

---

# When Attention Sink Emerges in Language Models: An Empirical View

---

Xiangming Gu<sup>1,2\*</sup>, Tianyu Pang<sup>1†</sup>, Chao Du<sup>1</sup>, Qian Liu<sup>1</sup>, Fengzhuo Zhang<sup>1,2</sup>,  
Cunxiao Du<sup>1</sup>, Ye Wang<sup>2†</sup>, Min Lin<sup>1</sup>

<sup>1</sup>Sea AI Lab, Singapore <sup>2</sup>National University of Singapore

{guxm, tianyupang, duchao, liuqian, zhangfz, ducx, linmin}@sea.com;  
wangye@comp.nus.edu.sg

## Abstract

Language Models (LMs) assign significant attention to the first token, even if it is not semantically important, which is known as **attention sink**. This phenomenon has been widely adopted in applications such as streaming/long context generation, KV cache optimization, inference acceleration, model quantization, and others. Despite its widespread use, a deep understanding of attention sink in LMs is still lacking. In this work, we first demonstrate that attention sinks exist universally in LMs with various inputs, even in small models. Furthermore, attention sink is observed to emerge during the LM pre-training, motivating us to investigate how *optimization*, *data distribution*, *loss function*, and *model architecture* in LM pre-training influence its emergence. We highlight that attention sink emerges after effective optimization on sufficient training data. The sink position is highly correlated with the loss function and data distribution. Most importantly, we find that attention sink acts more like key biases, *storing extra attention scores*, which could be non-informative and not contribute to the value computation. We also observe that this phenomenon (at least partially) stems from tokens' inner dependence on attention scores as a result of softmax normalization. After relaxing such dependence by replacing softmax attention with other attention operations, such as sigmoid attention without normalization, attention sinks do not emerge in LMs up to 1B parameters.

## 1 Introduction

Xiao et al. [35] showed that Large Language models (LLMs) allocate significant attention to the initial tokens, irrespective of their semantic relevance. This interesting phenomenon is termed as **attention sink** and has widespread applications, including streaming/long context generation [35, 12, 36], KV cache optimization [11, 32, 34], efficient inference [40, 5], model quantization [20, 15], and others.

A seminal of works attempted to understand attention sink. Among them, Cancedda [4] clarified that attention sink primarily appears only on the first token. They found that early FFNs in LLaMA2 blast off the large norm of hidden states of the first token, thus leading to the attention sink in later layers. This is referred to as *massive activations* (very few activations exhibit extremely large values compared to others) in Sun et al. [29]. Besides, Sun et al. [29], Yu et al. [37] observed that attention sink may also appear in several word tokens carrying limited semantic information and having no fixed position. Despite the above research efforts, a deep understanding of attention sink is still absent. Therefore, we conduct a comprehensive study to investigate when attention sink emerges.

Based on open-sourced LMs, we show that attention sink universally exists in auto-regressive LMs across different inputs, even in the small models or with random token sequences. Additionally,

---

\*Work done during Xiangming Gu's internship at Sea AI Lab.

†Correspondence to Tianyu Pang and Ye Wang.

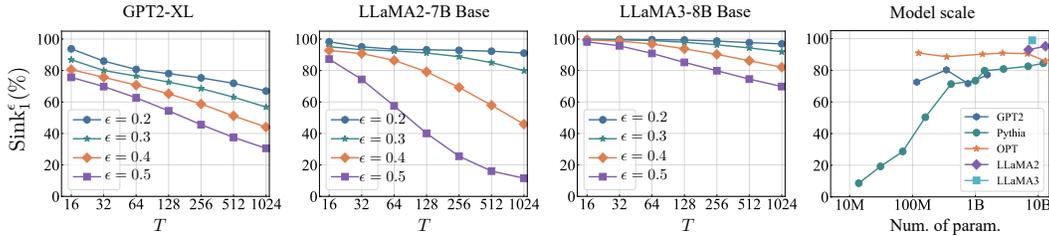


Figure 1: From left to right (A)-(D). The metric  $\text{Sink}_1^\epsilon$  (averaged on 100 sequences) tends to decrease with larger token lengths  $T$  in (A) GPT2-XL; (B) LLaMA2-7B Base; (C) LLaMA-8B Base. This tendency becomes more obvious with larger  $\epsilon$ . (D) Attention sink emerges in small LMs.

attention sink is observed to emerge during the LM pre-training before continual instruction tuning [22]. This motivates us to focus on the LM pre-training, whose objective can be formulated as:

$$\min_{\theta} \mathbb{E}_{\mathbf{X} \sim p_{\text{data}}} [\mathcal{L}(p_{\theta}(\mathbf{X}))]. \quad (1)$$

In the remaining part of this paper, we investigate how the **optimization** (Section 4), **data distribution** (Section 5), **loss function** (Section 6), and **model architecture** (Section 7) influence the emergence of attention sink. We have the following conclusions:

- Attention sink emerges after LMs are trained effectively on sufficient training data. It appears less obvious in LMs trained with small learning rates. While weight decay encourages the emergence of attention sink.
- The sink position is highly related to the loss function and data distribution and can be shifted to other positions rather than the first token.
- Attention sink acts more like key biases, storing extra attention and meanwhile not contributing to the value computation. This phenomenon (at least partially) stems from tokens’ inner dependence on attention scores due to the softmax normalization. After relaxing such dependence by replacing softmax attention with other attention operations, e.g., sigmoid attention without normalization, attention sinks do not emerge in LMs up to 1B parameters.

## 2 Preliminaries on LMs (full version in Appendix A)

**Auto-regressive modeling.** Let  $f_{\theta}$  be an LM with  $L$  transformer decoder blocks and  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$  are the input tokens, where each token belongs to the tokenizer vocabulary  $\mathbb{V}$ . The LM output is also a sequence  $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T\} = f_{\theta}(\mathbf{X})$ , where  $\mathbf{y}_t$  represents the predicted logits of  $p(\mathbf{x}_{t+1} | \mathbf{x}_{<t})$ . The pre-training objective of LMs is to maximize the likelihood of input data. Given a large corpus containing many documents, a packing strategy is adopted: documents are concatenated together and chunked into sequences with the context length  $C$ . Then the empirical loss function for each chunk is  $\mathcal{L} = \sum_{t=2}^C \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t})$ . We note that  $p_{\theta}(\mathbf{x}_1)$  is ignored since  $\mathbf{y}_1 = f_{\theta}(\mathbf{x}_1)$  is the prediction for the next token  $\mathbf{x}_2$ .

**Attention in LMs.** In the  $l$ -th multi-head self-attention (MHSA) layer, the input  $\mathbf{H}^{l-1}$  are first transformed into keys, queries, and values:  $\mathbf{K}^{l,h} = \mathbf{H}^{l-1} \mathbf{W}_K^{l,h}$ ,  $\mathbf{Q}^{l,h} = \mathbf{H}^{l-1} \mathbf{W}_Q^{l,h}$ ,  $\mathbf{V}^{l,h} = \mathbf{H}^{l-1} \mathbf{W}_V^{l,h}$  for each head  $1 \leq h \leq H$  (we omit the notation of LN when considering pre-norm design for simplicity). Here  $\mathbf{W}_K^{l,h}, \mathbf{W}_Q^{l,h}, \mathbf{W}_V^{l,h} \in \mathbb{R}^{d \times d_h}$ ,  $d_h = d/H$ . Then the attention output is computed as

$$\mathbf{A}^{l,h} = \text{Softmax}(\mathbf{Q}^{l,h} \mathbf{K}^{l,hT} / \sqrt{d_h} + M), \quad \mathbf{O}^l = \text{Concat}_{h=1}^H (\mathbf{A}^{l,h} \mathbf{V}^{l,h}) \mathbf{W}_O^l, \quad (2)$$

where  $M \in \mathbb{R}^{T \times T}$  is an attention mask, where  $M_{ij} = -\infty$  for  $i < j$  and  $M_{ij} = 0$  for  $i \geq j$ . Xiao et al. [35] revealed that LLMs allocate significant attention scores to specific token positions, e.g.  $\mathbf{x}_1$ , resulting in “vertical” attention patterns. This holds even when  $\mathbf{x}_1$  is not a BOS token [35, 29].

## 3 Measuring attention sink

**Threshold-based metrics.** Xiao et al. [35] showcased the appearance of attention sink by visualizing attention logits/scores in different heads/blocks. This leads to the intractability of measuring attention sink quantitatively due to the large number of attention heads and blocks. Therefore, we first explore the metrics to measure the attention sink. Within each head, we compute the importance scores for the  $k$ -th token  $\alpha_k^{l,h} = \frac{1}{T-k+1} \sum_{i=k}^T \mathbf{A}_{i,k}^{l,h}$ . We mainly focus on the first token  $\alpha_1^{l,h}$ . It is noted that

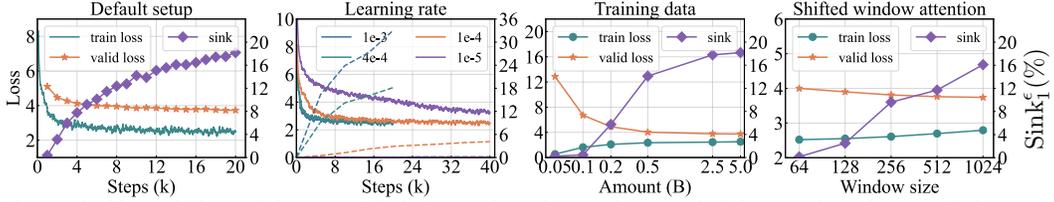


Figure 2: From left to right (A)-(D), left y-axis refers to loss and right y-axis refers to  $\text{Sink}_1^\epsilon$ . (A) Train/valid loss and  $\text{Sink}_1^\epsilon$  during LM pre-training under the default setup. (B) Training loss (solid lines)/attention sink (dashed lines) dynamics of LMs using different learning rates. Train loss/valid loss and  $\text{Sink}_1^\epsilon$  of LMs trained with (C) different amount of training data and (D) different window size.

$\frac{1}{T} \leq \alpha_1^{l,h} \leq 1$  since  $A_{1,1}^{l,h} = 1$  and  $0 \leq A_{i \neq 0,1}^{l,h} \leq 1$ . Then we adopt a threshold-based metric, we consider a head has attention sink in the first token if  $\alpha_1^{l,h} > \epsilon$ . Considering that the whole model has  $L$  blocks and each block has  $H$  heads, we use the following metric to measure the attention sink of the whole LM:  $\text{Sink}_k^\epsilon = \frac{1}{L} \sum_{l=1}^L \frac{1}{H} \sum_{h=1}^H \mathbb{I}(\alpha_k^{l,h} > \epsilon)$ .

**Selections of thresholds.** Generally, a larger  $\epsilon$  represents a strict definition for attention sink. There is no principal way to find an optimal threshold and we only use this metric to quantify the emergence of attention sink empirically. Based on Figure 1(A)-(C), we prefer to select a threshold that is both strict in quantifying attention sink and less sensitive to the token length  $T$ . This gives us the selection of  $\epsilon = 0.3$ . For fair comparisons, we need to fix  $T$  when computing the metric, e.g.,  $T = 64$ .

**Attention sink in open-sourced LMs.** In Appendix C, we first show that the first token acts as biases: the angles between the first key and queries of other tokens are typically small, leading to attention sink. Then we show that attention sink appears across different data domains, even on random token sequences. Besides, we find that instruction tuning has an insignificant impact on attention sink, which motivates us to focus on the LM pre-training. As visualized in Figure 1(D), attention sink emerges in small LMs, even in Pythia-14M.

#### 4 Effects of optimization on attention Sink.

We pre-train a series of LLaMA models to conduct our experiments. Due to the intractability of replicating LLaMA pre-training, we design small-sized models. Following Liu et al. [19], we set hidden dimension  $d = 768$ , block number  $L = 10$ , head number  $H = 8$ , intermediate size of FFN as 1536, resulting in approximately 60M parameters except for word embeddings and unembeddings. We keep the other design the same as LLaMA2 models.

For data distribution, we sample 5B tokens from the Pile dataset [10]. We set the context length to 2048 tokens, the batch size to 1M tokens, and the training step to 20k (including 100 steps for warm-up). We adopt a learning rate of  $4e-4$  with cosine scheduling. The optimizer is AdamW [21] with a weight decay ratio of 0.1. We use the Pile-CC validation loss [10, 19] to measure the model performance and sample 100 sequences with  $T = 64$  (no BOS) out of training data to measure  $\text{Sink}_k^\epsilon$  with  $\epsilon = 0.3$ .

**Optimization steps.** As visualized in Figure 2(A), under our default setup, attention sink emerges after certain optimization steps, e.g., between 1k and 2k steps. With the progression of pre-training, attention sink becomes more obvious.

**Learning rate.** With a smaller learning rate, it takes longer training steps to lower training loss, as present in Figure 2(B). Meanwhile, the emergence of attention sink is also delayed. Besides, we also find that a smaller learning rate ( $1e-4$ ) results in LMs with less obvious attention sink. But further decreasing learning rate significantly affects the optimization and model performance, thus affecting the emergence of attention sink.

**Batch size.** In Appendix D, we find that only modifying batch size has no effects on attention sink.

**Takeaways:** 1. Attention sink emerges after LMs are trained effectively. 2. Attention sink appears less obvious in LMs trained with small learning rates.

#### 5 Effects of data distribution $p_{\text{data}}$ on attention sink

**Training data amount.** In the default setup, we consider 5B tokens. We wonder whether the attention sink emerges if we further constrain the data within a fixed compute budget. Therefore, we constrain the training data to 5B, 2.5B, 500M, 200M, 100M, and 50M. Meanwhile, we fix the

Table 1: Larger weight decay ratios tend to induce more attention sink heads in LMs. But much larger values hurt the model performance and attention sink disappears.

$\gamma$	0.0	0.001	0.01	0.1	0.5	1.0	2.0	5.0
Sink $^c_1$ (%)	15.20	15.39	15.23	18.18	41.08	37.71	6.13	0.01
valid loss	3.72	3.72	3.72	3.73	3.80	3.90	4.23	5.24

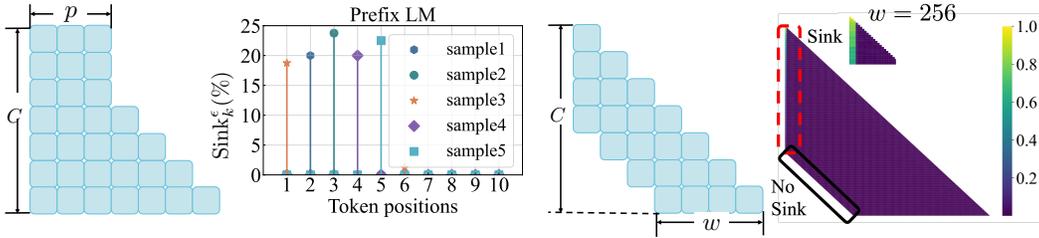


Figure 3: From left to right (A)-(D). (A) Attention pattern for prefix language modeling. (B) Attention sink appears among the prefix tokens rather than the first token. (C) Shifted window attention pattern. (D) Attention sink appears on the first token, but not on the “first token” within each window.

batch size and optimization steps. As visualized in Figure 2(C), with less training data, the trained LLMs tend to overfit. Meanwhile, attention sink also disappears.

**Randomness in data distribution.** After packing documents into chunks, we re-sample the first token within the chunk  $\mathbf{x}_1 \sim \text{Uniform}(\mathbb{V})$ . The trained LM has the metric  $\text{Sink}_1^c = 27.03\%$ , even larger than the default setup. This further validates the low semantic information of the sink token. We also consider  $\mathbf{x}_1, \mathbf{x}_2 \sim \text{Uniform}(\mathbb{V})$ , and we find attention sink shifts to the second token with  $\text{Sink}_2^c = 14.08\%$  while the attention sink on the first token is much less obvious  $\text{Sink}_1^c = 1.98\%$ . But when only sample  $\mathbf{x}_2 \sim \text{Uniform}(\mathbb{V})$ , the attention sink still always appears on the first token ( $\text{Sink}_1^c = 20.99\%$ ). Additionally, we find with more random tokens during pre-training, attention sink tends to disappear.

**Fixing token in a specific position.** Xiao et al. [35] considered a learnable token in the first token position within each chunk, which can be considered as  $\mathbf{x}_1 \sim \mathbb{I}(\mathbf{x} = \mathbf{x}_{\text{fix}})$ . We also consider fixing the token  $\mathbf{x}_{\text{fix}}$  in the second/third token position during pre-training. Consequently, the attention sink always appears in the fixed token instead of the first token, as shown in Table 7.

**Takeaways:** 1. Attention sink emerges after LMs are trained on sufficient training data. 2. Attention sink could be shifted to other positions rather than the first token if modifying  $p_{\text{data}}$ .

## 6 Effects of loss function $\mathcal{L}$ on attention sink

**Weight decay.** The loss function becomes  $\mathcal{L} = \sum_{t=2}^C \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t}) + \gamma \|\theta\|_2^2$  when introducing weight decay ratio  $\gamma$ . As indicated in Table 1, even  $\gamma = 0$  in the loss function, attention sink still emerges in LMs. Then a larger  $\gamma$  encourages more heads to have attention sink. But further increasing weight decay hurts the optimization, leading to less obvious or even no attention sink.

**Prefix language modeling.** Since the first token is not predicted in the auto-regressive loss function, it could be considered as the prefix token. Then the original auto-regressive loss can be generalized into the formula  $\mathcal{L} = \sum_{t=p+1}^C \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{p+1:t-1}, \mathbf{x}_{1:p})$ , with the prefix length  $p = 1$ . Motivated by Wang et al. [33], we consider  $p > 1$  and the casual mask visualized in Figure 3(A). Although this design does not affect the emergence of attention sink, it shifts the sink position. In Figure 3(B), the attention sink only appears on one token. But it appears among these prefix tokens instead of on the first token only. Massive activations also appear on the corresponding sink token.

**Shifted window attention.** Motivated by the shifted window attention adopted in Mistral-7B, we further explore the effects of window size on attention sink. With shifted window attention, the loss function becomes  $\mathcal{L} = \sum_{t=2}^C \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{t-w:t-1})$ , where  $w$  refers to the window size. As shown in Figure 3(C) and (D), with shifted window attention, we find that if  $t \leq w$ , the  $t$ -th token can still “look at” the first token, and LMs still have attention sink on the first token. When  $t > w$ , the  $t$ -th token can only attend up to the  $t - w + 1$ -th token. Although this token is the “first token” for the  $t$ -th token, typically it has no attention sink. We have similar observations in Mistral-7B. Additionally, from Figure 2(D), smaller window size prevents the emergence of attention sink.

Table 2: Positional embedding does not affect the emergence of attention sink.

PE	NoPE	Absolute PE	Learnable PE	Relative PE	ALiBi	Rotary
Sink <sub>1</sub> <sup>ε</sup> (%)	20.35	32.73	33.13	35.53	20.78	18.18
valid loss	3.81	3.74	3.79	5.45	3.71	3.73

Table 3: With comparable performance, LMs with sink token, KV biases, and K biases could shift attention sink from the first token to key biases’ position. Value biases cannot affect attention sink.

Attention in each head	Sink <sub>*</sub> <sup>ε</sup> (%)	Sink <sub>1</sub> <sup>ε</sup> (%)	valid loss
$\text{Softmax} \left( \frac{1}{\sqrt{d_h}} \mathbf{Q}^{l,h} \mathbf{K}^{l,hT} \sqrt{d_h} + \mathbf{M} \right) \mathbf{V}^{l,h}$	-	18.18	3.73
$\text{Softmax} \left( \frac{1}{\sqrt{d_h}} \begin{bmatrix} \mathbf{q}^{*,l,h} \\ \mathbf{Q}^{l,h} \end{bmatrix} \begin{bmatrix} \mathbf{k}^{*,l,hT} & \mathbf{K}^{l,hT} \end{bmatrix} + \mathbf{M} \right) \begin{bmatrix} \mathbf{v}^{*,l,h} \\ \mathbf{V}^{l,h} \end{bmatrix}$	74.12	0.00	3.72
$\text{Softmax} \left( \frac{1}{\sqrt{d_h}} \mathbf{Q}^{l,h} \begin{bmatrix} \mathbf{k}^{*,l,hT} & \mathbf{K}^{l,hT} \end{bmatrix} + \mathbf{M} \right) \begin{bmatrix} \mathbf{v}^{*,l,h} \\ \mathbf{V}^{l,h} \end{bmatrix}$	72.76	0.04	3.72
$\text{Softmax} \left( \frac{1}{\sqrt{d_h}} \mathbf{Q}^{l,h} \begin{bmatrix} \mathbf{k}^{*,l,hT} & \mathbf{K}^{l,hT} \end{bmatrix} \sqrt{d_h} + \mathbf{M} \right) \begin{bmatrix} \mathbf{0} \\ \mathbf{V}^{l,h} \end{bmatrix}$	73.34	0.00	3.72
$\text{Softmax} \left( \frac{1}{\sqrt{d_h}} \mathbf{Q}^{l,h} \mathbf{K}^{l,hT} + \mathbf{M} \right) \mathbf{V}^{l,h} + \mathbf{v}^{*,l,h}$	-	17.53	3.73

**Takeaways:** 1. Weight decay encourages the emergence of attention sink. 2. With prefix language modeling, attention sink appears among the prefix tokens rather than the first token only. 3. With shifted window attention, attention sink appears on the “absolute”, not the “relative” first token. Smaller window size prevents the emergence of attention sink.

## 7 Effects of model architecture $p_\theta$ on attention sink

Besides the following main factors, we also explore activation functions in the FFN (Appendix D.2), pre-norm or post-norm structure (Appendix D.3), multi-head design (including concatenation operation and number of heads, Appendix D.4). These designs do not affect the emergence of attention sink.

### 7.1 Positional embedding

Attention sink always appears on the first token, which motivates us to explore where such a position property is brought by positional embedding (PE). Then we replace the original Rotary with other PEs, as shown in Table 2. The detailed formulations for these PEs are in Appendix B. From the same Table, we observe all these LMs, even the one without explicit PE (NoPE), have attention sink.

### 7.2 Attention biases

**Implicit biases in attention.** In Appendix C, we show that  $\mathbf{x}_1$  acts as a bias: its key  $\mathbf{k}_1^{l,h}$  is distributed in a different manifold and its value  $\mathbf{v}_1^{l,h}$  has small  $\ell_2$ -norm. Xiao et al. [35] considered a learnable sink token before the input tokens during LM pre-training. As this token is fixed in the first position, this could be considered as implicitly introducing biases  $\mathbf{k}^{*,l,h}$ ,  $\mathbf{v}^{*,l,h}$ ,  $\mathbf{q}^{*,l,h}$  in attention, as shown in the second row in Table 3. As a result, attention sink appears on  $\mathbf{x}^*$  instead of  $\mathbf{x}_1$ .

**Explicit biases in attention.** Sun et al. [29] directly introducing  $\mathbf{k}^{*,l,h}$  and  $\mathbf{v}^{*,l,h}$  as learnable parameters in attention (*KV biases*). They found that this design could alleviate the massive activations. Considering the important role of  $\mathbf{k}^{*,l,h}$  and small  $\ell_2$ -norm of  $\mathbf{v}^{*,l,h}$ , we propose introducing only the key biases  $\mathbf{k}^{*,l,h}$  and fix value biases  $\mathbf{v}^{*,l,h} = \mathbf{0}$  (*K biases*). As a control, we also consider only adding value biases (*V biases*). The formulations of all these attention designs are shown in Table 3.

**LMs need key biases.** After evaluating the LMs with setups in Table 3, we first observe that these LMs have comparable model performance. Moreover, as long as there are key biases  $\mathbf{k}^{*,l,h}$ , attention sink disappears on the first token but on the biases. From the setup of K biases, we reaffirm that the sink token acts as key biases, storing extra attention scores, which could be completely non-informative and not contribute to the value computation. This is different from registers [6], which aggregates global information, in Vision Transformers (ViTs) [8]. In Figure 4(A), we also find that the LM with a learnable sink token has massive activations in  $\mathbf{x}^*$ . While LMs with KV biases and K biases have no massive activations. In Appendix D.5, we conduct more explorations.

Table 4: Normalization and selections of kernels in attention significantly affect the emergence of the attention sink. We use “\*” to mark that the metric  $\text{Sink}_1^\epsilon$  is computed by proxy attention scores.

$\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))$	$\mathbf{Z}_i$	$\text{Sink}_1^\epsilon$ (%)	loss
$\exp(\mathbf{q}_i^T \mathbf{k}_j / \sqrt{d_h})$	$\sum_{j'=1}^i \exp(\mathbf{q}_i \mathbf{k}_{j'}^T / \sqrt{d_h})$	18.18	3.73
$\text{sigmoid}(\mathbf{q}_i \mathbf{k}_j^T / \sqrt{d_h})$	1	0.44*	3.70
$\text{sigmoid}(\mathbf{q}_i \mathbf{k}_{j'}^T / \sqrt{d_h})$	$\sum_{j'=1}^i \text{sigmoid}(\mathbf{q}_i \mathbf{k}_{j'}^T / \sqrt{d_h})$	30.24	3.74
$\text{elu}(\mathbf{q}_i \mathbf{k}_j^T / \sqrt{d_h}) + 1$	1	0.80*	3.69
$(\text{elu}(\mathbf{q}_i) + 1)(\text{elu}(\mathbf{k}_j) + 1)^T$	$\sum_{j'=1}^i (\text{elu}(\mathbf{q}_i) + 1)(\text{elu}(\mathbf{k}_{j'}) + 1)^T$	53.65*	4.19
$\mathbf{q}_i \mathbf{k}_j^T / \sqrt{d_h}$	1	0.00*	3.99
$\text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_j)^T / \sqrt{d_h}$	$\max( \sum_{j'=1}^i \text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_{j'})^T / \sqrt{d_h} , 1)$	0.19*	3.85
$\text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_j)^T / \sqrt{d_h}$	1	0.74*	3.91

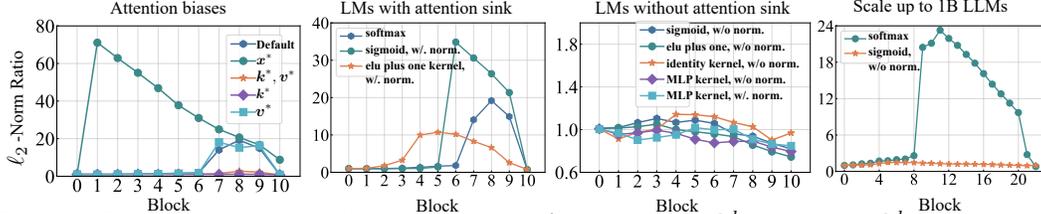


Figure 4: From left to right (A)-(D). We visualize the  $\ell_2$ -norm ratio of  $\mathbf{h}_1^l$  and mean of  $\mathbf{h}_{t \neq 0}^l$  in (A) LMs trained with various attention biases. (B) LMs that have attention sink when modifying attention operation. (C) LMs that have no attention sink when modifying attention operation. (D) When scaling the model size to 1B, LLMs with sigmoid attention (no normalization) still have no massive activations.

### 7.3 Attention operation

**General formulation of attention.** In the last section, we realize that LMs with softmax attention need key biases to save extra attention. This motivates us to explore whether such a property is related to the dependence among attention scores due to the softmax. First, the attention output for the  $i$ -th token can be generalized into:  $\mathbf{v}_i^\dagger = \left( \sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'})) \right)^{-1} \sum_{j=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j)) \mathbf{v}_j = \mathbf{Z}_i^{-1} \sum_{j=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j)) \mathbf{v}_j$ , where we omit the positional embedding, and block/head indexes for simplicity.  $\mathbf{Z}_i$  is a normalization term and  $\varphi(\cdot)$  is a kernel function.

**Normalization.** In Appendix D.6, we first show that only scaling the normalization term  $\mathbf{Z}_i \rightarrow \mathbf{Z}_i / \alpha$  with  $\alpha < 1$  results in less obvious attention sink but does not stop its emergence. So we consider removing the normalization. Since the exponential function in softmax tends to explode without normalization, we replace it with sigmoid or elu plus one. When evaluating the attention sink, we compute the *proxy* attention scores by using the term  $\mathbf{Z}_i = \sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))$  for attention sink metric  $\text{Sink}_1^\epsilon$ . As shown in Table 4, without normalization, LMs still have comparable validation loss but no attention sink. With normalization, attention sink also emerges in LMs with sigmoid attention.

**Kernel functions.** Motivated by linear attention [17], we consider different kernel functions  $\varphi(\cdot)$ , including elu plus one, identity, and MLP. It is noted that  $\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))$  could be minus for identity and MLP kernels. This brings intrinsic difficulty for normalization during the training and calculation of  $\text{Sink}_1^\epsilon$ . For normalization in the training, we consider  $\mathbf{Z}_i = \max(|\sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))|, 1)$ . When computing  $\text{Sink}_1^\epsilon$ , we consider the proxy attention scores  $|\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))| / \sum_{j'=1}^i |\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))|$ . From Table 4 (a full version in Table 13), we find that the LM with MLP kernel have no attention sink with or without normalization.

**Inner dependence on attention scores.** We note that the LMs with no attention sink typically relax tokens’ inner dependence on attention scores. Their attention scores during pre-training could be negative or not add up to one. This indicated that attention sink (at least partially) stems from such inner dependence. Besides the attention metric computed by proxy attention scores, we also observe that the above LMs also have no massive activations, as shown in Figure 4(C).

**Scale up to 1B parameters.** We compare model behaviors of 1B LMs with softmax attention and sigmoid attention (without normalization). Specifically, the latter achieves a validation loss of 3.10, slightly larger than the 3.07 achieved by the former. However, the attention sink metric significantly

drops from 45.11% to near zero:  $\text{Sink}_1^\epsilon = 2.46\%$  using the proxy attention scores. Meanwhile, as visualized in Figure 4(D), LLMs with sigmoid attention have no massive activations.

**Takeaways:** 1. Positional embedding, FFN design, LN location, and multi-head design do not affect the emergence of attention sink. 2. Attention sink acts more like key biases, storing extra attention and meanwhile not contributing to the value computation. 3. When relaxing tokens’ inner dependence on attention scores, attention sink does not emerge in LMs.

## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [2] Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR, 2023.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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. *Advances in neural information processing systems*, 33, 2020.
- [4] Nicola Cancedda. Spectral filters, dark signals, and attention sinks. *arXiv preprint arXiv:2402.09221*, 2024.
- [5] Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. *arXiv preprint arXiv:2403.06764*, 2024.
- [6] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
- [8] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [9] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [10] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [11] Suyu Ge, Yunan Zhang, Liyuan Liu, Minjia Zhang, Jiawei Han, and Jianfeng Gao. Model tells you what to discard: Adaptive kv cache compression for llms. *arXiv preprint arXiv:2310.01801*, 2023.
- [12] Chi Han, Qifan Wang, Hao Peng, Wenhan Xiong, Yu Chen, Heng Ji, and Sinong Wang. Lm-infinite: Zero-shot extreme length generalization for large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3991–4008, 2024.

- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [15] Wei Huang, Haotong Qin, Yangdong Liu, Yawei Li, Xianglong Liu, Luca Benini, Michele Magno, and Xiaojuan Qi. Slim-llm: Saliency-driven mixed-precision quantization for large language models. *arXiv preprint arXiv:2405.14917*, 2024.
- [16] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- [17] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020.
- [18] Amirhossein Kazemnejad, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Payel Das, and Siva Reddy. The impact of positional encoding on length generalization in transformers. *Advances in Neural Information Processing Systems*, 36, 2024.
- [19] Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. Regmix: Data mixture as regression for language model pre-training. *arXiv preprint arXiv:2407.01492*, 2024.
- [20] Ruikang Liu, Haoli Bai, Haokun Lin, Yuening Li, Han Gao, Zhengzhuo Xu, Lu Hou, Jun Yao, and Chun Yuan. Intactkv: Improving large language model quantization by keeping pivot tokens intact. *arXiv preprint arXiv:2403.01241*, 2024.
- [21] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [22] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- [23] Ofir Press, Noah A Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021.
- [24] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [25] 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. *Journal of machine learning research*, 21(140):1–67, 2020.
- [26] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [27] Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.
- [28] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- [29] Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. Massive activations in large language models. *arXiv preprint arXiv:2402.17762*, 2024.
- [30] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

- [31] 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.
- [32] Zhongwei Wan, Ziang Wu, Che Liu, Jinfa Huang, Zhihong Zhu, Peng Jin, Longyue Wang, and Li Yuan. Look-m: Look-once optimization in kv cache for efficient multimodal long-context inference. *arXiv preprint arXiv:2406.18139*, 2024.
- [33] Thomas Wang, Adam Roberts, Daniel Hesslow, Teven Le Scao, Hyung Won Chung, Iz Beltagy, Julien Launay, and Colin Raffel. What language model architecture and pretraining objective works best for zero-shot generalization? In *International Conference on Machine Learning*, pages 22964–22984. PMLR, 2022.
- [34] Haoyi Wu and Kewei Tu. Layer-condensed kv cache for efficient inference of large language models. *arXiv preprint arXiv:2405.10637*, 2024.
- [35] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023.
- [36] Shuai Yang, Yuying Ge, Yang Li, Yukang Chen, Yixiao Ge, Ying Shan, and Yingcong Chen. Seed-story: Multimodal long story generation with large language model. *arXiv preprint arXiv:2407.08683*, 2024.
- [37] Zhongzhi Yu, Zheng Wang, Yonggan Fu, Huihong Shi, Khalid Shaikh, and Yingyan Celine Lin. Unveiling and harnessing hidden attention sinks: Enhancing large language models without training through attention calibration. *arXiv preprint arXiv:2406.15765*, 2024.
- [38] Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [39] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [40] Zhenyu Zhang, Shiwei Liu, Runjin Chen, Bhavya Kailkhura, Beidi Chen, and Atlas Wang. Q-hitter: A better token oracle for efficient llm inference via sparse-quantized kv cache. *Proceedings of Machine Learning and Systems*, 6:381–394, 2024.

## A Detailed formulations of LMs

Let  $f_\theta$  be an auto-regressive LM with  $L$  transformer decoder blocks and  $\mathbf{X} \in \mathbb{R}^{T \times |\mathbb{V}|} := \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$  are the input tokens, where each token  $\mathbf{x}_t$  is a one-hot encoding and  $|\mathbb{V}|$  is the vocabulary size of tokenizer  $\mathbb{V}$ . The LM output is also a sequence  $\mathbf{Y} \in \mathbb{R}^{T \times |\mathbb{V}|} := \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T\} = f_\theta(\mathbf{X})$ , where  $\mathbf{y}_t$  represents the predicted logits of  $p(\mathbf{x}_{t+1} | \mathbf{x}_{\leq t})$ .

**Transformer blocks.** In the forward pass,  $\mathbf{X}$  is first embedded as  $\mathbf{H}^0 \in \mathbb{R}^{T \times d} := \mathbf{X}\mathbf{W}_E + \mathbf{P}$ , where  $\mathbf{W}_E \in \mathbb{R}^{|\mathbb{V}| \times d}$  is the learnable word embedding,  $\mathbf{P} \in \mathbb{R}^{T \times d}$  is the positional embedding, and  $d$  is the hidden dimension. We denote  $\mathbf{H}^l \in \mathbb{R}^{T \times d} := \{\mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_T^l\}$ ,  $1 \leq l \leq L$  to be the output of the  $l$ -th block. Each block comprises a multi-head self-attention (MHSA) operation and a feed-forward network (FFN). The block has either a pre-norm or post-norm structure according to the location of layer normalization (LN) [1, 38]. Most of LLMs consider a pre-norm block, as also shown in Figure 5(Left):

$$\mathbf{H}^l = \text{FFN}(\text{LN}(\mathbf{O}^l + \mathbf{H}^{l-1})) + \mathbf{O}^l + \mathbf{H}^{l-1}, \quad \mathbf{O}^l = \text{MHSA}(\text{LN}(\mathbf{H}^{l-1})), \quad (3)$$

while the post-norm transformer block is

$$\mathbf{H}^l = \text{LN}(\text{FFN}(\text{LN}(\mathbf{O}^l + \mathbf{H}^{l-1})) + \text{LN}(\mathbf{O}^l + \mathbf{H}^{l-1})), \quad \mathbf{O}^l = \text{MHSA}(\mathbf{H}^{l-1}). \quad (4)$$

**MHSA layers.** In the MHSA layer, the input  $\mathbf{H}^{l-1}$  are first transformed into keys, queries, and values:  $\mathbf{K}^{l,h} = \mathbf{H}^{l-1}\mathbf{W}_K^{l,h}$ ,  $\mathbf{Q}^{l,h} = \mathbf{H}^{l-1}\mathbf{W}_Q^{l,h}$ ,  $\mathbf{V}^{l,h} = \mathbf{H}^{l-1}\mathbf{W}_V^{l,h}$  for each head  $1 \leq h \leq H$  (we omit the notation of LN when considering pre-norm design for simplicity). Here  $\mathbf{W}_K^{l,h}, \mathbf{W}_Q^{l,h}, \mathbf{W}_V^{l,h} \in \mathbb{R}^{d \times d_h}$ ,  $d_h = d/H$ . Then the attention output is computed as

$$\mathbf{A}^{l,h} = \text{Softmax}(\mathbf{Q}^{l,h}\mathbf{K}^{l,hT} / \sqrt{d_h} + \mathbf{M}), \quad \mathbf{O}^l = \text{Concat}_{h=1}^H(\mathbf{A}^{l,h}\mathbf{V}^{l,h})\mathbf{W}_O^l, \quad (5)$$

where  $\mathbf{M} \in \mathbb{R}^{T \times T}$  is an attention mask. For vanilla causal attention,  $M_{ij} = -\infty$  for  $i < j$  and  $M_{ij} = 0$  for  $i \geq j$ . Finally, the output of final transformer block  $\mathbf{H}^L$  is fed into an unembedding layer for prediction:  $\mathbf{Y} = \text{LN}(\mathbf{H}^L)\mathbf{W}_{\text{cls}}$ , where  $\mathbf{W}_{\text{cls}} \in \mathbb{R}^{d \times |\mathbb{V}|}$ .

**Positional embedding.** NoPE [18] considered no explicit positional embedding (PE) in LMs, where  $\mathbf{P} = \mathbf{0}$ . When using absolute PE [31],  $\mathbf{P}$  is a periodic function of token positions. Devlin et al. [7], Brown et al. [3] adopted a learnable PE, which means  $\mathbf{P}$  is a learnable embedding of token positions. The dot product between each query and key meets  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i \mathbf{k}_j^T$  when using the above three PEs. While for relative PE [25], ALiBi [23], Rotary [28], they have  $\mathbf{P} = \mathbf{0}$ . Instead, they modify the dot product  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle$ . For relative PE and ALiBi,  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i \mathbf{k}_j^T + g(i-j)$ , where  $g(\cdot)$  is pre-defined function of the distance between two token positions. For Rotary,  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i \mathbf{R}_{\Theta, j-i} \mathbf{k}_j^T$ , where  $\mathbf{R}_{\Theta, (\cdot)}$  is a pre-defined rotation matrix. We include detailed formulations of the above PEs in Appendix B.

**Auto-regressive objective.** The pre-training objective of LMs is to maximize the likelihood of input data:  $\theta^* = \arg \max_{\theta} \mathbb{E}_{\mathbf{X} \sim p_{\text{data}}} \left[ \sum_{t=1}^T \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t}) \right]$ , where  $p_{\text{data}}$  refers to the data distribution.

**Packing documents in pre-training.** Given a large corpus  $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$ , where each  $d_i$  represents a document containing a sequence of tokens. A packing strategy is adopted in the LM pre-training, as present in Figure 5(Right). All documents are concatenated and chunked into sequences with a context length of  $C$ . Each chunk could start with any token within one document or the BOS/EOS token. Then the empirical loss function for each chunk is  $\mathcal{L} = \sum_{t=2}^C \log p_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t})$ . We note that  $p_{\theta}(\mathbf{x}_1)$  is ignored since  $\mathbf{y}_1 = f_{\theta}(\mathbf{x}_1)$  is the prediction for the next token  $\mathbf{x}_2$ .

**LM inference.** During the inference, a BOS token is fed into the model as the prefix for unconditional generation:  $\mathbf{x}'_t \sim p_{\theta}(\mathbf{x}'_t | \mathbf{x}'_{<t}, \mathbf{x}, \text{BOS})$ , where  $\mathbf{x}'_t$  is the  $t$ -th generated token,  $\mathbf{x}$  is the optional prompt. If there are no BOS tokens in the pre-training, the EOS token is considered as the BOS.

## B Detailed formulations of positional embedding

In this section, we provide detailed formulations of positional embedding (PE) in LMs. PEs could be classified into two categories. NoPE, absolute positional embedding, and learnable positional embedding belong to the same category since they are added to the initial hidden states:  $\mathbf{H}^0 =$

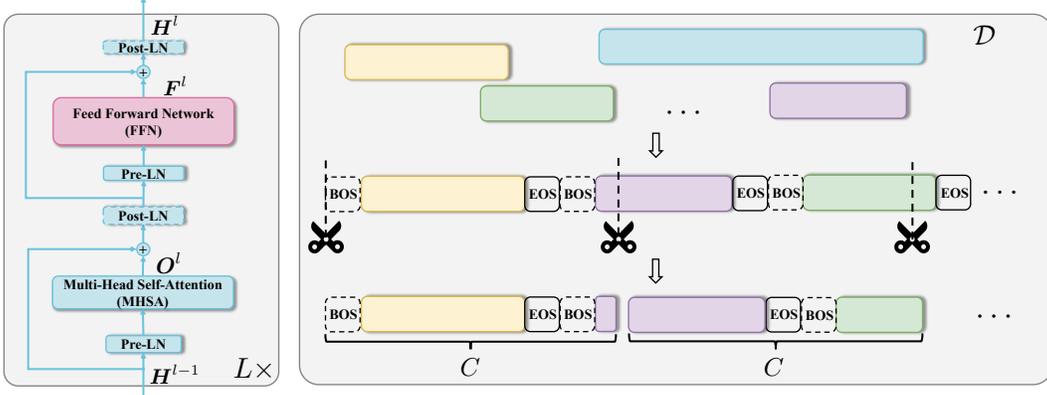


Figure 5: (Left) Architecture of pre-norm transformer block (we highlight the location of post-norm LN using dashed lines). We denote the output of MHSA as  $O^l$  and the output of FFN as  $F^l$ . (Right) The packing strategy in the LM pre-training. All documents are concatenated with BOS (optional) and EOS tokens as the boundaries. Then it is chunked into equal-sized sequences with context length  $C$ .

$\mathbf{XW}_E + \mathbf{P}$ . Here  $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_T\} \in \mathbb{R}^{T \times d}$ . Meanwhile, the dot product between each query and key is computed as  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i^T \mathbf{k}_j$ .

**NoPE.** NoPE [18] refers to no positional embedding. Therefore,  $\mathbf{P} = \mathbf{0}$ .

**Absolute PE.** Each position vector  $\mathbf{p}_t$  in absolute positional embedding is a periodic function of token position  $t$  following Vaswani et al. [31]:

$$\mathbf{p}_t = [\sin(\omega_1 t) \quad \cos(\omega_1 t) \quad \sin(\omega_2 t) \quad \cos(\omega_2 t) \quad \dots \quad \sin(\omega_{d/2} t) \quad \cos(\omega_{d/2} t)], \quad (6)$$

where  $\omega_i = 1/10000^{2(i-1)/d}$ .

**Learnable PE.** Each position vector  $\mathbf{p}_t$  in learnable positional embeddings is a learnable parameter.

Relative positional embedding, ALiBi, and rotary belong to another category since they consider the relative distance among tokens. Therefore, the initial hidden states are  $\mathbf{H}^0 = \mathbf{XW}_E$  but how to compute the dot product between each query and key is modified.

**Relative PE.** The relative positional embeddings are adopted in T5 models [25]. A bias term is adopted for the dot product:  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i \mathbf{k}_j^T + g(i - j)$ , where the definition for distance function  $g(\cdot)$  is:

$$g(i - j) = \begin{cases} i - j & \text{if } i - j < \mathcal{B}/2 \\ \frac{\mathcal{B}}{2} + \lfloor \frac{\log(\frac{i-j}{\mathcal{B}/2})}{\log(\frac{\mathcal{D}}{\mathcal{B}/2})} \rfloor \times \frac{\mathcal{B}}{2} & \text{if } \mathcal{B}/2 \leq i - j < \mathcal{D} \\ \mathcal{B} - 1 & \text{if } i - j \geq \mathcal{D} \end{cases} \quad (7)$$

Here  $\mathcal{B}$  and  $\mathcal{D}$  refer to the number of buckets and maximum distance, respectively. In T5 models,  $\mathcal{B} = 32$  and  $\mathcal{D} = 128$ .

**ALiBi.** Similarly, ALiBi [23] also adds a bias term to the dot product:  $\langle \mathbf{q}_i, \mathbf{k}_j \rangle = \mathbf{q}_i^T \mathbf{k}_j + g(i - j)$ , where  $g(i - j) = -(i - j) \cdot m$ .  $m$  is a head-specific slope fixed:

$$m = 2^{-h \cdot 2^{-\log_2 H + 3}}, \quad (8)$$

where  $1 \leq h \leq H$  is the head index and  $H$  is the number of heads in the multi-head self-attention (MHSA). This slope  $m$  is a geometric sequence. For instance, when  $H = 8$ , the sequence is  $\frac{1}{2^1}, \frac{1}{2^2}, \dots, \frac{1}{2^8}$ . When  $H = 16$ , the sequence is  $\frac{1}{2^{0.5}}, \frac{1}{2^1}, \frac{1}{2^{1.5}}, \dots, \frac{1}{2^8}$ .

**Rotary.** Rotary [28] is the most adopted position encoding approach in the LLM community. It projects queries and keys into another space through rotations:

$$\langle \mathbf{q}_i, \mathbf{k}_j \rangle = (\mathbf{q}_i \mathbf{R}_{\Theta, -i}) (\mathbf{k}_j \mathbf{R}_{\Theta, -j})^T = \mathbf{q}_i \mathbf{R}_{\Theta, -i} \mathbf{R}_{\Theta, j} \mathbf{k}_j^T = \mathbf{q}_i \mathbf{R}_{\Theta, j-i} \mathbf{k}_j^T, \quad (9)$$

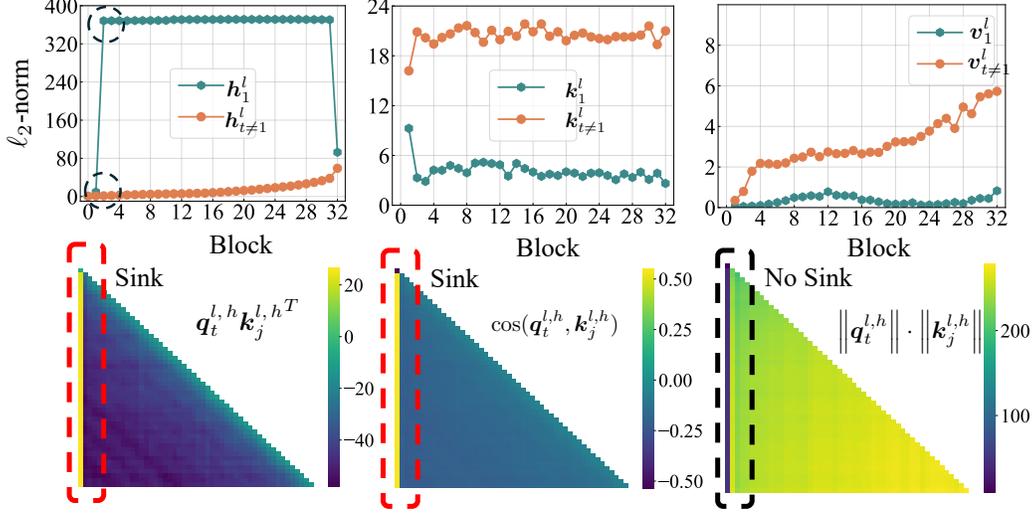


Figure 6: In LLaMA3-8B Base, (Top) the first token has significantly larger  $\ell_2$ -norm of hidden states, but much smaller  $\ell_2$ -norm of keys and values than the mean of other tokens; (Bottom) cosine similarity instead of norm product contributes to attention sink.

where  $\mathbf{R}_{\Theta,(\cdot)}$  is a pre-defined rotation matrix:

$$\mathbf{R}_{\Theta,m} = \begin{pmatrix} \cos m\omega_1 & -\sin m\omega_1 & 0 & 0 & \cdots & 0 & 0 \\ \sin m\omega_1 & \cos m\omega_1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cos m\omega_2 & -\sin m\omega_2 & \cdots & 0 & 0 \\ 0 & 0 & \sin m\omega_2 & \cos m\omega_2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & \cos m\omega_{d_h/2} & -\sin m\omega_{d_h/2} \\ 0 & 0 & 0 & 0 & \cdots & \sin m\omega_{d_h/2} & \cos m\omega_{d_h/2} \end{pmatrix}. \quad (10)$$

Here  $\omega_i = 1/10000^{2(i-1)/d_h}$  and  $d_h = d/H$  is the hidden dimension in each head. From the above definition, it is noted that the rotation matrix  $\mathbf{R}_{\Theta,m}^{d_h}$  satisfies  $\mathbf{R}_{\Theta,m}^T = \mathbf{R}_{\Theta,-m}$  and  $\mathbf{R}_{\Theta,i}\mathbf{R}_{\Theta,j} = \mathbf{R}_{\Theta,i+j}$ .

## C Properties of attention sink

### C.1 The first token acts as biases

**Uniqueness of the first token.** It is noted that the calculation of hidden states for the first token has no involvement of self-attention  $\mathbf{h}_1^l = \text{FFN}(\text{LN}(\mathbf{o}_1^l + \mathbf{h}_1^{l-1})) + \mathbf{o}_1^l + \mathbf{h}_1^{l-1}$ , where  $\mathbf{o}_1^l = \text{LN}(\mathbf{h}_1^{l-1}) [\mathbf{W}^{l,1} \ \mathbf{W}^{l,2} \ \cdots \ \mathbf{W}^{l,H}] \mathbf{W}_O^l$ . Therefore,  $\mathbf{h}_1^l$ , and corresponding queries/keys/values  $\mathbf{k}_1^{l,h} = \text{LN}(\mathbf{h}_1^{l-1}) \mathbf{W}_K^{l,h}$ ,  $\mathbf{q}_1^{l,h} = \text{LN}(\mathbf{h}_1^{l-1}) \mathbf{W}_Q^{l,h}$ ,  $\mathbf{v}_1^{l,h} = \text{LN}(\mathbf{h}_1^{l-1}) \mathbf{W}_V^{l,h}$  could be considered as the MLP output of input word embedding  $\mathbf{x}_1 \mathbf{W}_E$ . Cancedda [4], Sun et al. [29] showed that  $\mathbf{h}_1^l$  has large norm or massive activations, thus behaving as a bias to absorb the attention. Using LLaMA3-8B Base model [9], we show that from certain transformer block, e.g.,  $l = 2$ , the  $\ell_2$ -norm of  $\mathbf{h}_1^l$  is significantly larger than that of other tokens  $\mathbf{h}_{t \neq 1}$  in Figure 6(Top). Despite the large  $\ell_2$ -norm of hidden states, we observe that the  $\ell_2$ -norm of keys and values of the first token is significantly smaller than that of other tokens in the same figure. This motivates us to further explore the functionality of massive activations in hidden states.

**QK angles contribute to attention sink.** In the  $l$ -th transformer block, we consider the keys and queries after adding PE (Rotary in LLaMA3-8B Base):  $\mathbf{k}_t^{l,h} = \text{LN}(\mathbf{h}_t^{l-1}) \mathbf{W}_K^{l,h} \mathbf{R}_{\Theta,-t}$ ,  $\mathbf{q}_t^{l,h} = \text{LN}(\mathbf{h}_t^{l-1}) \mathbf{W}_Q^{l,h} \mathbf{R}_{\Theta,-t}$ , where LN is RMSNorm [38]:  $\text{LN}(\mathbf{h}) = \frac{\mathbf{h}}{\text{RMS}(\mathbf{h})} \odot \mathbf{g}$  and  $\text{RMS}(\mathbf{h}) =$

$\sqrt{\frac{1}{d} \sum_{i=1}^d h_i^2}$ . Here  $g$  is a learnable gain parameter. Suppose that  $h_1^{l-1}$  already has massive activations. Since  $h_1^l$  has a massive magnitude in specific dimensions, the LN operation retains the magnitude in these dimensions and further reduces the magnitude in other dimensions, leading to that  $q_1^{l,h}$ ,  $k_1^{l,h}$ , and  $v_1^{l,h}$  are distributed on different manifolds, especially for  $k_1^{l,h}$ .

For the  $t$ -th query, we know that  $q_t^{l,h} k_1^{l,h^T}$  typically has much larger values than  $q_t^{l,h} k_{j \neq 1}^{l,h^T}$ , as visualized in Figure 6(Bottom). We further show that due to the different manifold of  $k_1^{l,h}$ , the angles between  $k_1^{l,h}$  and  $q_t^{l,h}$  play an important role. Considering  $q_t^{l,h} k_j^{l,h^T} = \|q_t^{l,h}\| \cdot \|k_j^{l,h}\| \cdot \cos(q_t^{l,h}, k_j^{l,h})$ , we visualize the cosine similarity between keys and values, and the product of  $\ell_2$ -norm between keys and values in Figure 6(Bottom). Although  $\|q_t^{l,h}\| \cdot \|k_1^{l,h}\|$  is comparatively small,  $\cos(q_t^{l,h}, k_1^{l,h})$  is significantly large, leading to attention sink. This explains why attention sink exists despite the small  $\ell_2$ -norm of keys of the first token. To conclude, the first token leverages its keys to act as biases, thus minimizing the angles between  $k_1^{l,h}$  and  $q_t^{l,h}$  and exhibiting attention sink.

## C.2 Attention sink under different inputs

**Different data domains.** We first explore the effects of input domains on attention sinks. The pile dataset [10], a regular dataset for LM pretraining, has 17 available data domains. We sample 100 data from each domain and then evaluate the attention sink metric for GPT2-XL/Mistral-7B/LLaMA2-7B Base/LLaMA3-8B base. As shown in Table 7, the evaluated attention sink metrics are similar across different domains.

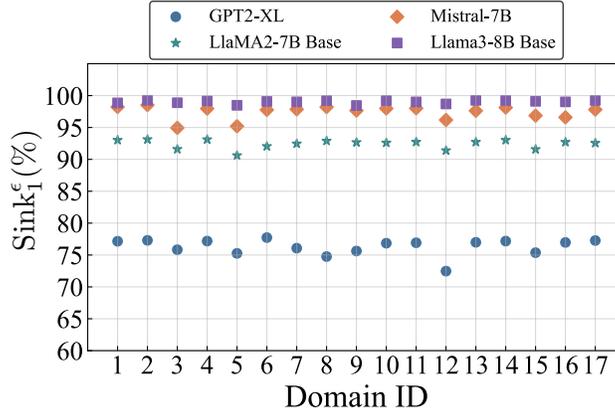


Figure 7: Input domains have negligible effects on attention sink.

**Beyond natural languages.** We also consider two ideal scenarios: (i) randomly sample  $T$  tokens from the tokenizer vocabulary  $\mathbb{V}$  to construct a sequence and (ii) randomly sample 1 token from the tokenizer  $\mathbb{V}$  and repeat it  $T$  times. We exclude the BOS token. As present in Table 5(Left), compared to natural language, random tokens have insignificant effects on the existence of attention sink. However, with repeated tokens, attention sink in Mistral [16] and LLaMA models disappears. In Appendix C.4, we prove that for LMs with NoPE/relative PE/ALiBi/Rotary, if the first  $T$  tokens are the same, their corresponding hidden states are the same. They all have massive activations, thus dispersing the attention sink. We also provide the closed form/upper bound for attention scores in these LMs through Propositions 1-4.

## C.3 Attention sink under different LMs

**Base vs. chat model.** Compared with base models, chat models are typically continually trained through instruction tuning [22]. From Table 5(Right), instruction tuning has an insignificant impact on attention sink, which motivates us to focus on the LM pre-training.

Table 5: (Left) Even with random sequence as input, there still exists an obvious attention sink. But with repeated tokens, the attention sink disappears for Mistral/LLaMA models. (Right) Chat models have comparable attention sink metrics with base models.

LLM	Sink <sub>1</sub> <sup>ε</sup> (%)			LLM	Sink <sub>1</sub> <sup>ε</sup> (%)	
	natural	random	repeat		Base	Chat
GPT2-XL	77.00	70.29	62.28	Mistral-7B	97.49	88.34
Mistral-7B	97.49	75.21	0.00	LLaMA2-7B	92.47	92.88
LLaMA2-7B Base	92.47	90.13	0.00	LLaMA2-13B	91.69	90.94
LLaMA3-8B Base	99.02	91.23	0.00	LLaMA3-8B	99.02	98.85

**Model scale.** We evaluate the metric Sink<sub>1</sub><sup>ε</sup> of LLaMA2 Base [30], LLaMA3 Base [9], Pythia [2], GPT2 [24], OPT [39] families. As visualized in Figure 1(Left), attention sink emerges in small LMs, even in Pythia-14M. Only in Pythia family, larger-sized LMs tend to have more obvious attention sink.

#### C.4 How positional embedding links with attention sink

In Section C.2, we have shown that Mistral-7B, LLaMA2-7B Base, and LLaMA3-8B Base, which adopt rotary as PE, have no attention sink for repeated token sequences. While GPT2, which adopts learnable PE, has the attention sink in such a scenario. To further understand this, we explore how positional embedding plays a role through a theoretical perspective.

**Activations.** Suppose the repeated sequence is  $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$  and each  $\mathbf{x}_t = \mathbf{x}$ . For LMs with NoPE/relative PE/ALiBi/Rotary, we have the initial hidden states  $\mathbf{h}_t^0 = \mathbf{x}\mathbf{W}_E$ , which are the same among all the token positions since  $\mathbf{P} = \mathbf{0}$ . Then for the first transformer block  $l = 1$ , we know that  $\mathbf{k}_t^{1,h} = \text{LN}(\mathbf{h}_t^0)\mathbf{W}_K^{1,h} = \text{LN}(\mathbf{x}\mathbf{W}_E)\mathbf{W}_K^{1,h}$ ,  $\mathbf{q}_t^{1,h} = \text{LN}(\mathbf{h}_t^0)\mathbf{W}_Q^{1,h} = \text{LN}(\mathbf{x}\mathbf{W}_E)\mathbf{W}_Q^{1,h}$ ,  $\mathbf{v}_t^{1,h} = \text{LN}(\mathbf{h}_t^0)\mathbf{W}_V^{1,h} = \text{LN}(\mathbf{x}\mathbf{W}_E)\mathbf{W}_V^{1,h}$ . Then all tokens have the same  $\mathbf{k}_t^{1,h}$ ,  $\mathbf{q}_t^{1,h}$ , and  $\mathbf{v}_t^{1,h}$ . Then the attention output is

$$\mathbf{v}_t^{\prime 1,h} = \sum_{i=1}^t \mathbf{A}_{ii}^{1,h} \mathbf{v}_i^{1,h} = \mathbf{v}_t^{1,h} \quad (11)$$

Then  $\mathbf{o}_t^1 = \text{Concat}_{h=1}^H(\mathbf{v}_t^{\prime 1,h})\mathbf{W}_O^1$ , we have the hidden states after the first block:

$$\mathbf{h}_t^1 = \text{FFN}(\text{LN}(\mathbf{o}_t^1 + \mathbf{h}_t^0)) + \mathbf{o}_t^1 + \mathbf{h}_t^0 \quad (12)$$

$$= \text{FFN}(\text{LN}(\text{Concat}_{h=1}^H(\mathbf{v}_t^{\prime 1,h})\mathbf{W}_O^1 + \mathbf{h}_t^0)) + \text{Concat}_{h=1}^H(\mathbf{v}_t^{\prime 1,h})\mathbf{W}_O^1 + \mathbf{h}_t^0 \quad (13)$$

$$= \text{FFN}(\text{LN}(\text{Concat}_{h=1}^H(\text{LN}(\mathbf{h}_t^0)\mathbf{W}_V^{1,h})\mathbf{W}_O^1 + \mathbf{h}_t^0)) \quad (14)$$

$$+ \text{Concat}_{h=1}^H(\text{LN}(\mathbf{h}_t^0)\mathbf{W}_V^{1,h})\mathbf{W}_O^1 + \mathbf{h}_t^0. \quad (15)$$

Since  $\mathbf{h}_1^0 = \mathbf{h}_2^0 = \dots = \mathbf{h}_T^0$ , we have  $\mathbf{h}_1^1 = \mathbf{h}_2^1 = \dots = \mathbf{h}_T^1$  based on the above equation. Using this induction, we could prove that

$$\mathbf{h}_t^l = \text{FFN}(\text{LN}(\text{Concat}_{h=1}^H(\text{LN}(\mathbf{h}_t^{l-1})\mathbf{W}_V^{l,h})\mathbf{W}_O^l + \mathbf{h}_t^{l-1})) \quad (16)$$

$$+ \text{Concat}_{h=1}^H(\text{LN}(\mathbf{h}_t^{l-1})\mathbf{W}_V^{l,h})\mathbf{W}_O^l + \mathbf{h}_t^{l-1}, \quad \forall 1 \leq l \leq L. \quad (17)$$

$$\mathbf{h}_1^l = \mathbf{h}_2^l = \dots = \mathbf{h}_T^l, \quad \forall 0 \leq l \leq L. \quad (18)$$

Typically, the hidden states of the first token  $\mathbf{h}_1^l$  in specific blocks have massive activations. Due to the above equality, all the repeated tokens have massive activations.

For LMs with absolute PE/learnable PE, the initial hidden states  $\mathbf{h}_t^0 = \mathbf{x}\mathbf{W}_E + \mathbf{p}_t$ . Although the word embeddings are the same for repeated tokens,  $\mathbf{p}_t$  for different token positions  $t$  is different. Therefore, GPT2 models have no the above equality.

**Proposition 1.** For LMs with NoPE, the attention scores for  $t$  repeated tokens are  $t^{-1}$  uniformly, i.e., there is no attention sink.

*Proof.* We have that

$$\mathbf{A}_{t,i}^{l,h} = \frac{e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle}}{\sum_{j=1}^t e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_j^{l,h} \rangle}} = \frac{e^{\mathbf{q}_t^{l,h} \mathbf{k}_i^{l,h T}}}{\sum_{j=1}^t e^{\mathbf{q}_t^{l,h} \mathbf{k}_j^{l,h T}}} = \frac{e^{\mathbf{q}_t^{l,h} \mathbf{k}_i^{l,h T}}}{t e^{\mathbf{q}_t^{l,h} \mathbf{k}_i^{l,h T}}} = \frac{1}{t}. \quad (19)$$

Therefore, the attention scores follow a uniform distribution over all previous tokens.  $\square$

**Proposition 2.** For LMs with relative PE, there is no attention sink for  $t$  repeated tokens.

*Proof.* For LMs with relative PE, the dot product between each query and key is

$$\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle = \mathbf{q}_t^{l,h} \mathbf{k}_i^{l,h T} + g_{\text{rel}}(t-i) = \mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{rel}}(t-i), \quad (20)$$

then we have the attention scores

$$\mathbf{A}_{t,i}^{l,h} = \frac{e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle}}{\sum_{j=1}^t e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_j^{l,h} \rangle}} = \frac{e^{\mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{rel}}(t-i)}}{\sum_{j=1}^t e^{\mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{rel}}(t-j)}} = \frac{e^{g_{\text{rel}}(t-i)}}{\sum_{j=1}^t e^{g_{\text{rel}}(t-j)}}. \quad (21)$$

Considering  $g_{\text{rel}}(t-i)$  is a monotonic non-increasing function of  $t-i$  and  $g_{\text{rel}}(t-i) = \mathcal{B} - 1$  when  $t-i > \mathcal{D}$ , then  $\mathbf{A}_{t,1}^{l,h} = \mathbf{A}_{t,1}^{l,h} = \dots = \mathbf{A}_{t,t-\mathcal{D}}^{l,h}$  are the largest values. Therefore, there is no attention sink on the first token.  $\square$

**Proposition 3.** For LMs with ALiBi, there is no attention sink for  $t$  repeated tokens.

*Proof.* For LMs with ALiBi, similar to relative PE, the dot product between each query and key is

$$\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle = \mathbf{q}_t^{l,h} \mathbf{k}_i^{l,h T} + g_{\text{alibi}}^h(t-i) = \mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{alibi}}^h(t-i), \quad (22)$$

then we have the attention scores

$$\mathbf{A}_{t,i}^{l,h} = \frac{e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle}}{\sum_{j=1}^t e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_j^{l,h} \rangle}} = \frac{e^{\mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{alibi}}^h(t-i)}}{\sum_{j=1}^t e^{\mathbf{q}^{l,h} \mathbf{k}^{l,h T} + g_{\text{alibi}}^h(t-j)}} = \frac{e^{g_{\text{alibi}}^h(t-i)}}{\sum_{j=1}^t e^{g_{\text{alibi}}^h(t-j)}}. \quad (23)$$

Here  $g_{\text{alibi}}^h(t-i)$  is monotonic decreasing function of  $t-i$ , so there is no attention sink on the first token.  $\square$

**Proposition 4.** For LMs with Rotary, there is no attention sink for  $t$  repeated tokens when  $t$  is large if  $\|\mathbf{q}^{l,h}\| \cdot \|\mathbf{k}^{l,h}\| \leq \xi$  for a constant  $\xi$ .

*Proof.* For LMs with Rotary, the dot product between each query and key is

$$\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle = \mathbf{q}_t^{l,h} \mathbf{R}_{\Theta, i-t} \mathbf{k}_i^{l,h T} \quad (24)$$

$$= \mathbf{q}^{l,h} \mathbf{R}_{\Theta, i-t} \mathbf{k}^{l,h T} \quad (25)$$

$$= \|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h} \mathbf{R}_{\Theta, i-t}\| \cos\left(\frac{\mathbf{q}^{l,h} \mathbf{R}_{\Theta, i-t} \mathbf{k}^{l,h T}}{\|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h} \mathbf{R}_{\Theta, i-t}\|}\right) \quad (26)$$

$$= \|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h}\| \cos(\beta_{t-i}), \quad (27)$$

where  $\beta_{j-t}$  is the angle between the rotated query and the rotated key. Then the attention scores are

$$\mathbf{A}_{t,i}^{l,h} = \frac{e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_i^{l,h} \rangle}}{\sum_{j=1}^t e^{\langle \mathbf{q}_t^{l,h}, \mathbf{k}_j^{l,h} \rangle}} = \frac{e^{\mathbf{q}^{l,h} \mathbf{R}_{\Theta, j-t} \mathbf{k}^{l,h T}}}{\sum_{j=1}^t e^{\mathbf{q}^{l,h} \mathbf{R}_{\Theta, j-t} \mathbf{k}^{l,h T}}} = \frac{e^{\|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h}\| \cos(\beta_{t-i})}}{\sum_{j=1}^t e^{\|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h}\| \cos(\beta_{t-j})}}. \quad (28)$$

Suppose the norm of multiplication for query and key  $\|\mathbf{q}^{l,h}\| \|\mathbf{k}^{l,h}\| = \xi$ . Considering  $-1 \leq \cos(\beta_{t-j}) \leq 1$ , then we have

$$\mathbf{A}_{t,i}^{l,h} = \frac{e^{\xi \cos(\beta_{t-i})}}{\sum_{j=1}^t e^{\xi \cos(\beta_{t-j})}} = \frac{1}{1 + \frac{\sum_{j \neq i} e^{\xi \cos(\beta_{t-j})}}{e^{\xi \cos(\beta_{t-i})}}} \leq \frac{e^{2\xi}}{e^{2\xi} + (t-1)} \quad (29)$$

Then the attention scores for each token are upper-bounded and decrease to 0 as  $t$  grows.  $\square$

## D More experiments and results

### D.1 Optimization and data distribution

**Batch size.** During the pre-training, we consider different batch sizes with other hyper-parameters fixed. As shown in Table 6, batch size does not affect the emergence of attention sink.

Table 6: Batch size has no effect on the emergence of attention sink.

batch size	0.25M	0.5M	1M	2M
Sink <sub>1</sub> <sup>ε</sup> (%)	16.45	17.19	18.18	16.19
valid loss	3.92	3.78	3.73	3.68

**Fixing token in a specific position.** During the pre-training, we consider fixing the token  $x_{\text{fix}}$  in the first/second/third position. Consequently, when evaluating the attention sink metric, in Table 7, we find that attention sink appears in the appears in the fixed token instead of the first token.

Table 7: Attention sink appears in the appears in the fixed token instead of the first token.

Token position	$x_1 = x_{\text{fix}}$	$x_2 = x_{\text{fix}}$	$x_3 = x_{\text{fix}}$
Sink <sub>1</sub> <sup>ε</sup> (%)	74.11	0.00	0.00
Sink <sub>2</sub> <sup>ε</sup> (%)	0.00	69.03	0.00
Sink <sub>3</sub> <sup>ε</sup> (%)	0.00	0.00	69.64
Sink <sub>4</sub> <sup>ε</sup> (%)	0.01	0.01	0.00

### D.2 Activation functions in FFN

Since FFN in the earlier transformer block blasts off the  $\ell_2$ -norm of the first token, we are wondering whether the choices of activation functions in FFN will affect the attention sink. Besides the SwiGLU activations used in our default setup, we also consider other activation functions, as present in Table 8. We observe that different FFN designs do not affect the emergence of attention sink.

Table 8: Modifying the activation functions in FFN does not affect the emergence of attention sink.

Activation functions	$F$	Sink <sub>1</sub> <sup>ε</sup> (%)	valid loss
ReLU	$\text{ReLU}(\mathbf{h}\mathbf{W}_1)\mathbf{W}_2$	16.90	3.82
GeLU [14]	$\text{GeLU}(\mathbf{h}\mathbf{W}_1)\mathbf{W}_2$	14.76	3.79
Swish [26]	$\text{Swish}(\mathbf{h}\mathbf{W}_1)\mathbf{W}_2$	17.89	3.80
ReGLU [27]	$(\text{ReLU}(\mathbf{h}\mathbf{W}_1) \odot \mathbf{h}\mathbf{W}_2)\mathbf{W}_3$	13.88	3.75
GeGLU [27]	$(\text{GeLU}(\mathbf{h}\mathbf{W}_1) \odot \mathbf{h}\mathbf{W}_2)\mathbf{W}_3$	17.86	3.73
SwiGLU [27]	$(\text{Swish}(\mathbf{h}\mathbf{W}_1) \odot \mathbf{h}\mathbf{W}_2)\mathbf{W}_3$	18.18	3.73

### D.3 Pre-norm and post-norm structure

Layer normalization (LN) [1, 38] regularizes the hidden states in LMs by re-centering and re-scaling, which may affect the massive activations. This motivates us to explore the effects of LN location on attention sink. In the pre-norm structure, as stated in Equation 3, hidden states from the earlier blocks could be retained by the later blocks through residual connections [13]. Therefore, if massive activations appear in a specific block, they will likely be retained in the subsequent blocks. Within a post-norm structure, the hidden states will be normalized before being fed into the following blocks, as present in Figure 5(Left).

When replacing the pre-norm structure with the post-norm structure, Sink<sub>1</sub><sup>ε</sup> becomes 13.54%. This indicates that the attention sink still exists in post-norm LMs. After further investigations, as visualized in Figure 8(Left), massive activations exist in the hidden states before the post LN instead of  $h_1^1$ .

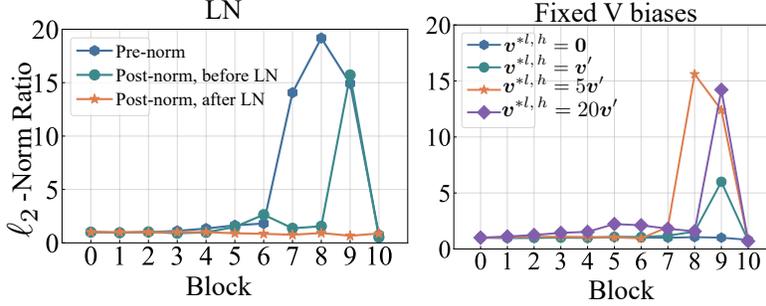


Figure 8:  $\ell_2$ -norm ratio of  $\mathbf{h}_1^l$  and mean of  $\mathbf{h}_{t \neq 0}^l$ . (Left) Massive activations exist not in hidden states in post-norm LMs, but in the features before LN. (Right) Massive activations emerge when increasing the  $\ell_2$ -norm of fixed  $\mathbf{v}^{*,h}$  for the setup of K biases.

#### D.4 Multi-head design in Attention

We consider two perspectives in multi-head design in attention. Firstly, we explore whether the number of heads has an impact on attention sink, especially for  $H = 1$ , which refers to single-head self-attention. Second, attention output from each head is concatenated:  $\mathbf{O}^l = \text{Concat}_{h=1}^H (\mathbf{A}^{l,h} \mathbf{V}^{l,h}) \mathbf{W}_O^l$ . We replace such concatenation operation with addition operation:  $\mathbf{O}^l = \sum_{h=1}^H (\mathbf{A}^{l,h} \mathbf{V}^{l,h}) \mathbf{W}_O^l$ . As shown in Table 9, multi-head design does not affect the emergence of attention sink.

Table 9: Modifying multi-head design does not affect the emergence of attention sink.

multi-head design	$H = 8$	$H = 4$	$H = 2$	$H = 1$	multi-head addition
Sink <sub>1</sub> <sup>ξ</sup> (%)	18.18	10.68	12.95	19.50	21.76
Valid loss	3.73	3.74	3.76	3.78	3.74

#### D.5 Attention biases

**Beyond all zeros in V biases.** In the setup of K biases, we fix  $\mathbf{v}^{*,h} = \mathbf{0}$ . We wonder whether the fixed values of  $\mathbf{v}^{*,h}$  could affect the attention sink. We consider,  $\mathbf{v}^{*,h} = m\mathbf{v}'$  or  $\mathbf{v}^{*,h} = m\mathbf{v}''$ , where  $m$  is the controllable  $\ell_2$  norm and  $\mathbf{v}' = [1, 0, 0, \dots, 0]$  and  $\mathbf{v}'' = [1, 1, 1, \dots, 1]/\sqrt{d_h}$ . As shown in Table 10, with larger  $\ell_2$ -norm of  $\mathbf{v}^{*,h}$ , attention sink shifts from  $\mathbf{k}^{*,h}$  to the first token. Meanwhile, massive activations on the first token appear, as visualized in Figure 8(Right). Intuitively, it is difficult for LMs to remove the effects of  $\mathbf{v}^{*,h}$  with larger  $\ell_2$ -norm in model predictions. Then they opt to optimize the keys and values of the first token to save extra attention.

Table 10: Larger  $\ell_2$ -norm of fixed  $\mathbf{v}^{*,h}$  results in LMs allocating more attention on  $\mathbf{x}_1$  instead of  $\mathbf{k}^{*,h}$ .

$\mathbf{v}^{*,h}$	$\mathbf{0}$	$\mathbf{v}'$	$5\mathbf{v}'$	$20\mathbf{v}'$	$\mathbf{v}''$	$5\mathbf{v}''$	$20\mathbf{v}''$
Sink <sub>1</sub> <sup>ξ</sup> (%)	73.34	70.03	44.43	1.51	69.74	27.99	0.00
Sink <sub>1</sub> <sup>ξ</sup> (%)	0.00	0.06	3.71	25.88	2.15	5.93	11.21
valid loss	3.72	3.72	3.72	3.71	3.72	3.72	3.73

**Head-sharing K/KV Biases in Attention.** In the main paper, we consider both KV biases and V biases in attention. These biases are not shared by different heads in each block. We further explore whether head-sharing patterns will affect their functionality. As present in Table 11, LMs with KV biases are more likely affected by the head-sharing pattern: attention sink shifts from the K biases to the first token. While LMs with K biases are less affected.

Table 11: Sharing KV biases across heads in each block results in attention sink shifting back to the first token from K biases.

Head-sharing Biases type	✓ KV	✓ KV	× K	× K
Sink <sub>*</sub> <sup>ε</sup> (%)	72.76	56.61	73.34	68.31
Sink <sub>1</sub> <sup>ε</sup> (%)	0.04	12.44	0.00	0.23
Valid loss	3.72	3.72	3.72	3.72

## D.6 Scaling the normalization in Softmax attention.

Before our experiments about replacing softmax attention, we first explore the effects of normalization scales in softmax attention. In the main paper, the attention output for the  $i$ -th output is

$$\mathbf{v}_i^\dagger = \sum_{j=1}^i \frac{\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))}{\sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))} \mathbf{v}_j = \sum_{j=1}^i \frac{\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))}{\mathbf{Z}_i} \mathbf{v}_j. \quad (30)$$

We consider a scale factor  $\alpha$ , the normalization term is  $\mathbf{Z}_i = \frac{1}{\alpha} \sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))$ , and then the attention score are sum up to  $\alpha$ . For default setup,  $\alpha = 1$ . As shown in Table 12, with a smaller normalization scale, attention sink tends to appear in fewer heads. From another perspective,

$$\mathbf{v}_i^\dagger = \sum_{j=1}^i \frac{\alpha \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))}{\sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))} \mathbf{v}_j = \sum_{j=1}^i \frac{\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))}{\sum_{j'=1}^i \text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_{j'}))} \mathbf{h}_j(\alpha \mathbf{W}_V), \quad (31)$$

$$\mathbf{o}'_i = \text{Concat}_{h=1}^H(\mathbf{v}_i^h) \mathbf{W}_O. \quad (32)$$

Therefore, this normalization scale can be regarded as the scale for  $\mathbf{W}_V$  or  $\mathbf{W}_O$ . We show that this normalization scaling could be implemented by scaling learning rates and initialization. We use the  $s$  to represent the optimization step, and  $s = 0$  refers to the initialization. When scaling the normalization, we have the following update rule (take  $\mathbf{W}_O$  for example):

$$\mathbf{W}_O^{s+1} = \mathbf{W}_O^s - \eta \nabla \mathcal{L}(\alpha \mathbf{W}_O^s) \quad (33)$$

$$= \mathbf{W}_O^s - \alpha \eta \nabla \mathcal{L}(\mathbf{W})|_{\mathbf{W}=\alpha \mathbf{W}_O^s}, \quad (34)$$

where  $\eta$  is the original learning rate. Suppose we only modify the learning rate and initialization, to ensure each optimization step  $\hat{\mathbf{W}}_O^s = \mathbf{W}_O^s$ , we need to first ensure  $\hat{\mathbf{W}}_O^0 = \alpha \mathbf{W}_O^0$ . Suppose that we have  $\hat{\mathbf{W}}_O^s = \mathbf{W}_O^s$ , then the update rule is:

$$\hat{\mathbf{W}}_O^{s+1} = \hat{\mathbf{W}}_O^s - \eta' \nabla \mathcal{L}(\hat{\mathbf{W}}_O^s) \quad (35)$$

$$= \alpha \mathbf{W}_O^s - \eta' \nabla \mathcal{L}(\mathbf{W})|_{\mathbf{W}=\alpha \mathbf{W}_O^s}, \quad (36)$$

To ensure  $\hat{\mathbf{W}}_O^{s+1} = \mathbf{W}_O^{s+1}$ , we need the new learning rate  $\eta'$  meets  $\eta' = \alpha^2 \eta$ .

Table 12: Scaling the normalization in Softmax attention to less than one can mitigate attention sink but not prevent its emergence.

$\alpha$	2.0	1.0	0.2	0.1	0.02
Sink <sub>1</sub> <sup>ε</sup> (%)	29.10	18.18	9.41	3.59	4.53
Valid loss	3.72	3.73	3.72	3.76	3.78

## D.7 More results about attention operations

Firstly, we present all attempted attention operations in Table 13. It is noted that several setups lead to training failure. For LMs with sigmoid attention without normalization, we vary the learning rates or weight decay ratios  $\gamma$ . Consequently, as shown in Table 14, the trained LMs still have no attention sink, which further confirms our conclusion in the main paper.

Table 13: Normalization and selections of kernels in attention significantly affect the emergence of the attention sink. We use “\*” to mark that the metric Sink<sub>1</sub><sup>ε</sup> is computed by proxy attention scores. We use “-” to represent the training failure under the setup.

$\text{sim}(\varphi(\mathbf{q}_i), \varphi(\mathbf{k}_j))$	$\mathbf{Z}_i$	Sink <sub>1</sub> <sup>ε</sup> (%)	Valid loss
$\exp\left(\frac{\mathbf{q}_i^T \mathbf{k}_j}{\sqrt{d_h}}\right)$	$\sum_{j'=1}^i \exp\left(\frac{\mathbf{q}_i \mathbf{k}_{j'}^T}{\sqrt{d_h}}\right)$	18.18	3.73
$\text{sigmoid}\left(\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}\right)$	1	0.44*	3.70
$\text{sigmoid}\left(\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}\right)$	$\sum_{j'=1}^i \text{sigmoid}\left(\frac{\mathbf{q}_i \mathbf{k}_{j'}^T}{\sqrt{d_h}}\right)$	30.24	3.74
$\text{elu}\left(\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}\right) + 1$	1	0.80*	3.69
$\text{elu}\left(\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}\right) + 1$	$\sum_{j'=1}^i \text{elu}\left(\frac{\mathbf{q}_i \mathbf{k}_{j'}^T}{\sqrt{d_h}}\right) + 1$	-	-
$\frac{(\text{elu}(\mathbf{q}_i)+1)(\text{elu}(\mathbf{k}_j)+1)^T}{\sqrt{d_h}}$	$\sum_{j'=1}^i \frac{(\text{elu}(\mathbf{q}_i)+1)(\text{elu}(\mathbf{k}_{j'}+1)^T}{\sqrt{d_h}}$	53.65*	4.19
$\frac{(\text{elu}(\mathbf{q}_i)+1)(\text{elu}(\mathbf{k}_j)+1)^T}{\sqrt{d_h}}$	1	-	-
$\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}$	$\max\left(\left \sum_{j'=1}^i \frac{\mathbf{q}_i \mathbf{k}_{j'}^T}{\sqrt{d_h}}\right , 1\right)$	-	-
$\frac{\mathbf{q}_i \mathbf{k}_j^T}{\sqrt{d_h}}$	1	0.00*	3.99
$\frac{\text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_j)^T}{\sqrt{d_h}}$	$\max\left(\left \sum_{j'=1}^i \frac{\text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_{j'})^T}{\sqrt{d_h}}\right , 1\right)$	0.19*	3.85
$\frac{\text{mlp}(\mathbf{q}_i) \text{mlp}(\mathbf{k}_j)^T}{\sqrt{d_h}}$	1	0.74*	3.91

Table 14: LMs with sigmoid attention (without sigmoid attention) trained by different learning rates and weight decay ratios have no attention sink.

Learning rate	4e-4	4e-4	4e-4	4e-4	1e-3	1e-4
Weight decay	0.0	0.1	0.5	1.0	0.1	0.1
Sink <sub>1</sub> <sup>ε</sup> (%)	0.64	0.44	0.18	0.30	0.81	0.36
Valid loss	3.77	3.70	3.76	4.06	3.68	4.08