatively small language model with suboptimal performance and continual pre-training may result in positional vectors exhibiting properties distinct from those obtained through

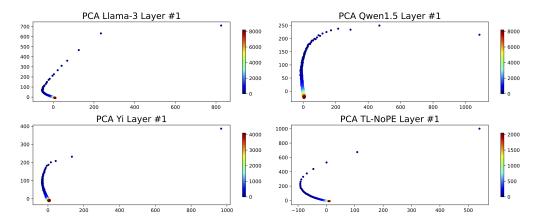


Figure 10: PCA visualization of positional vectors from the 1-st layer of Llama-3-8B, Qwen1.5-7B, Yi-9B, and TL-NoPE-new.

training from scratch. Therefore, we selected three mainstream LLMs: Llama-3-8B [34], Yi-9B [35], and Qwen1.5-7B [36], for comparison. Additionally, we trained a new LLM, TL-NoPE-new, from scratch under the same conditions as TL-NoPE. In a similar vein, we extracted positional vectors using 32K samples from the RedPajama dataset.

F.1 Formation of Positional Vectors within Context Window

Through principal component analysis (PCA), we first visualize the positional vectors from the initial layer of these LLMs, as illustrated in Figure 10. Consistent with our expectations, the initial tokens exhibit distinct positional information, while the subsequent tokens display a high degree of similarity. This observation supports the conclusion that the first-layer attention mechanism makes the initial tokens form unique positional information, as discussed in Section 3.2.1.

Furthermore, we remove different components of the value vectors at different positions across all attention heads after the first layer. We then evaluate the impact of these modifications on both the positional vectors and the perplexity (PPL). As shown in Table 6, removing the positional basis of the initial tokens significantly degrades the model's performance. Conversely, removing the components from subsequent tokens has a relatively smaller effect, highlighting the pivotal role of initial token positioning in influencing later tokens. However, we observe that larger LLMs are more attuned to semantic information and less affected by the removal of positional vectors compared to smaller LLMs.

Table 6: Results of removing different components in attention.

		original	w/o	value w/o pos		v/o positional vector		w/o positional basis		w/o semantic basis	
		-	$0\sim4$	32-256	$0 \sim 4$	$32 \sim 256$	$0\sim4$	32-256	$0\sim4$	32-256	
Llama-3	simi	1	0.75	0.9995	0.75	0.9583	0.2059	0.9997	0.9259	0.8666	
	ppl	6.74	16.27	6.63	17.20	8.4	>1000	6.60	17.6	15.18	
Yi-9B	simi	1	0.98	0.9999	0.92	0.9998	0.5368	1	0.91	0.9996	
	ppl	7.08	8.03	6.56	37.92	6.62	>1000	6.52	42.271	7.08	
Qwen-1.5-7B	simi	1	0.98	0.9997	0.9847	0.9986	0.7382	0.9993	0.9998	0.9951	
	ppl	7.97	9.51	8.03	9.51	8.04	217.13	7.98	8.09	8.68	
TL-NoPE-new	simi	1	0.70	0.95	0.69	0.95	0.41	0.93	0.99	1.0	
	ppl	11.03	224.74	22.36	263.53	20.91	>1000	21.78	11.66	12.699	

F.2 Effect of Positional Vectors on Attention

In line with the experiments conducted in Section 3.2.3, we extract various components from the keys and queries to assess their impact on attention scores. The attention maps for the first 50 tokens are illustrated in Figure 11. When both the positional vectors and positional basis are removed, attention

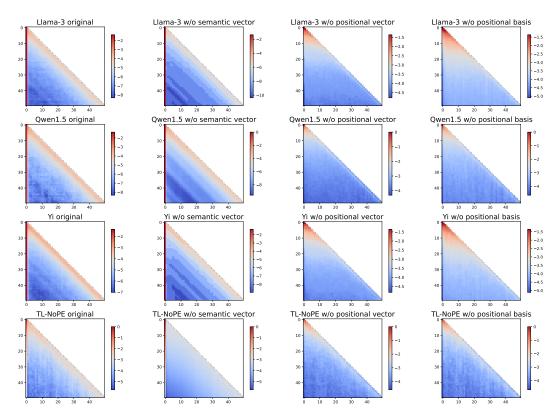


Figure 11: Logarithmic attention maps of Llama-3-8B, Qwen1.5-7B, Yi-9B, and new TL-NoPE.

sinks disappear, and long-term decay is observed, which is consistent with the behavior seen in TinyLlama.

F.3 Effect of Positional Vectors Beyond Context Window

Table 7: Resuls of PPL and change of positional vectors during direct extrapolation.

model	context window(C)	PPL(C)	PPL(2C)	Simi(2C)
Llama-3-8B	8192	6.74	>1000	0.30
Yi-9B	4096	7.08	102.58	0.24
TL-NoPE-new	2048	11.75	351.49	0.71

Table 8: Change of attention sinks and output logits beyond context window.

model	context window (C)	property	0~C	C~1.5C	1.5C~2C
Llama-3-8B	8192	attention sink logits similarity	0.467 1	0.1 0.9	0.005 0.88
Yi-9B	4096	attention sink logits similarity	0.68 1	0.344 0.98	0.056 0.97
TL-NoPE-new	2048	attention sink logits similarity	0.17 1	0.02 0.97	0.0006 0.9

We investigate the behavior of positional vectors under direct extrapolation scenarios. Specifically, we evaluate LLMs on sequences twice the length of their maximum context window, comparing the variations in PPL and positional vectors. The results, as shown in Table 7, indicate a sharp

10338

increase in PPL once the context window is exceeded. Additionally, the similarity between positional vectors within and beyond the context window decreases, highlighting the relationship between length extrapolation and the stability of positional representations.

Furthermore, we assess the impact of OOD positional vectors. We analyze the attention scores of initial tokens and the positional vectors of output logits within and beyond the context window. As demonstrated in Table 8, exceeding the context window results in the loss of the attention sinking property, and the positional vectors of the logits exhibit a different distribution compared to those within the context window. This suggests that OOD positional vectors influence the model's distribution. Moreover, OOD positional encodings play a crucial role in large language models (LLMs) that leverage RoPE [15].

Experiments of Positional Vector Replacement

To assess whether our positional vector replacement method is specific to TL-NoPE, we applied it to TL-NoPE-new, which was trained from scratch. We evaluated its language modeling performance on the PG-19 dataset, using the same experimental setup as TL-NoPE. As shown in Table 9, our approach effectively extends the context window to 8K tokens and achieves performance comparable to attention scaling.

Table 9: Results of language modeling in PG-19 with TL-NoPE-new								
Interpolation Method	Factor	2048	4096	6144	8192			
-	-	29.0	110.7	634.8	1078.5			
Attention Scaling	1.2 1.5	29.0 36.8	27.3 41.6	- 44.4	46.2			
Positional Vector Replacement(ours)	$r=2.5, \alpha=0.8$ $r=5, \alpha=0.8$	32.8 57.4	29.5 50.8	40.0	- 44.2			

Visualization of Positional Vectors

We visualize all models listed in Table 1 with PCA. The subsequent figures present the visualization of TL-NoPE, TL-RoPE, TL-ALiBi, TL-Window, TL-Window-RoPE, and TL-Window-80.

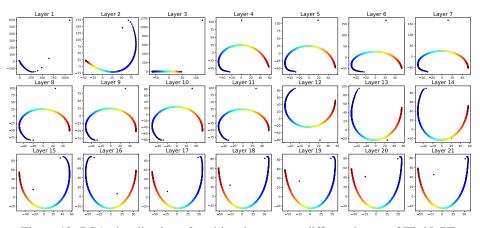


Figure 12: PCA visualization of positional vectors at different layers of TL-NoPE.

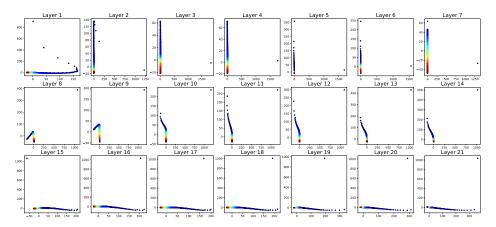


Figure 13: PCA visualization of positional vectors at different layers of TL-RoPE.

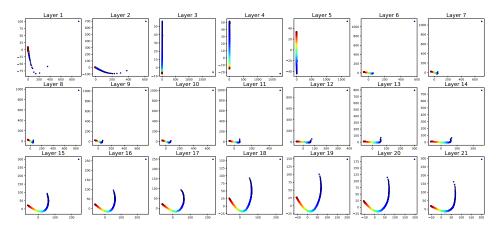


Figure 14: PCA visualization of positional vectors at different layers of TL-ALiBi.

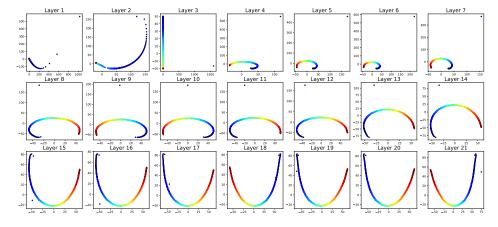


Figure 15: PCA visualization of positional vectors at different layers of TL-Window.

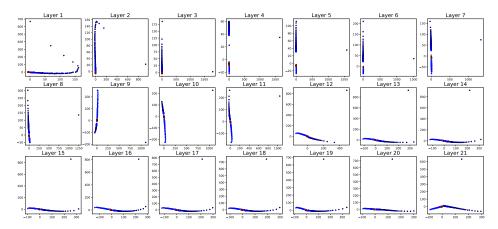


Figure 16: PCA visualization of positional vectors at different layers of TL-Widow-RoPE.

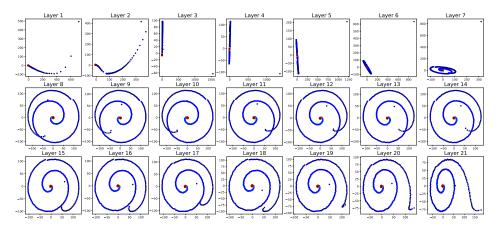


Figure 17: PCA visualization of positional vectors at different layers of TL-Window-80.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We provide the main claims in abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitation in Appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We have provided proof in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have provide the information for reproduce the results.

Guidelines:

• The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We only provide the Pytorch implementation of our methods in the Appendix.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

 Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have provided the training details in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The paper does not provide this.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the computer resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The paper respect with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

• The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite existing papers and url.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [No]

Justification: The models trained by us can not be launched. However, we give the details of training these models.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent)
 may be required for any human subjects research. If you obtained IRB approval, you
 should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.