

Dissecting Transformer Length Extrapolation via The Lens of Receptive Field Analysis

Ta-Chung Chi

Carnegie Mellon University
tachungc@andrew.cmu.edu

Alexander I. Rudnicky
Carnegie Mellon University
air@cs.cmu.edu

Ting-Han Fan

Princeton University
tinghanf@princeton.edu

Peter J. Ramadge
Princeton University
ramadge@princeton.edu

Abstract

Length extrapolation permits training a transformer language model on short sequences that preserves perplexities when tested on substantially longer sequences. A relative positional embedding design, ALiBi, has had the widest usage to date. We dissect ALiBi via the lens of receptive field analysis empowered by a novel cumulative normalized gradient tool. The concept of receptive field further allows us to modify the vanilla Sinusoidal positional embedding to create **Sandwich**, the first parameter-free relative positional embedding design that truly length information uses longer than the training sequence. Sandwich shares with KERPLE and T5 the same logarithmic decaying temporal bias pattern with learnable relative positional embeddings; these elucidate future extrapolatable positional embedding design.

1 Introduction

The length of input sequences is an important hyperparameter choice for pretraining a transformer language model. A vanilla transformer language model has a quadratic training cost w.r.t L_{tr} , the training sequence length. As the value of L_{tr} increases, cost becomes impractical. However, we can use the model for substantially longer evaluation sequence lengths $L_{ex} \gg L_{tr}$ as gradients no longer need to be recorded. The discrepancy between L_{tr} and L_{ex} motivates the task of **length extrapolation** (Press et al., 2022): *Can a transformer language model maintain equally good, if not better, perplexities when longer sequences are used in the testing stage?*

Several extrapolatable transformer language models have been proposed including ALiBi (Press et al., 2022) and KERPLE (Chi et al., 2022), of which the relative positional embedding design is hypothesized to be critical to success. Empirically, they extrapolate to $L_{ex} \gg L_{tr}$ much better than other absolute and relative positional embeddings

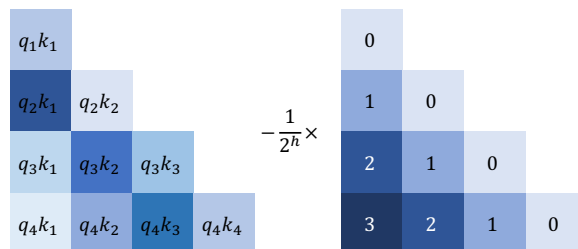


Figure 1: ALiBi. For a transformer language model with H attention heads, the range of h is $n \cdot \frac{8}{H}$, where $n = \{1 \dots H\}$. Left = self-attention matrix, right = temporal biases matrix.

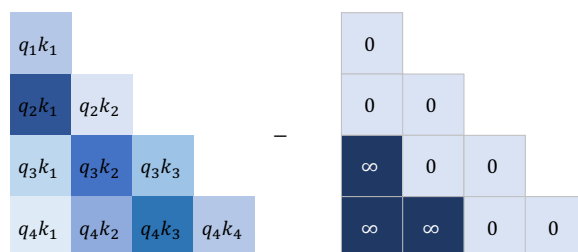


Figure 2: Windowed Attention. This is the same design as Longformer (Beltagy et al., 2020). We limit the context window size to $w = 2$ in this example. Left = self-attention matrix, right = temporal biases matrix.

including Sinusoidal (Vaswani et al., 2017), Rotary (Su et al., 2021), and T5 (Raffel et al., 2020), resulting in the adoption of ALiBi for the recently released Bloom (Scao et al., 2022) model. Despite the significant empirical success of ALiBi, there is still a lack of fundamental understanding of why it works.¹

Figure 1 shows the implementation of ALiBi. We hereinafter refer to the coefficient $\frac{1}{2^h}$ as *slope*. Intuitively, ALiBi encourages a token to focus on neighbors based on its temporal biases matrix. When two tokens are distant, ALiBi becomes highly similar to windowed attention, shown in Figure 2. Experiments in §4 will further establish the

¹https://github.com/ofirpress/attention_with_linear_biases#why-do-you-think-alibi-works

connection between the two.

Windowed attention allows the easy derivation of a theoretical (maximum) receptive field: wR for an R layer transformer model with windowed attention size w . A windowed attention model can extrapolate if $L_{tr} > wR$ because 1) wR is fully covered by L_{tr} during the training stage, and 2) it simply ignores the additional $L_{ex} - wR$ tokens during the testing stage. Surprisingly, a model can still extrapolate when $L_{tr} < wR$ which we show in §4. This calls for the need for empirical receptive field measurement and motivates our model-agnostic cumulative normalized gradient tool. The tool we develop can be applied back on ALiBi to show that L_{tr} covers most of its empirical receptive field.

Our analysis tool also provides critical context for explaining the length extrapolation failure (Press et al., 2022; Chi et al., 2022) of Sinusoidal (Vaswani et al., 2017) and Rotary (Su et al., 2021) by showing their violation of the empirical receptive field coverage principle. Sinusoidal can be fixed by dropping the intermediate terms and keeping only the decay-with-distance biases; this leads to the creation of **Sandwich**, the first parameter-free relative positional embedding that uses information beyond L_{tr} . Sandwich shares a similar temporal bias pattern with trainable positional embeddings such as KERPLE (Chi et al., 2022) and T5 (Raffel et al., 2020), and they jointly suggest the future design of extrapolatable transformer positional embeddings.

2 Related Work

2.1 Length Extrapolation

In the context of language modeling, we expect token-level perplexities to remain at least the same, if not lower (i.e. better), when $L_{ex} \gg L_{tr}$ sequences are provided. Recurrent neural networks (Mikolov et al., 2010; Mikolov and Zweig, 2012; Zaremba et al., 2014) can easily perform length extrapolation. But this is not an easy task for transformer language models, among which only those equipped with special relative positional embeddings (Press et al., 2022; Chi et al., 2022) are length extrapolatable.

2.2 Positional Embeddings

It is widely believed that the design of positional embeddings is the key to successful length extrapolation of transformer language models (Press et al., 2022; Chi et al., 2022). We can roughly catego-

rize existing positional embeddings into absolute (APE) (Vaswani et al., 2017) and relative (RPE) (Su et al., 2021; Raffel et al., 2020; Press et al., 2022; Chi et al., 2022) variants. APE often assigns one positional embedding per token and combines them directly with input embeddings. In contrast, RPE adds temporal bias terms to the self-attention matrix to encode the relative distance between token pairs. For example, the right triangular matrix in Figure 1 shows the set of temporal bias terms. It is challenging for APE to extrapolate well without any further fine-tuning since either the beyond L positional embeddings do not exist, or the model needs to process unseen positional embeddings (e.g. unseen sinusoidal embeddings). (Press et al., 2022; Chi et al., 2022). In contrast, RPE usually performs better length extrapolation since it is easier to construct the additional temporal bias terms.

2.3 Windowed and Sparse Attention

We will see later that ALiBi can be viewed as imposing a windowed attention mask on the self-attention matrix, similar to previous transformer models with sparse attention (Beltagy et al., 2020; Zaheer et al., 2020; Ainslie et al., 2020; Gupta and Berant, 2020). Interpreting ALiBi from the perspective of windowed attention allows us to easily calculate the theoretical receptive field of a model.

2.4 Receptive Field

A model’s receptive field is defined as the size of the input region that contributes the most to model outputs. It is often measured in the context of convolution neural networks (Luo et al., 2016; Dai et al., 2017; Araujo et al., 2019; Raghu et al., 2021; Dosovitskiy et al., 2021) and their dilated variants (Oord et al., 2016; Yu and Koltun, 2016; Chang et al., 2017; Beltagy et al., 2020) with the ultimate goal of receptive field size maximization. Even though we focus on transformer language models, we borrow the idea to show that the empirical receptive field coverage of a model is crucial to its length extrapolation performance.

3 Background and Notations

3.1 Transformer Language Model

Given a sequence of $L \in \{L_{tr}, L_{ex}\}$ input embeddings $\{e_m\}_{m=1}^L$ in \mathbb{R}^d , an R layer transformer language model with H attention heads converts each e_m into its corresponding query, key, and value

vectors in $\mathbb{R}^{\frac{d}{H}}$ at each layer:

$$\mathbf{q}_m = \mathbf{W}_q \mathbf{e}_m, \quad \mathbf{k}_m = \mathbf{W}_k \mathbf{e}_m, \quad \mathbf{v}_m = \mathbf{W}_v \mathbf{e}_m,$$

where $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \in \mathbb{R}^{\frac{d}{H} \times d}$ are learnable matrices. The resulting vectors are processed by the self-attention module for pre-Softmax logits:

$$l_{mn} = \begin{cases} \langle \mathbf{q}_m, \mathbf{k}_n \rangle, & \text{if } m \geq n \\ -\text{inf}, & \text{otherwise} \end{cases}$$

followed by the scaled softmax normalization:

$$a_{m,n} = \frac{\exp(l_{m,n}/\sqrt{d/H})}{\sum_{i=1}^L \exp(l_{m,i}/\sqrt{d/H})} \quad (1)$$

To be precise, the matrices ($\mathbf{W}_q^{(h)}, \mathbf{W}_k^{(h)}, \mathbf{W}_v^{(h)}$), vectors ($\mathbf{q}_m^{(h)}, \mathbf{k}_m^{(h)}, \mathbf{v}_m^{(h)}, \mathbf{o}_m^{(h)}$), and scalars ($l_{mn}^{(h)}, a_{mn}^{(h)}$) are associated with a head number h . For notation simplicity, we only show the dependency on h when we need it. For example, the output vector $\mathbf{o}_m^{(h)}$ at position m for head h is:

$$\mathbf{o}_m^{(h)} = \sum_{n=1}^L a_{m,n}^{(h)} \mathbf{v}_n^{(h)}$$

All the H output vectors are concatenated, denoted by \oplus , and transformed by $\mathbf{W}_o \in \mathbb{R}^{d \times d}$ to obtain $\mathbf{o}_m \in \mathbb{R}^d$:

$$\mathbf{o}_m = \mathbf{W}_o (\mathbf{o}_m^{(1)} \oplus \mathbf{o}_m^{(2)} \oplus \dots \oplus \mathbf{o}_m^{(H)})$$

A layer normalization (Ba et al., 2016) on \mathbf{o}_m , i.e. $\text{LayerNorm}(\mathbf{o}_m)$, gives the input embedding to the next layer. After R layers of propagation, the last \mathbf{o}_m is transformed by $\mathbf{V} \in \mathbb{R}^{v \times d}$ and normalized by Softmax to get the distribution $\mathbf{p} \in \mathbb{R}^v$ over vocabulary size v :

$$\mathbf{p} = \text{Softmax}(\mathbf{V} \mathbf{o}_m) \quad (2)$$

We set $R = 12$, $H = 12$, $d = 768$, and $L_{tr} = 512$ for all experiments reported in this paper.

3.2 ALiBi

ALiBi modifies $l_{m,n}$ to be:

$$l_{mn} = \begin{cases} \langle \mathbf{q}_m, \mathbf{k}_n \rangle - \frac{1}{2^h} (m - n), & \text{if } m \geq n \\ -\text{inf}, & \text{otherwise} \end{cases} \quad (3)$$

The range of h is $n \cdot \frac{8}{H}$, where $n = \{1 \dots H\}$.

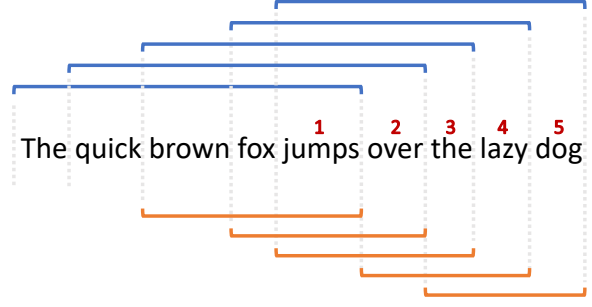


Figure 3: We always evaluate the perplexities of the 5 tokens numbered from 1 to 5. The upper brackets represent $L_{ex} = 5$. The lower brackets represent $L_{ex} = 3$. This formulation ensures the same 5 tokens are always evaluated with different numbers of previous tokens.

3.3 Windowed Attention

If the windowed attention has a size w , then:

$$l_{mn} = \begin{cases} \langle \mathbf{q}_m, \mathbf{k}_n \rangle, & \text{if } n + w > m \geq n \\ -\text{inf}, & \text{otherwise} \end{cases}$$

3.4 Evaluation of Length Extrapolation

We prepare $N = 1000$ text segments of length $L_{ex} > L_{tr}$ from the evaluation dataset. For each segment, we alter the number of previous tokens ranging from 1 to $L_{ex} - 1$ of the last token and only calculate its perplexity:

$$\text{PPL} = \exp \left(\frac{1}{N} \sum_{i=1}^N -\log p_i \right),$$

where p_i is the predicted probability from Eq. (2) of the last (L_{ex} -th) token in the i -th segment. This ensures that the same set of tokens is always used for perplexity calculation and only their number of previous tokens is varied, see Figure 3.²

4 ALiBi and Windowed Attention

Here, we alter the slope ($\frac{1}{2^h}$) of ALiBi to check if the length extrapolation property persists and reveal the connection between ALiBi and windowed attention. We present three experiments on two datasets, ArXiv and OpenWebText2 (Appendix A), to ensure that the observations are consistent across different text domains, shown in Table 1 and 4.

²There exists another evaluation protocol named non-overlapping subsequences adopted in the main experiment tables of ALiBi (Press et al., 2022). It is not the most suitable protocol for length extrapolation evaluation as it suffers from the ‘‘early token’’ curse. Please refer to Appendix B of ALiBi (Press et al., 2022) for details.

L_{ex}	Shift all h by Δ						Same h for all heads					Windowed Attention with Size w					
	$\Delta:-3$	0	2	4	6	8	$h:0$	2	4	6	8	$w:40$	80	100	120	160	320
512	5.76	5.57	5.50	5.63	5.70	5.70	9.45	6.65	5.85	5.60	5.70	8.27	7.28	7.04	6.77	6.41	6.04
1024	7.15	5.64	5.31	5.81	55.4	55.4	9.20	7.01	8.66	25.4	55.4	8.27	7.29	7.02	8.90	67.4	178
2048	7.15	5.94	5.89	6.92	94.4	94.4	9.21	7.08	8.66	31.7	94.4	8.27	7.29	7.03	8.90	67.5	202
4096	7.15	5.95	5.92	6.94	96.0	96.0	9.21	7.08	8.66	31.8	96.0	8.27	7.29	7.02	8.90	67.5	202
8192	7.15	5.95	5.92	6.94	96.0	96.0	9.21	7.08	8.66	31.8	96.0	8.27	7.29	7.02	8.90	67.5	202

Table 1: The three experiments on the Arxiv dataset.

4.1 Slope Shift (Shift all h by Δ)

We first investigated whether slope diversity (each attention head has one slope) is the key to length extrapolation. We shift h by a fixed amount Δ and find that the model, unfortunately, fails to extrapolate beyond a certain quantity. This implies that diversity itself might not be the deciding factor, but that the actual slope value is more important.

4.2 Slope Equalization (Same h for all heads)

To identify the slope magnitude that enables length extrapolation, we set all slopes to be the same instead of the original geometric sequence. We then steadily increase the slope value from 0 to 8 and find that only large slopes ($\frac{1}{2^h}$), or equivalently small h , allow a model to extrapolate well. Large slopes implicitly enforce a narrow windowed bias on the self-attention matrix such that distant tokens cannot interact with each other.

4.3 Windowed Attention (Size w)

We make the implicit window effect explicit as shown by Eq. (3), which is also adopted by Longformer (Beltagy et al., 2020). We define the windowed attention size to be w . The model underperforms at small w and diverges on long L_{ex} at large w . The same trend holds in the first two experiments when h is too small or large.

4.4 Other Observations

First, ALiBi does not in fact extrapolate since its perplexities all increase instead of staying the same when $L_{ex} > L_{tr}$. In contrast, windowed attention models are extrapolatable up to $w = 100$. Second, we can clearly see that once L_{ex} passes a certain threshold, the perplexity either remains the same or explodes. This suggests that the model is either ignoring tokens beyond a certain length (same)³ or not using it properly (explosion). In the next

³A limited but similar observation was made in Appendix B.2 of ALiBi (Press et al., 2022).

section, we will use the concept of receptive field to explain these observations.

5 Receptive Field Measurement

Following the definition of windowed attention size w , an R layer transformer has a theoretical receptive field (TRF) of wR , which is the maximum number of tokens that contribute to the prediction of the next token. In practice, a neural model often uses a subset of TRF, named empirical receptive field (ERF). While previous work (Luo et al., 2016; Dai et al., 2017; Araujo et al., 2019; Raghu et al., 2021; Dosovitskiy et al., 2021; Beltagy et al., 2020) aims to increase ERF to match TRF, we show that decreasing ERF could serve as one feasible approach to enable successful length extrapolation.

Consider the case where $\text{TRF} \leq L_{tr}$: This model can extrapolate easily because its TRF is fully covered and trained. Concretely, if we set $R = 12$, $L_{tr} = 512$ in Table 1 and 4, we know that as long as $w < 42.6 = 512/12$, TRF will be fully covered by L_{tr} . Surprisingly, the model is still able to extrapolate up to $w = 100$, leading to a TRF of $100 * 12 = 1200 \gg 512$. This can be explained by the ERF and TRF discrepancy discussed above; this calls for the need to quantify ERF.

5.1 Quantifying Empirical Receptive Field

We first calculate the normalized gradient (Luo et al., 2016) of each input token w.r.t the prediction of the next token:

$$s_m = \frac{\|\mathbf{g}_m\|_2}{\sum_{n=1}^{L_{ex}} \|\mathbf{g}_n\|_2},$$

where \mathbf{g}_m is the gradient vector of the input embedding e_m . We then calculate the cumulative sum as:

$$c_m = \sum_{n=m}^{L_{ex}} s_n, \quad 0 \leq c_m \leq 1,$$

Visualizations of c_m for the slope shift and windowed attention experiments are shown in Fig-

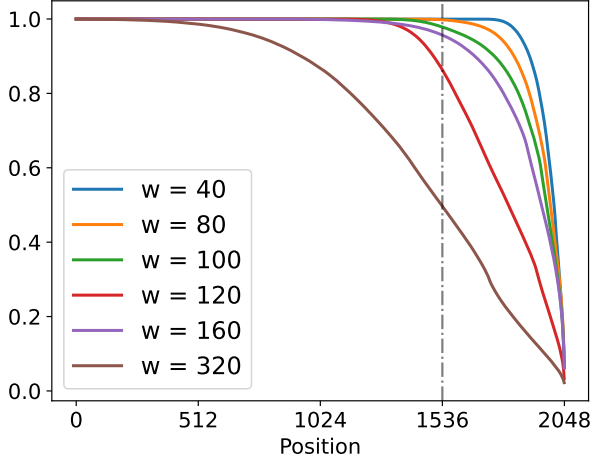


Figure 4: Cumulative normalized gradient on ArXiv when predicting the next (2048-th) token.

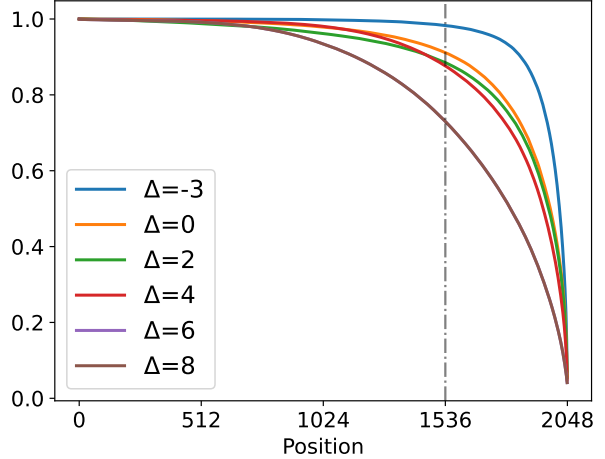


Figure 5: Cumulative normalized gradient on ArXiv when predicting the next (2048-th) token.

ures 4 and 5. We define the ERF of a model as:

$$\text{ERF} = \min\{m \mid c_m > 0.99\}.$$

Figure 4 demonstrates how we derive the model’s ERF when it is predicting the 2048-th token. For models with $w \in [40, 80, 100]$, the most recent $L_{ex} = L_{tr} = 512$ (1536-th to 2047-th) covers more than 99% of the total (1.0) normalized gradient, so their ERF is smaller than 512. In contrast, models with $w \in [120, 160, 320]$ have ERF = 768, 1024, and 1536 tokens, respectively. Since $L_{tr} = 512$ does not fully cover their ERFs, they fail to extrapolate well.

We next focus on the more complex Figure 5, in which neither of the configurations reaches 0.99 within the most recent $L_{tr} = 512$ tokens. Generally, this explains why the perplexity often bumps up when L_{ex} goes from 512 to 1024: Models cannot perfectly process more tokens than they were trained on. If we take a closer look, the $\Delta = -3$ model has the strongest windowing effect and the smallest ERF=768 tokens, therefore its perplexity plateaus the soonest at $L_{ex} = 1024$ in Table 1. The remaining models all need ERF=2048 tokens to reach $c_m = 0.99$, which explains why their perplexities become stable only after $L_{ex} = 2048$ (Table 1). For $\Delta \in [6, 8]$ models specifically, the difference between L_{tr} and ERF is too large to be handled, resulting in exploded perplexities.

5.2 Fixing Failed Cases

We fix the failed cases in Table 1 section 1 (varying Δ) and section 3 (varying w) by increasing L_{tr} to cover their ERFs. We increase L_{tr} to 1024 for

windowed attention with $w = 160$; For shifted ALiBi with $\Delta = 6$, we need $L_{tr} = 2048$ tokens. Table 2 shows that both are now able to maintain stable perplexities.

L_{ex}	Shift all h by $\Delta = 6$		Windowed Attention $w = 160$	
	Arxiv	OpenWebText2	Arxiv	OpenWebText2
2048	4.4	15.2	6.2	19.9
4096	6.2	19.8	6.2	19.9
8192	6.2	19.9	6.2	19.9

Table 2: Fixing failed cases with longer L_{tr} : $L_{tr} = 2048$ for ALiBi with $\Delta = 6$ and $L_{tr} = 1024$ for windowed attention with $w = 160$.

5.3 Analyses of Sinusoidal and Rotary

Sinusoidal (Vaswani et al., 2017) constructs the positional embedding at position m and $\forall i \in [1, d/2]$ as:

$$\begin{aligned} \mathbf{p}_{m,2i} &= \sin\left(\frac{m}{10000^{2i/d}}\right), \\ \mathbf{p}_{m,2i+1} &= \cos\left(\frac{m}{10000^{2i/d}}\right) \end{aligned} \quad (5)$$

They will be added with the input embeddings $\{e_m\}_{m=1}^L$ followed by the query and key transformations as shown in Eq. (4). Unlike addition, Rotary (Su et al., 2021) multiplies each token embedding e_m with a position-specific rotation matrix $\mathbf{R}_m e_m$.

What could c_m tell us when it is applied to the non-extrapolatable Sinusoidal and Rotary positional embeddings? As we can see in Figure 6 and 7, they both fail to focus on the most recent L_{tr}

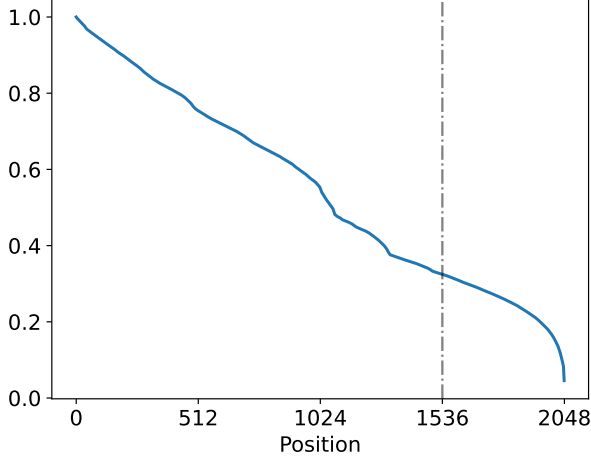


Figure 6: Cumulative normalized gradient of Rotary on ArXiv when predicting the last (2048-th) token with $L_{tr} = 512$.

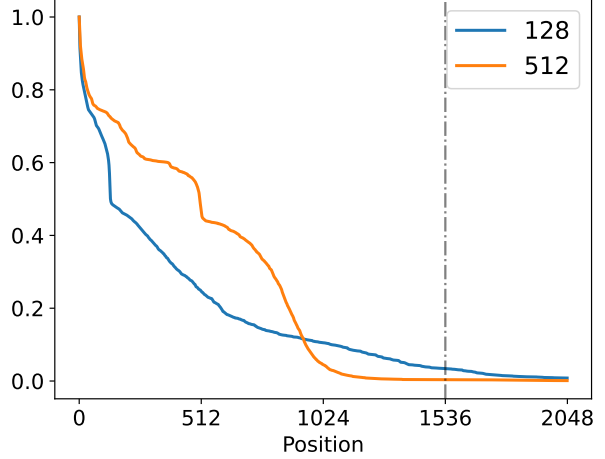


Figure 7: Cumulative normalized gradient of Sinusoidal on ArXiv when predicting the last (2048-th) token with $L_{tr} \in [128, 512]$.

$$\begin{aligned}
 (\mathbf{W}_q(\mathbf{e}_m + \mathbf{p}_m))^\top (\mathbf{W}_k(\mathbf{e}_n + \mathbf{p}_n)) &= \\
 \underbrace{\mathbf{e}_m^\top \mathbf{W}_q^\top \mathbf{W}_k \mathbf{e}_n}_{\text{semantic info.}} + \underbrace{\mathbf{e}_m^\top \mathbf{W}_q^\top \mathbf{W}_k \mathbf{p}_n + \mathbf{p}_m^\top \mathbf{W}_q^\top \mathbf{W}_k \mathbf{e}_n + \mathbf{p}_m^\top \mathbf{W}_q^\top \mathbf{W}_k \mathbf{p}_n}_{\text{mixture of semantic and positional info.}} &\approx \underbrace{\mathbf{e}_m^\top \mathbf{W}_q^\top \mathbf{W}_k \mathbf{e}_n}_{\text{semantic info.}} + \underbrace{\mathbf{p}_m^\top \mathbf{p}_n}_{\text{positional info.}}
 \end{aligned} \tag{4}$$

tokens because neither of their formulations guarantees a L_{tr} -bounded receptive field. Figure 7 tells additional stories: To predict the last token (2048-th), Sinusoidal focuses on the 512-th token when $L_{tr} = 512$ and the 128-th token when $L_{tr} = 128$ as indicated by the sudden jump on their normalized gradient plots. This is because the model has only seen at most L_{tr} positional embeddings and overfitted on them, which provides explicit evidence to the Sinusoidal, or APE in general, overfitting hypothesis made by the author of ALiBi⁴. It also explains why RPE is a better choice for length extrapolatable transformers: They cannot overfit on the positional embeddings.

6 A New RPE for Length Extrapolation

6.1 Introduction to Sandwich

We fix the overfitting issue of Sinusoidal by transforming it into a new RPE, **Sandwich**, shown in Eq. (4). Specifically, we drop the cross terms and keep only the inner product of two positional embeddings⁵ at m and n . Now $\mathbf{p}_m^\top \mathbf{p}_n$ with $m, n \in$

$[1, L]$ become the temporal bias terms of Sandwich:

$$\begin{aligned}
 \mathbf{p}_m^\top \mathbf{p}_n &= \sum_{i=1}^{\bar{d}/2} \sin\left(\frac{m}{10000^{2i/\bar{d}}}\right) \sin\left(\frac{n}{10000^{2i/\bar{d}}}\right) + \\
 &\quad \cos\left(\frac{m}{10000^{2i/\bar{d}}}\right) \cos\left(\frac{n}{10000^{2i/\bar{d}}}\right) \\
 &= \sum_{i=1}^{\bar{d}/2} \cos\left(\frac{m-n}{10000^{2i/\bar{d}}}\right)
 \end{aligned}$$

A similar observation was previously made in a context different from length extrapolation (Yan et al., 2019).

The largest value of $\mathbf{p}_m^\top \mathbf{p}_n$ happens at the point where $m - n = 0$, which gives the maximum value of $\bar{d}/2$. To align L_{tr} with the ERF of Sandwich, we need to further check that $\mathbf{p}_m^\top \mathbf{p}_n$ demonstrates a similar windowed attention effect as ALiBi. This can be done by subtracting all $\mathbf{p}_m^\top \mathbf{p}_n$ by $\bar{d}/2$ and further dividing them by a set of predefined compression ratios. for the sake of simplicity, we set the compression ratios to be the same as ALiBi’s $h = n \cdot \frac{8}{H}$ with $n \in \{1 \dots H\}$:

$$\frac{\mathbf{p}_m^\top \mathbf{p}_n - \bar{d}/2}{h} \tag{6}$$

Eq. (6) is added after the scaled softmax is done in Eq. (1). Figures 8 and 9 show a visualization of

⁴<https://twitter.com/OfirPress/status/1435690039925567489>

⁵We set $\mathbf{p}_{m,n}$ to $2d$ as doing so gives better empirical performance; it only needs to be computed once before training.

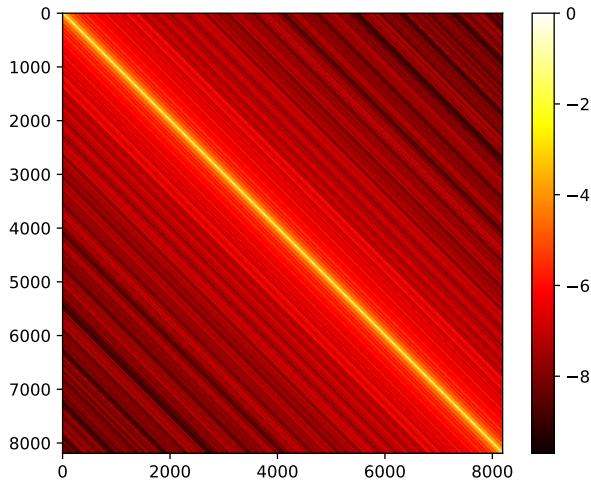


Figure 8: The visualization of Eq. (6) when the compression ratio $h = 8$ and $\bar{d} = 128$.

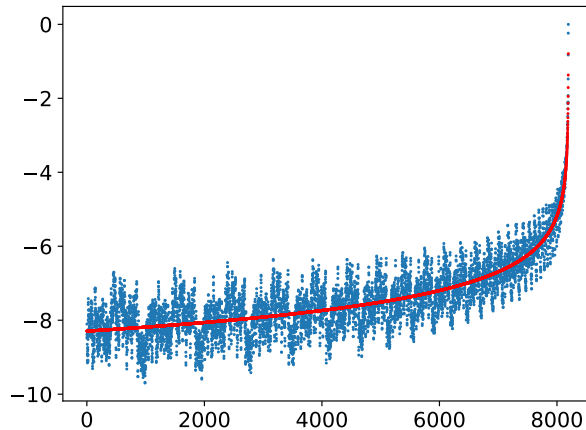


Figure 9: We plot the last row in Figure 8. The red curve is the least-squared fitted log function: $y = -0.825 \cdot \log(|m - n| + 1) - 0.8$ with $m = 8192$ in this example.

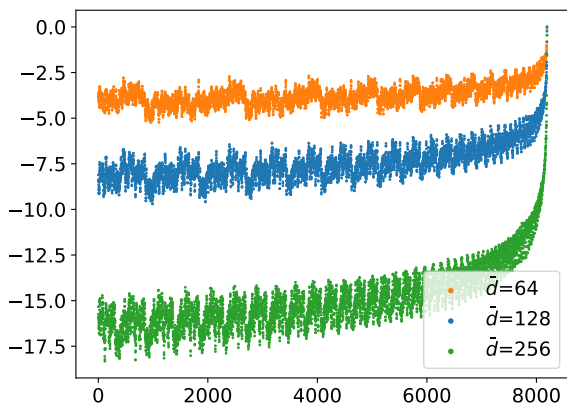


Figure 10: We experiment with different \bar{d} and find they can create different windowed attention effect.

Sandwich when $h = 8$. Sandwich indeed has the same decay-with-distance pattern as ALiBi.⁶

Note that we deliberately decouple this \bar{d} from d in Eq. (5) since we treat \bar{d} as a hyperparameter that controls the shape of Sandwich. A larger \bar{d} leads to a stronger windowed attention effect as shown in Figure 10. We set $\bar{d} = 128$ in this work for all the experiments. We also experiment with smaller and larger \bar{d} and only find worse performance. Finally, readers can find the reference Python implementation in Appendix E.

6.2 Experiments and Discussion

To verify the performance of Sandwich, we train a transformer language model following previous

⁶Fun fact: We imagine different compression ratios as the ways we eat sandwiches: For a huge sandwich, we have to squeeze it more to fit in our mouths!

work (Press et al., 2022; Chi et al., 2022). Table 3 presents the results; the left part contains all models without learnable parameters, and the right part contains models with learnable parameters. These numbers should not be compared across sections.

In general, models on the right achieve lower perplexities across the three datasets. This is expected as they can adapt to individual datasets more easily thanks to the additional learnable parameters. However, there is no free lunch: They often consume more GPU memory and run much slower. For example, T5 is 10% slower than Sandwich during the training stage. Note that Sandwich can also be equipped with learnable parameters such as learnable compression ratios h ; this is left to future work. We now shift our focus to the left section. When $L_{ex} = L_{tr} = 512$, Sandwich is comparable to other models except that Rotary performs a bit better on OpenWebText2. Once we increase L_{ex} , Sandwich begins to reveal its advantages: On ArXiv and GitHub, it is consistently better than all the baselines but only marginally worse than ALiBi when $L_{ex} \geq 4096$ on OpenWebText2.

It is worth mentioning that Sandwich is the first parameter-free RPE that truly makes use of distant token information beyond $L_{tr} = 512$. To see this, notice that lower (better) perplexities occur at $L_{ex} > L_{tr} = 512$. The gradient analysis tool in §5.1 further corroborates this in Figure 11, which reveals a receptive field pattern distinct from that of ALiBi and windowed attention. Even though Sandwich allocates about 60% of the total cumulative gradient on the most recent $L_{tr} = 512$ tokens, distant tokens beyond L_{tr} still contribute substantially

OpenWebText2							
L_{ex}	Sandwich	Smoothed	ALiBi	Sinusoidal	Rotary	KERPLE	T5
512	23.5 ± 3.8	23.2 ± 3.7	22.8 ± 3.3	26 ± 1 [†]	23.0 ± 3.4*	22.6 ± 3.5*	22.6 ± 3.6*
1024	23.0 ± 3.6	23.1 ± 3.6	23.3 ± 3.4	14168 [†]	61 [†]	22.0 ± 3.3*	22.2 ± 3.3*
2048	23.3 ± 3.5	23.2 ± 3.2	23.5 ± 3.3	20370 [†]	96 [†]	21.9 ± 3.1*	23.0 ± 3.1
4096	23.8 ± 3.3	23.6 ± 3.0	23.5 ± 3.3*	42003 [†]	232 [†]	22.1 ± 2.9*	26.8 ± 3.2 [†]
8192	24.7 ± 3.4	24.0 ± 2.9	23.5 ± 3.3*	67869 [†]	343 [†]	22.3 ± 2.9*	38.6 ± 7.2 [†]
ArXiv							
L_{ex}	Sandwich	Smoothed	ALiBi	Sinusoidal	Rotary	KERPLE	T5
512	5.27 ± 0.33	5.33 ± 0.32	5.25 ± 0.33	5.8 [†]	5.25 ± 0.33	5.22 ± 0.37	5.16 ± 0.37*
1024	5.05 ± 0.33	5.13 ± 0.32	5.41 ± 0.36 [†]	1070 [†]	16.02 [†]	4.95 ± 0.34*	4.91 ± 0.35*
2048	5.02 ± 0.34	5.15 ± 0.36	5.58 ± 0.40 [†]	1784 [†]	33.76 [†]	4.83 ± 0.35*	4.92 ± 0.35*
4096	5.15 ± 0.39	5.33 ± 0.39	5.58 ± 0.40 [†]	18050 [†]	71.96 [†]	4.84 ± 0.34*	5.35 ± 0.36
8192	5.28 ± 0.44	5.45 ± 0.42	5.58 ± 0.40 [†]	44100 [†]	111 [†]	4.90 ± 0.33*	6.74 ± 0.90 [†]
GitHub							
L_{ex}	Sandwich	Smoothed	ALiBi	Sinusoidal	Rotary	KERPLE	T5
512	2.88 ± 0.12	2.88 ± 0.17	2.83 ± 0.11 [†]	4 [†]	2.82 ± 0.11	2.81 ± 0.14*	2.76 ± 0.14*
1024	2.71 ± 0.09	2.70 ± 0.07	2.97 ± 0.11 [†]	8342 [†]	3.86 ± 0.25 [†]	2.67 ± 0.10*	2.61 ± 0.08*
2048	2.69 ± 0.11	2.74 ± 0.08	3.01 ± 0.10 [†]	9179 [†]	5.94 ± 0.64 [†]	2.65 ± 0.10*	2.65 ± 0.05
4096	2.73 ± 0.12	2.78 ± 0.08	3.01 ± 0.10 [†]	11017 [†]	11.1 ± 1.55 [†]	2.70 ± 0.09	2.91 ± 0.12
8192	2.79 ± 0.15	2.83 ± 0.08	3.01 ± 0.10 [†]	11270 [†]	20.2 ± 2.75 [†]	2.75 ± 0.08	3.68 ± 0.50 [†]

Table 3: **Perplexity Comparison on the OpenWebText2, GitHub, and ArXiv datasets.** All models are trained for 50k steps with a training length of 512 and five random seeds. The models in the left section have parameter-free positional embeddings. In contrast, both KERPLE and T5 are equipped with learnable parameters. A fair comparison should only be made within the same section. x^\dagger means sandwich is statistically significantly *better* than x . x^* means sandwich is statistically significantly *worse* than x . The test used is paired two-sided t-test with $\alpha = 0.05$. More details about the datasets and hyperparameters are provided in Appendix C and D.

to the model prediction.

Why do ALiBi and windowed attention need to have their ERFs covered by L_{tr} while Sandwich does not? To answer this question, we revisit Figure 9 and approximate (least-squared) the original temporal bias pattern using a log curve, which gives a snug fit⁷: $y = -0.825 \cdot \log(1 + |m - n|) - 0.8$. Table 3 shows its language modeling performance under the “smoothed” column. Pictorially, the log curve decays relatively fast when two tokens are nearby and plateaus when the distance between them increases. In other words, tokens that are far away from the last one ($m = 8192$) share similar temporal biases, possibly leading to beneficial averaging and denoising effects. Note that the averaging effect does not come out of thin air during the extrapolation stage: The almost linear segment ranging from 1536 to 1792 suggests that Sandwich was trained to perform averaging within L_{tr} ; it just needs to average over more historical tokens when

⁷In the actual implementation, we fit the curve using the most recent 50 points of Sandwich. The reason is because the most recent tokens are more important, and we want them to be closer to the original Sandwich.

it extrapolates to longer L_{ex} . In contrast, ALiBi’s linear bias lacks the middle ground to learn the averaging behavior: It either decays so fast that distant tokens are masked out or so slow that the ERF becomes much greater than L_{tr} . The averaging hypothesis also explains why Sandwich, KERPLE, and T5’s perplexities go up in Table 3 instead of continuing to decrease after some L_{ex} (4096 on ArXiv for example): While averaging and denoising improve performance, doing so over too many historical tokens (very large L_{ex}) will reintroduce noises.

6.3 Connection to KERPLE and T5

KERPLE (Chi et al., 2022) has the formulation of $c - r_1 \cdot \log(1 + r_2|m - n|)$. The -0.8 in our fitted log curve term can be absorbed by c , as Softmax is shift-invariant, and if we set $r_1 = 0.825$ and $r_2 = 1$, Sandwich becomes a special case of KERPLE. T5 (Raffel et al., 2020) adopts the log-binning strategy that assigns distinct bins to nearby tokens whereas distant tokens all share the same bin. In spirit, T5 treats distant tokens similarly to Sandwich. Figure 11 verifies that all three of them

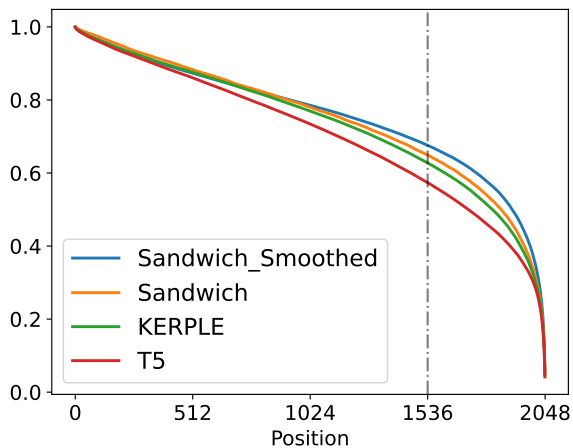


Figure 11: Cumulative normalized gradient of Sandwich, Smoothed Sandwich, KERPLE, and T5 on ArXiv when predicting the last (2048-th) token with $L_{tr} = 512$.

share a similar empirical receptive field pattern.

7 Conclusion

In this paper, we first establish the connection between ALiBi and windowed attention through their constructions and language modeling performance. We then develop a cumulative normalized gradient tool to measure the empirical receptive field. It shows that length extrapolation of ALiBi and windowed attention is possible when the training sequence length covers the empirical receptive field. It also reveals the models’ limitation of not utilizing information beyond the training sequence length. Fortunately, this is overcome by our new relative positional embedding, Sandwich, which is simplified from the earliest proposed Sinusoidal positional embedding. Finally, Sandwich demonstrates a log-decaying temporal bias pattern similar to that previously seen in the design of KERPLE and T5, and such pattern is likely to be the secret to successful length extrapolation. Together these findings supports more effective design of future extrapolatable transformer language models.

Limitations

Although Sandwich, KERPLE, and T5 use information beyond training sequence length, their receptive fields still highly favor the most recent tokens. While this recency bias is beneficial to the modeling of human-written text, it is problematic in other scenarios.

Let us consider the task of *parity* prediction: A model needs to predict whether a bit string has

an even or odd number of ones. For example, the parity of [1, 1, 0, 1] is odd (or 1) and the parity of [1, 0, 1, 0] is even (or 0). Unlike human-written text, every single bit is equally important. Transformer language models with current RPEs still struggle on this simple task (Anil et al., 2022). Its difficulty can be explained by the recency bias effect that we described. Devising a new positional embedding or transformer model architecture that solves this problem is a promising direction for future work.

Ethics Statement

Our work advances the understanding of positional embeddings adopted in almost all transformer models. In addition, our proposed new positional embedding significantly reduces energy consumption and training cost thanks to its length extrapolation property. Finally, our work lays the groundwork for developing future transformers that are greener and more cost-efficient enabled by improved length extrapolation. Inappropriate usage of our technique might have negative societal impacts. These include the ethical challenges of improper text generation and privacy issues inherent in the data collection process. These implications apply to any natural language processing research and are not unique to this specific work.

Acknowledgment

The authors acknowledge the support from Boeing (2019-STU-PA-259), Amazon (CC ADV 00474341 2021 TR), NSF MRI Award 1919452, and Princeton Research Computing.

References

- Joshua Ainslie, Santiago Ontanon, Chris Alberti, Václav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. [ETC: Encoding long and structured inputs in transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 268–284, Online. Association for Computational Linguistics.
- Alex Andonian, Quentin Anthony, Stella Biderman, Sid Black, Preetham Gali, Leo Gao, Eric Hallahan, Josh Levy-Kramer, Connor Leahy, Lucas Nestler, Kip Parker, Michael Pieler, Shivanshu Purohit, Tri Songz, Wang Phil, and Samuel Weinbach. 2021. [GPT-NeoX: Large Scale Autoregressive Language Modeling in PyTorch](#).
- Cem Anil, Yuhuai Wu, Anders Johan Andreassen, Aitor Lewkowycz, Vedant Misra, Vinay Venkatesh

- Ramasesh, Ambrose Slone, Guy Gur-Ari, Ethan Dyer, and Behnam Neyshabur. 2022. [Exploring length generalization in large language models](#). In *Advances in Neural Information Processing Systems*.
- Andre Araujo, Wade Norris, and Jack Sim. 2019. [Computing receptive fields of convolutional neural networks](#). *Distill*. <https://distill.pub/2019/computing-receptive-fields>.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Shiyu Chang, Yang Zhang, Wei Han, Mo Yu, Xiaoxiao Guo, Wei Tan, Xiaodong Cui, Michael Witbrock, Mark A Hasegawa-Johnson, and Thomas S Huang. 2017. [Dilated recurrent neural networks](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ta-Chung Chi, Ting-Han Fan, Peter J Ramadge, and Alexander I Rudnicky. 2022. Kerple: Kernelized relative positional embedding for length extrapolation. *arXiv preprint arXiv:2205.09921*.
- Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. 2017. [Deformable convolutional networks](#). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 764–773.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Ankit Gupta and Jonathan Berant. 2020. [GMAT: global memory augmentation for transformers](#). *CoRR*, abs/2006.03274.
- Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. 2016. Understanding the effective receptive field in deep convolutional neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 4905–4913, Red Hook, NY, USA. Curran Associates Inc.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Interspeech*, volume 2, pages 1045–1048. Makuhari.
- Tomas Mikolov and Geoffrey Zweig. 2012. [Context dependent recurrent neural network language model](#). In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 234–239.
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. [Wavenet: A generative model for raw audio](#). Cite arxiv:1609.03499.
- Ofir Press. 2022. [The use case for relative position embeddings](#).
- Ofir Press, Noah Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *International Conference on Learning Representations*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. 2021. [Do vision transformers see like convolutional neural networks?](#) In *Advances in Neural Information Processing Systems*.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Hang Yan, Bocao Deng, Xiaonan Li, and Xipeng Qiu. 2019. Tener: adapting transformer encoder for named entity recognition. *arXiv preprint arXiv:1911.04474*.
- Fisher Yu and Vladlen Koltun. 2016. [Multi-scale context aggregation by dilated convolutions](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). *CoRR*, abs/2007.14062.

Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. [Recurrent neural network regularization](#). Cite arxiv:1409.2329.

A Results on OpenWebText2

Table 4 includes the three experiments conducted in §4 on OpenWebText2. Their corresponding receptive field plots are shown in Figure 12 and 13.

B Efficient Inference

Although ALiBi might not be using token information further than L_{tr} , it has the nice property of efficient inference (Press, 2022). Tables 1 and 4 show that ALiBi perplexities stay constant when $L_{ex} \geq 2048$. This suggests a cache window size $\bar{w} = 2048$ for inference. The generation of the first \bar{w} tokens remains the same, and we can still cache all \mathbf{q}_m , \mathbf{k}_m , and \mathbf{v}_m vectors for $m \in [1, 2048]$. When it comes to generating the $\bar{w} + 1$ -th token, we simply discard the first cached \mathbf{q}_1 , \mathbf{k}_1 , and \mathbf{v}_1 and use the rest of $\bar{w} - 1$ tokens along with the newly added token to perform self-attention. If we want to generate a length L_{ex} text snippet, the complexity is $O(\bar{w} \times L_{ex})$ instead of $O(L_{ex}^2)$. This complexity is also better than that of an APE model, which is $O(\bar{w}^2 \times L_{ex})$ since an APE model needs to completely re-encode the previous \bar{w} vectors when generating new tokens following the first \bar{w} ones.

We implement the process discussed above to verify that ALiBi indeed allows for efficient inference. The results, along with ones for Sandwich, are presented in Table 5. Both ALiBi and Sandwich permit efficient inference by setting $\bar{w} = 2048$. It is worth pointing out that the performance of Sandwich at $L_{ex} = 4096$ becomes a bit worse compared to that in Table 3. This is more evidence that Sandwich is using longer than L_{tr} token information.

C Scientific Artifacts

We use the gpt-neox library (Andonian et al., 2021) under Apache-2.0 license and the datasets (Gao et al., 2020) released by the authors of gpt-neox. The codebase and datasets (Table 6) are publicly released for research purposes. The steps taken to protect the privacy and anonymization are discussed in Gao et al. (2020) section 6 and 7. Finally, Gao et al. (2020) section 5 also discusses the distribution and statistics of the datasets used in this work.

D Implementation Details

The configurations and hyperparameters are outlined in Table 7. The pretraining takes 5 hours on

a single NVIDIA A-100 GPU. We do not tune any hyperparameters and just use the default ones.

L_{ex}	Shift all h by Δ						Same h for all heads					Windowed Attention Size w					
	$\Delta:-3$	0	2	4	6	8	$h:0$	2	4	6	8	$w:40$	80	100	120	160	320
512	18.6	19.0	19.5	20.0	20.5	20.5	32.7	22.2	19.7	19.7	20.5	25.3	23.7	23.1	24.0	22.9	21.9
1024	21.6	19.3	19.6	24.8	232	232	32.8	23.2	24.9	146	232	25.3	23.7	23.2	137	234	353
2048	21.6	19.7	20.5	29.3	299	299	32.8	23.2	24.9	165	299	25.3	23.7	23.2	137	236	408
4096	21.6	19.7	20.5	29.4	299	299	32.9	23.2	24.9	165	299	25.3	23.7	23.2	137	236	408
8192	21.6	19.7	20.5	29.4	299	299	32.9	23.2	24.9	165	299	25.3	23.7	23.2	137	236	408

Table 4: The three experiments on the OpenWebText2 dataset.

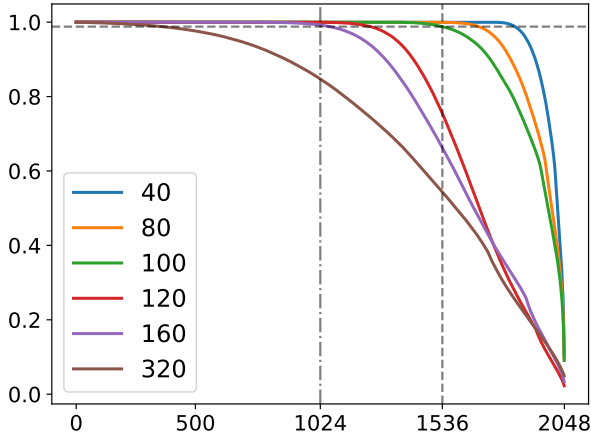


Figure 12: Cumulative normalized gradient on OpenWebText2 when predicting the last (2048-th) token. Windowed Attention Size $w =$

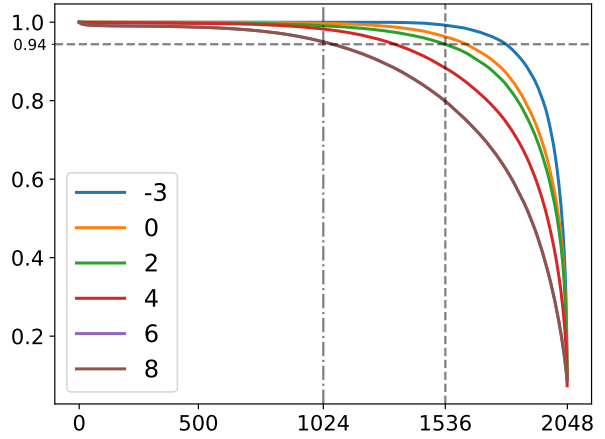


Figure 13: Cumulative normalized gradient on OpenWebText2 when predicting the last (2048-th) token. Shift all h by $\Delta =$

L_{ex}	OpenWebText2		Arxiv		GitHub	
	Sandwich	ALiBi	Sandwich	ALiBi	Sandwich	ALiBi
4096	23.9	23.5	5.31	5.59	2.79	3.01
8192	24.1	23.5	5.35	5.59	2.81	3.01
16384	24.1	23.5	5.35	5.59	2.81	3.01

Table 5: Efficient Inference with $\bar{w} = 2048$.

	OpenWebText2	GitHub	ArXiv
Raw Size	66.77 GB	95.16 GB	56.21 GB
Type	Internet	Coding	Academic

Table 6: **Dataset Overview.** Raw Size is the size before any up- or down-sampling.

# Layers	Hidden Size	# Attention Heads	Train Seq. Len.	# Trainable Params.
12	64	12	512	162M
Optimizer	Batch Size	Train Steps	Precision	# Trainable Params. for RPEs
Adam (lr 6e-4)	32	50,000	bfloat16	0

Table 7: 162M Model Configurations.

E Python Implementation of Sandwich

```
import numpy as np

base = 1e4
heads = 12
seq_len = 8192
positions = np.arange(seq_len)[..., None]
bar_d = 128 # This is the hyperparameter of Sandwich
i = np.arange(bar_d // 2)

pos_embs = np.concatenate([np.sin(positions / base ** (2 * i / bar_d)),
                           np.cos(positions / base ** (2 * i / bar_d))],
                           axis=-1)
sandwich = np.matmul(pos_embs, pos_embs.T)
compression_ratio = np.arange(1, heads + 1) * 8 / heads
multi_head_sandwich = sandwich[None, ...] / compression_ratio[None, None]
```

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations section
- A2. Did you discuss any potential risks of your work?
Ethics Statement
- A3. Do the abstract and introduction summarize the paper’s main claims?
abstract and section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

section 6

- B1. Did you cite the creators of artifacts you used?
appendix C
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
appendix C
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
appendix C
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
appendix C
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
appendix C
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
appendix C

C Did you run computational experiments?

section 6

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
section 6 and appendix D

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
section 6 and appendix D

C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
section 6

C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Not applicable. Left blank.

D **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Left blank.

D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
No response.

D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
No response.

D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
No response.

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
No response.

D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
No response.