

Extending Input Contexts of Language Models through Training on Segmented Sequences

Petros Karypis
UC San Diego
pkarypis@ucsd.edu

Julian McAuley
UC San Diego
jmcauley@ucsd.edu

George Karypis
University of Minnesota
karypis@umn.edu

Abstract

Effectively training language models on long inputs poses many technical challenges. As a cost consideration, language models are pre-trained on a fixed sequence length before being adapted to longer sequences. We explore various methods for adapting models to longer inputs by training on segmented sequences and an interpolation-based method for extending absolute positional embeddings. We develop a training procedure to extend the input context size of pretrained models with no architectural changes and no additional memory costs than training on the original input lengths. By sub-sampling segments from long inputs while maintaining their original position the model is able to learn new positional interactions. Our method benefits both models trained with absolute positional embeddings, by extending their input contexts, as well as popular relative positional embedding methods showing a reduced perplexity on sequences longer than they were trained on. We demonstrate our method can extend input contexts by a factor of $4\times$ while improving perplexity.

1 Introduction

Transformer-based models (Vaswani et al., 2017) capture sequence information through positional embeddings (PE). There are two types of PEs: absolute and relative. Absolute positional embeddings (APE) learn a separate embedding for each position in a sequence; these embeddings are added to the input of the first layer. Relative positional embeddings (RPE) encode the relative distance between positions, often by weighting attention score of positions further away less.

The ability for models to process long sequences efficiently is of growing importance as models become more capable. Increased input context allows for more complex in-context learning examples (Li et al., 2023a; Sun et al., 2023). Additionally, they allow for question answering and summarization

over scientific papers and patents (Dasigi et al., 2021; Koh et al., 2022; Sharma et al., 2019). Due to RPE’s positional information only being a function of relative distance these methods can be applied to any input sequence length. In practice, popular RPE methods fail to generalize to sequences longer than they were trained on. Furthermore, self-attention’s memory cost is quadratic meaning training on long sequences becomes prohibitively expensive as the sequence length grows.

In this work, we study the problem of extending the input context of pre-trained decoder-only transformer-based models, considering those that use either absolute or relative positional embeddings. We show that an interpolation-based approach allows APE models to extrapolate to sequence lengths longer than they were trained on—matching or outperforming the extrapolation ability of RPE methods like ALiBi (Press et al., 2021) and RoPE (Su et al., 2021). To further improve the ability of these models to take advantage of the longer input context, we present resource-efficient methods that continuously pre-train APE- and RPE-based models on carefully sampled segmented sub-sequences of long sequences. Doing so simulates training on long sequences while remaining within a fixed input length. This allows the models to efficiently learn the embeddings of the newly created absolute positions or the relative embeddings associated with the longer pairwise distances.

We experiment with models trained with APEs, RoPE, and ALiBi to verify our method improves the extrapolation performance independent of the choice of positional embeddings. Results show that interpolating the embedding matrix of absolute positional embeddings without any additional training allows for extrapolation to sequences $5\times$ the original input context. Furthermore, our segment-based methods are able to increase the extrapolation ability of all positional embedding approaches. When applied to APEs this method achieves 87% the per-

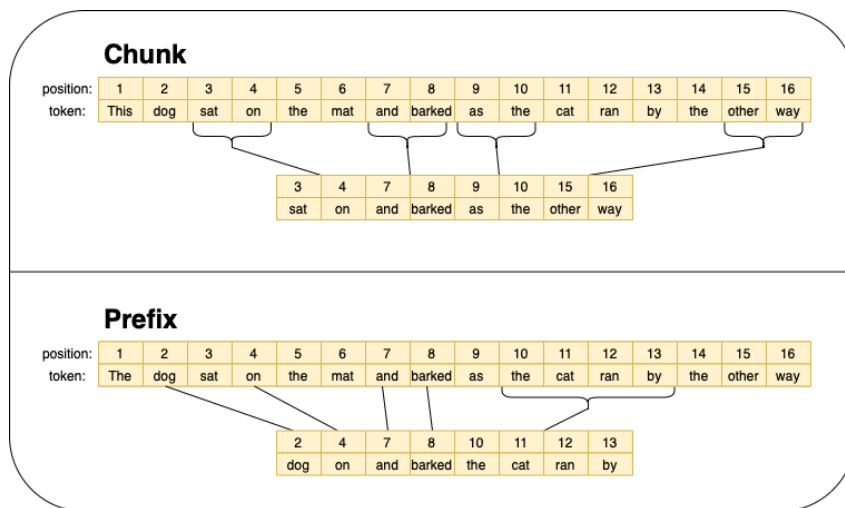


Figure 1: Visualization of our various segment-based methods. We sub-sampling tokens from the original sequence while maintaining the original positions.

formance of training on sequences twice as long at no extra memory footprint.

The paper is organized as follows: first, we conduct a review of various existing literature that motivated our approach. Second, we formally define the problem of length extrapolation and propose our methods for efficiently extending a model’s input context. Third, we provide a detailed breakdown of our experimental setup and methodology to enable reproducibility. Finally, we present our results along with a thorough discussion and analysis.

2 Related Work

2.1 Positional embeddings

Language is inherently sequential and Transformers are positional-agnostic, to account for this, positional information is often introduced to the architecture. The original authors Vaswani et al. (2017) suggested adding a positional embedding to the input of the first layer and offered two methods, absolute positional embeddings and sinusoidal embeddings. Absolute positional embeddings consist of a learnable embeddings matrix where each embedding corresponds to a position. While common, this method has an important limitation: it only allows for a fixed maximum input length determined during training. Sinusoidal embeddings did not have this limitation but performed worse in practice and the relative embeddings that came after were difficult to parallelize (Shaw et al., 2018) leading to APEs being the *de facto* method in early models, eg. BERT (Devlin et al., 2019) and GPT-3 (Brown

et al., 2020).

To address the limited input context size of APE researchers explored other relative positional embedding methods (Chi et al., 2022; Wennberg and Henter, 2021; Likhomanenko et al., 2021; Haviv et al., 2022). Most notable are rotary embeddings (RoPE) (Su et al., 2021), T5 (Raffel et al., 2019), and ALiBi (Press et al., 2021). RoPE rotates the query and the key embeddings as a function of their position; this method allowed for easier parallelization compared to previous relative embeddings. T5 bias (Raffel et al., 2019) adds a positional embedding for each relative distance instead of absolute position. ALiBi subtracts a linear bias from the query-key matrix product in the attention calculations. While T5 bias extrapolated to long contexts well it is too inefficient to scale, taking twice as long to train as sinusoidal (Press et al., 2021). RoPE and ALiBi have been widely adopted in various LLMs with LLaMA (Touvron et al., 2023), GPT-J (Wang and Komatsuzaki, 2021), and PaLM (Chowdhery et al., 2022) using RoPE and BLOOM (Scao et al., 2022) using ALiBi.

2.2 Length generalization

The choice of positional embeddings (PE) has been documented to be one of the leading factors in a Transformer based model’s ability to generalize to variable sequence lengths. The authors of ALiBi (Press et al., 2021) identified that RoPE and sinusoidal embeddings failed to generalize on sequence lengths greater than those they were trained on. Numerous new positional embedding methods with more favorable length generalization abilities

have been proposed (Sun et al., 2022; Chi et al., 2022; Li et al., 2023b) but these are required to be incorporated during pre-training.

There is a sizable body of work on methods for extending the input context of language models pre-trained with RoPE (Chen et al., 2023; Jin et al., 2024; Peng et al., 2023; Ding et al., 2024). These approaches map the positional information of long sequences into ranges seen during training through positional interpolation. In practice, these methods requires fine-tuning the models on long sequences to adjust to the new granularity of relative positional distance which is computational expensive.

2.3 Computationally efficient training

Numerous works have explored efficiency based modifications to the standard Transformer architecture (Xiong et al., 2021b; Choromanski et al., 2020; Kitaev et al., 2020; Qiu et al., 2019). These methods either modify the base architecture or rely on fast self-attention approximations.

While these methods all aim to reduce the memory cost of the Transformer architecture and allow for training on longer sequences, our work is orthogonal to these methods. Our approach can be used in conjunction with these existing methods since we do not rely on any specific architecture. We instead change the positional information of the input sequences.

2.4 Sparse input sequences

A number of works have explored training language models on sparse inputs. APEs have been shown to overfit to certain positions. To address this, Kiyono et al. (2021) proposed randomly padding or offsetting the positions during fine-tuning. This simple method led to better downstream performance on question answering and machine translation (Tao et al., 2023) and general length extension (Zhu et al., 2023; Ruoss et al., 2023). Another work proposed Forgetful Causal Masking (FCM) (Liu et al., 2022), a simple modification to the next token prediction task with a randomly selected fraction of previous tokens masked out. They demonstrated this method led to improvements in both few-shot and fine-tuned performance compared to standard causal masking. Most similar to ours, RandomPos (Ruoss et al., 2023) proposed sampling randomized, ordered positional embeddings to replace the sequential positional embeddings normally used. They sampled from a range of absolute positions much longer than the input se-

quence length. Results demonstrated this led to an increase in extrapolation performance. The authors argued this was due to exposure to longer relative pair-wise distances than those normally seen during training.

These results indicate that not only can language models be trained with heavily obfuscated sequences but can also benefit from doing so in some cases. This idea is the intuition behind our method.

3 Methods

There are three reasons that motivate this work. First, there exist numerous high-quality pre-trained models whose input context is limited to 1K–2K tokens. Extending the input context of these models will further increase their applicability. Second, even though methods that rely on relative positional embeddings can operate on input contexts that are longer than what they were trained on, their out-of-the-box extrapolation performance is not good (Press et al., 2021). Third, due to the quadratic complexity of self-attention and the linear compute/memory complexity of transformers w.r.t. sequence length, direct training on long input contexts is resource intensive. This limits the input context that we can directly train on.

3.1 Problem Statement

Let p_θ be a transformer-based language model trained to maximize the next-token-probabilities over a set of sequences \mathcal{D} of length L_t ; i.e.,

$$\arg \max_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \sum_i^{L_t} \log p_\theta(x_i | x_{<i}). \quad (1)$$

We will refer to L_t as the model’s *training input context length*.

We define *extrapolation* as the language model’s ability to improve its next-token-prediction by using input contexts that are longer than those it trained on. Specifically, for $k > L_t$, we will consider that a model can extrapolate successfully if

$$\sum_{i \geq k} \log p_\theta(x_i | x_{>k}) > \sum_{i \geq k} \log p_\theta(x_i | x_{>L_t}),$$

where $p_\theta(x_i | x_{>j}) = p_\theta(x_i | x_{i-1}, \dots, x_{i-j+1})$. In practice, we consider the average perplexity on sequences of different lengths from the same dataset a suitable proxy for this.

Given p_θ and L_t , the problem that we want to solve is to develop resource efficient methods that

allow p_θ to extrapolate to input contexts of length L_e that are longer than L_t . We refer to L_e as the *extended input context length*.

3.2 Extending APE via interpolation

APEs learn an embedding vector for each position up to a pre-specified maximum position. The fixed nature of the embedding matrix does not allow for inputs longer than the maximum pre-specified length. A necessary first step when training on longer sequences is to increase the size of the embedding matrix.

We use linear interpolation to extend the embedding matrix from the training input context length L_t to the new input context length L_e (Dehghani et al., 2023). Let E and E' be the old and new embedding matrices, respectively and assume that $\beta = L_e/L_t$ is integral. Then the embedding for position i ($0 \leq i < L_e$) is given by:

$$e'_i = \frac{\beta - i \% \beta}{\beta} e_{\lfloor i/\beta \rfloor} + \frac{i \% \beta}{\beta} e_{\lfloor i/\beta \rfloor + 1},$$

where ‘%’ is the modulo operation. This process retains the original embeddings but results in $\beta(L_t - 1) + 1$ embeddings. In practice, we set the remaining $\beta - 1$ embeddings to e_{L_t} .

3.3 Efficient input context extension

Pairwise attention is the mechanism by which transformer models incorporate information from other tokens. Positional embeddings are how attention takes into account the absolute or relative positions of the token-pairs. To fully take advantage of an increased input context, a model needs to learn the embeddings of the newly created absolute positions or the relative embeddings associated with the longer pairwise distances created with the increased input context. Thus, the model needs to be further pre-trained with input sequences that also include the new positions—in the case of absolute positional embeddings, or the longer pairwise distances—in the case of relative positional embeddings.

The key insight behind our efficient approaches is that we can meet the above requirements without directly training on long input sequences. Instead, we create short input sequences by sampling segments from the long sequences, *keep the original positional information*, concatenate them, and use them to further pre-train the language model. Since this approach retains the original positional information, the models see the new positions/distances

and learn how to use them. Though the length of the short sequence is a hyper-parameter of our approach, in all of our experiments we keep it the same as that of the original input context length; i.e., L_t .

We develop two different subsequence sampling approaches that we refer to as *chunk* and *prefix* which are defined as follows:

- **chunk- α :** This approach creates a short sequence by sampling a small number of equal-length contiguous subsequences from the long sequence. Specifically, given $0 < \alpha < 1$ and an L_e -long input sequence \mathbf{x} , this approach samples $1/\alpha$ contiguous non-overlapping subsequences of length αL_t from \mathbf{x} . The reason that we keep the sampled segments contiguous is to preserve the local context information, which is important for next-token prediction (Xiong et al., 2021a) and we do not want our model to ‘unlearn’ it.
- **prefix- α :** This approach creates a short sequence by randomly sampling a set of tokens that forms a prefix and a contiguous segment to form its associated suffix. Specifically, given $0 < \alpha < 1$ and an input sequence \mathbf{x} of length L_e , it randomly selects an index i with $(1 - \alpha)L_t < i < L_e - \alpha L_t$. It creates the suffix by taking the αL_t contiguous tokens starting at position i and creates the prefix by randomly sampling $(1 - \alpha)L_t$ tokens from the positions preceding i . In this method we only compute the loss over the continuous suffix in order to preserve the model’s ability to incorporate local context.

A visualization of the different sampling methods can be found in Figure 1.

While these methods can introduce discontinuities in the causal language modeling objective we argue that maintaining their original positional embedding on top of the fact they happen infrequently limits the harm they may cause. In practice we use α ’s small enough that discontinuities occurs approximately 2% of the time in *chunk* and never in *prefix*.

4 Experimental setup

4.1 Dataset

Since we are comparing the performance of various methods on long sequences we chose to use the scientific papers section of the arXiv dataset released

by Cohan et al. (2018). Scientific papers are a common choice for reporting results on long sequence modeling performance (Beltagy et al., 2020). This dataset consists of 215K scientific papers, split into 205K train and 7K test, with a total token count of approximately 1.6 billion and an average document length of 4,938 tokens. We do not pack our batches (Kosec et al., 2021), meaning each sequence contains only text from a single document at a time. If documents are longer than L_e we split them into non-overlapping sequences with length L_e and discard the remainder; if documents are shorter than L_e we discard them as well. We feel that ensuring each input only corresponds to one source text is an important factor when reporting performance on long sequences.

4.2 Models

To evaluate our methods we fine-tune three different classes of pretrained language models, one for each of the popular positional embedding methods: absolute, RoPE, and ALiBi. We use models with approximately 1.5 billion parameters; for absolute positional embeddings we use GPT-2 (Radford et al., 2019), for rotary embeddings we use Pythia (Biderman et al., 2023), and for ALiBi we use Bloom (Scao et al., 2022). In addition to these three models we use a smaller GPT-2 and Pythia checkpoint (approx. 10% the size), which we will refer to as GPT-2 Small and Pythia Small, and together as our development models. Due to a lack of small models trained with ALiBi we do not have a development model for ALiBi. Key information about these models can be found in Table 1. Note that besides the positional encoding schemes, these models also differ in other ways including training data and model parameters. As a result, a direct comparison of these models will be confounded by these additional factors. For this reason our evaluation only focuses on measuring how the different continuous pre-training approaches help in improving each model’s extrapolation capabilities against themselves and we never compare across models.

4.3 Domain adaptation

The perplexity on arXiv for these models is relatively high as arXiv is considered out of domain. In order to differentiate between gains attributed to adapting to the domain versus improving extrapolation performance we perform one full epoch of continual pre-training with a sequence length of L_t for each model.

Table 1: Key model characteristics.

	# of params	PE	L_t
GPT-2 Small	170M	APE	1024
Pythia Small	140M	RoPE	2048
GPT-2	1.64B	APE	1024
Pythia	1.4B	RoPE	2048
Bloom	1.45B	ALiBi	2048

Table 2: Perplexity on sequences of the model’s original input length, L_t , after domain adaptation.

	ppl.
GPT-2 Small	9.311
Pythia Small	8.609
GPT-2	6.675
Pythia	6.677
Bloom	7.217

We refer to the checkpoints after domain-adaptation as "out-of-the-box" models. All experiments start from the OOTB models unless otherwise mentioned. The perplexity of the models after domain adaptation can be found in Table 2.

4.4 Segmented pre-training

For training we use the causal language modeling objective with a cross entropy loss. All experiments on the same model are done in a compute equivalent manner unless stated otherwise. To ensure compute equivalence when training our models we fix the number of tokens as well as the input length, L_t , of the model.

Due to segmentation, one epoch of training on different sequence lengths results in a different number of tokens actually processed. For example, training with sequences of length $2L_t$ results in half the total number of tokens. To ensure an equal number of tokens across experiments we set the total number of epochs for each experiment to be:

$$\# \text{ epochs} = \frac{L_e}{L_t}. \quad (2)$$

4.5 Performance assessment

To evaluate the performance of our models on different sequence lengths we report the mean perplexity on sequence of length L_e from our test set. Perplexity measures the exponentiated average negative log likelihood over a sequence of tokens and

is a common evaluation metric for language models. We define the perplexity of a sequence of tokens \mathbf{x} of length L_t as:

$$ppl(\mathbf{x}) = \exp\left(-\frac{1}{L_t} \sum_i \log p_\theta(x_i | x_{<i})\right). \quad (3)$$

Note that unlike previous work, we do not perform sliding window evaluation (Baeviski and Auli, 2018).

5 Results

We conduct our experiments and present results in such a way to answer the following questions:

- How well do absolute positional embeddings extrapolate with interpolation of the embeddings matrix?
- Which of our proposed subsequence sampling methods performs the best and with what parameters?
- How does our approach compare with continual pre-training on sequences of the original length?

5.1 Out-of-the-box extrapolation

We begin by examining each model’s ability to extrapolate to sequences longer than they were trained on without any further pre-training. We report the perplexity on the test set with sequence lengths starting from L_t up to $5L_t$, depending on the memory constraints of each. Previous length extrapolation work did not include absolute positional embeddings due to their fixed nature (Press et al., 2021). To increase the input context size we interpolated the positional embedding matrix as described in Section 3.2. Results are shown in Figure 2 and the corresponding numbers can be found in Table 6 in Appendix A.

RoPE fails to extrapolate to sequences longer than originally trained on while ALiBi generalizes well. These findings about RPEs agree with those previously observed in Press et al. (2021). Our results show that interpolation works well until at least $5L_t$. This suggests that with linear interpolation APEs generalize better than RoPE and are comparable to ALiBi.

5.2 Comparison of segmented methods

We compare the performance of the various methods discussed in Section 3.3 on our development




Figure 2: Perplexity of "out-of-the-box" extrapolation. With interpolation of the positional embeddings, absolute positional embeddings (APE) extrapolate as well as ALiBi.

models. We train models on two separate extension sizes, $L_e = 2L_t$ and $L_e = 4L_t$. For each we use chunk with $\alpha = \{0.125, 0.25, 0.5\}$ and prefix with $\alpha = \{0.25, 0.5\}$. Furthermore, we train a models on sequences of $2L_t$ and $2L_t$ without any segmentation. We refer to these models as full, and they provide a point of comparison between our methods versus training on the full L_e sequence. The complete set of results can be found in Table 3.

The different segment-based methods work well to extend the input context of these models. We observe a decrease in perplexity when evaluating on sequences longer than originally trained on. Overall, chunk performs better than prefix on both models, prefix fails to improve extrapolation when extending RoPE to sequences $4\times$ in length. While the full approach has the lowest perplexity in most cases the relative loss in performance for chunk is low. One notable case is extending RoPE to $4L_t$, there we observe chunk outperforming full. Given that chunk requires half the sequence length of full it remains a competitive option due to its memory efficiency.

Comparing the performance of different chunk lengths, controlled by the parameter α , both models display similar trends. For chunk, there appears to be sweet-spot between the number of segments and each segment’s individual length (see Table 3). An α of 0.125 translates to chunks of 128 tokens for APE and 256 for RoPE. In most cases this α performed the worst amongst chunk, as the segments may be too short or lead to too many discontinuities in the sequence. For prefix, there is less of

Table 3: Perplexity of different input context length extension methods on the development sets.

	method	$2L_t$	$4L_t$
APE	OOTB	9.322	13.275
	full	8.287	7.819
	chunk-0.125	8.521	8.307
	chunk-0.25	8.471	7.989
	chunk-0.5	8.420	8.259
	prefix-0.25	8.757	8.826
	prefix-0.5	8.672	9.304
RoPE	OOTB	30.686	176.244
	full	7.403	7.353
	chunk-0.125	7.476	7.239
	chunk-0.25	7.447	7.210
	chunk-0.5	7.461	7.461
	prefix-0.25	9.543	25.539
	prefix-0.5	10.119	33.375

a concrete pattern. This could be due to the higher level of randomness in the prefix as tokens were sampled randomly. Between chunk and prefix, chunk computes loss over twice as many tokens, this could be a contributing factor to the gap in performance between the two.

Between RoPE and APE, RoPE benefits the most from segmented pre-training. After training on segmented sequences the perplexity on extensions of $2L_t$ and $4L_t$ decreases by a factor of $4\times$ and $24\times$ respectively. While our method still improves over the "out-of-the-box" performance of APEs, interpolation is a competitive approach for length extension.

5.3 Results on larger models

Based off the findings in Section 5.2 we use chunk-0.25 for our experiments on GPT-2 1.5B, Pythia-1.4B, and, Bloom-1.1B. As before, we continually pre-train the models as detailed in Section 4.4 and expand to $L_e = 2L_t$ and $L_e = 4L_t$.

Overall, chunk works for all three models on both expansion lengths. All models extrapolated better than their "out-of-the-box" performance. Again, RoPE was able to extrapolate to sequences it previously was not able to. Our method also demonstrated the ability to further increase the extrapolation ability of ALiBi. Results can be found in Table 4.

Table 4: Perplexity results for the 1.x billion parameter models.

	method	$2L_t$	$4L_t$
APE	OOTB	6.326	7.099
	DA	6.125	7.050
	chunk-0.25	6.314	6.425
RoPE	OOTB	16.428	52.644
	DA	16.285	50.652
	chunk-0.25	5.448	5.278
ALiBi	OOTB	7.295	7.773
	DA	6.887	7.417
	chunk-0.25	6.773	7.295

5.4 Comparison with further pre-training

Given that ALiBi and APE-based models already extrapolate well (see Figure 2), a natural question is whether the performance gains on longer sequences come from our segmented method or additional domain adaption. To ablate this, we perform another epoch of domain adaptation as described in Section 4.3. This isolates the benefit of our method versus further domain adaptation as the total number of tokens seen by all models are the same. Results can be found in Table 4.

For models that extrapolate well (ALiBi and APE), further domain adaptation also improves the extrapolation ability however the gains are less than our segmented training. The exception here is when extending APE to lengths $2\times$, in this case domain adaption performs slightly better. This result indicates that the interpolation-based extension method we propose works well for APEs. Overall, this demonstrates that while some of the gains may be due to further domain adaptation our method is still beneficial for models that extrapolate well "out-of-the-box".

5.5 Comparison with RandomPos

The authors of RandomPos (Ruoss et al., 2023) proposed a similar method for simulating training on long sequences within a fixed input context window. Instead of subsampling sequences of length L_e , RandomPos randomized the positional ids of sequences of length L_t selecting positions ranging from $[0, L_e - 1]$ while maintaining the causal ordering. Similar to our approach, RandomPos exposes the model to extrapolated pairwise relative distances but the key difference is content used. Whereas RandomPos only presents local context to

Table 5: Comparison with RandomPos. Numbers reported are perplexity.

method		$2L_t$	$4L_t$
APE	OOTB	9.322	13.275
	RandomPos	9.018	11.534
	chunk-0.25	8.420	7.989
RoPE	OOTB	30.686	176.244
	RandomPos	8.021	11.692
	chunk-0.25	7.447	7.210
ALiBi	OOTB	7.295	7.773
	DA	6.816	7.352
	chunk-0.25	6.773	7.295

the model, chunk exposes the model to distant content and encourages the model to learn to leverage distant contexts.

To verify the exposure to distant content is an important step in improving extrapolation we implement a version of RandomPos and extend our models to $2\times$ and $4\times$ the original input sizes. We keep all settings and models the same as Section 5.2 with the exception of including the ALiBi model. In all cases, chunk outperforms RandomPos indicating the inclusion of distant context valuable to length extrapolation. Results can be found in Table 5.

6 Analysis

Our results demonstrate that segmented training is a viable approach to extend the input context size of language models. It is not immediately intuitive why, especially given that the relative positional embeddings methods are not learned.

For absolute positional embeddings the reasoning is fairly straightforward. First, in Section 5.1 we demonstrated interpolating the embedding matrix led to reasonable extrapolation without any training. Before any training occurs the model already has some extrapolation ability. The segmented sequences allow for positions further away than the input size normally allows to interact and learn how to incorporate information.

In the case of relative positional embedding methods these results are less intuitive. Both RPE methods penalize the attention scores of positions as a function of their relative distance, meaning that initially there is not much attention across chunk boundaries. We hypothesize that through training on segmented sequences the model learns to at-




Figure 3: Histogram of median attention weights for positions past the original input length before and after our segmented training on models with RoPE. After adaptation, the distribution of attention weights becomes more uniform.

tend to longer-range interactions. There is a lack of nearby positions for the model to attend to so it learns to incorporate information from further away. In doing so it adjusts the weights to penalize further positions less. This counteracts the RPE’s inductive bias towards nearby positions.

To attempt to visualize this we plot the distribution of median attention weights for positions past L_t . In both cases, the medians are well below the mean suggesting that a few positions account for the majority of the attention weight. After segmented training, we observe the average median increases as well as become more evenly distributed. This suggests that more positions are being attended to as well as the model attending to more or less positions depending on the context. The plot can be found in Figure 3. This hypothesis is also supported by a recent work that analyzes the failure of RoPE to generalize to long sequences (Xiong et al., 2023). The observed that simply reducing the decaying effect of RoPE distant tokens lead to strong extrapolation performance.

7 Conclusion

In this work we proposed a simple and memory efficient approach to extend the effective input context size of models through training on sequences created by sampling segments from long documents. We demonstrated our method is robust to the choice of positional embeddings and allows models to be trained on sequences at least $4\times$ their original input length. Furthermore, our results on extending absolute positional embeddings through interpolation demonstrated they can extrapolate better than

RoPE and provide a method to extend the context of models trained with APEs at no additional cost.

8 Limitations

In this work we explore various computationally efficient methods for pre-training on long sequences. Due to the compute limitations we only verify our method’s performance on models up to 1.4 billion parameters. Current state of the art models are orders of magnitudes larger. While our results indicate the success of our method there is always the chance that results do not transfer to different model sizes. We believe these methods will hold as model size increases since the extrapolation problem is fundamentally an artifact of the positional embeddings and not model size. Additionally, the models we used were originally only trained with a maximum sequence length up to 2048 tokens and only extended to a maximum 8192 tokens. Even though this is a $4\times$ extension, this is much lower than the input size of some production models.

Inline with previous work on encoding positional information (Press et al., 2021; Su et al., 2021), we use perplexity as our method for evaluating a model’s extrapolation performance. Some recent work has shown that this may not always be a strong signal for downstream performance (Shaham et al., 2022). A more thorough evaluation on downstream benchmarks would be insightful, unfortunately the majority of our models were too weak to produce competitive performance on zero-shot or few-shot long sequence tasks.

9 Ethics statement

When working with language models and large, web-crawled datasets it is important to remain cognizant of some of the potential ethical concerns. We trained on scientific papers which are voluntarily posted by users.

References

Alexei Baevski and Michael Auli. 2018. Adaptive input representations for neural language modeling. *ArXiv*, abs/1809.10853.

Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*, abs/2004.05150.

Stella Rose Biderman, Hailey Schoelkopf, Quentin G. Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron,

Lintang Sutawika, and Oskar van der Wal. 2023. Pythia: A suite for analyzing large language models across training and scaling. *ArXiv*, abs/2304.01373.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *ArXiv*, abs/2005.14165.

Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. [Extending context window of large language models via positional interpolation](#). *ArXiv*, abs/2306.15595.

Ta-Chung Chi, Ting-Han Fan, Peter J. Ramadge, and Alexander I. Rudnicky. 2022. Kerple: Kernelized relative positional embedding for length extrapolation. *ArXiv*, abs/2205.09921.

Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás Sarról, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy J. Colwell, and Adrian Weller. 2020. Rethinking attention with performers. *ArXiv*, abs/2009.14794.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Benton C. Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier García, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Díaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#). *ArXiv*, abs/2204.02311.

Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, W. Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *North American Chapter of the Association for Computational Linguistics*.

- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. *ArXiv*, abs/2105.03011.
- Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Steiner, Mathilde Caron, Robert Geirhos, Ibrahim M. Alabdulmohsin, Rodolphe Jenatton, Lucas Beyer, Michael Tschannen, Anurag Arnab, Xiao Wang, Carlos Riquelme, Matthias Minderer, Joan Puigcerver, Utku Evci, Manoj Kumar, Sjoerd van Steenkiste, Gamaleldin F. Elsayed, Aravindh Mahendran, Fisher Yu, Avital Oliver, Fantine Huot, Jasmijn Bastings, Mark Collier, Alexey A. Gritsenko, Vighnesh Birodkar, Cristina Nader Vasconcelos, Yi Tay, Thomas Mensink, Alexander Kolesnikov, Filip Pavetić, Dustin Tran, Thomas Kipf, Mario Luvcić, Xiaohua Zhai, Daniel Keysers, Jeremiah Harmsen, and Neil Houlsby. 2023. Scaling vision transformers to 22 billion parameters. *ArXiv*, abs/2302.05442.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [Longrope: Extending llm context window beyond 2 million tokens](#). *ArXiv*, abs/2402.13753.
- Adi Haviv, Ori Ram, Ofir Press, Peter Izsak, and Omer Levy. 2022. Transformer language models without positional encodings still learn positional information. In *Conference on Empirical Methods in Natural Language Processing*.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia yuan Chang, Huiyuan Chen, and Xia Hu. 2024. [Llm maybe longlm: Self-extend llm context window without tuning](#). *ArXiv*, abs/2401.01325.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *ArXiv*, abs/2001.04451.
- Shun Kiyono, Sosuke Kobayashi, Jun Suzuki, and Kentaro Inui. 2021. Shape: Shifted absolute position embedding for transformers. In *Conference on Empirical Methods in Natural Language Processing*.
- Huan Yee Koh, Jiaxin Ju, Ming Liu, and Shirui Pan. 2022. An empirical survey on long document summarization: Datasets, models, and metrics. *ACM Computing Surveys*, 55:1 – 35.
- Matej Kosec, Shengyu Fu, and Mario Michael Krell. 2021. [Packing: Towards 2x nlp bert acceleration](#). *ArXiv*, abs/2107.02027.
- Mukai Li, Shansan Gong, Jiangtao Feng, Yiheng Xu, Jinchao Zhang, Zhiyong Wu, and Lingpeng Kong. 2023a. In-context learning with many demonstration examples. *ArXiv*, abs/2302.04931.
- Shanda Li, Chong You, Guru Guruganesh, Joshua Ainslie, Santiago Ontanon, Manzil Zaheer, Sumit K. Sanghai, Yiming Yang, Sanjiv Kumar, and Srinadh Bhojanapalli. 2023b. [Functional interpolation for relative positions improves long context transformers](#). *ArXiv*, abs/2310.04418.
- Tatiana Likhomanenko, Qiantong Xu, Ronan Collobert, Gabriel Synnaeve, and Alexey Rogozhnikov. 2021. Cape: Encoding relative positions with continuous augmented positional embeddings. In *Neural Information Processing Systems*.
- Hao Liu, Xinyang Geng, Lisa Lee, Igor Mordatch, Sergey Levine, Sharan Narang, and P. Abbeel. 2022. Fcm: Forgetful causal masking makes causal language models better zero-shot learners. *ArXiv*, abs/2210.13432.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2023. [Yarn: Efficient context window extension of large language models](#). *ArXiv*, abs/2309.00071.
- Ofir Press, Noah A. Smith, and Mike Lewis. 2021. Train short, test long: Attention with linear biases enables input length extrapolation. *ArXiv*, abs/2108.12409.
- Jiezhong Qiu, Hao Ma, Omer Levy, Scott Yih, Sinong Wang, and Jie Tang. 2019. Blockwise self-attention for long document understanding. *ArXiv*, abs/1911.02972.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, abs/1910.10683.
- Anian Ruoss, Grégoire Del’etang, Tim Genewein, Jordi Grau-Moya, R. Csordás, Mehdi Abbana Bennani, Shane Legg, and Joel Veness. 2023. [Randomized positional encodings boost length generalization of transformers](#). *ArXiv*, abs/2305.16843.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elizabeth-Jane Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Rose Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurenceon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa Etxabe, Alham Fikri Aji, Amit Alfassy, Anna Rogers,

Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris C. Emezue, Christopher Klamm, Colin Leong, Daniel Alexander van Strien, David Ifeoluwa Adelani, Dragomir R. Radev, Eduardo Gonz'alez Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady ElSahar, Hamza Benyamina, Hieu Trung Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jorg Frohberg, Josephine L. Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro von Werra, Leon Weber, Long Phan, Loubna Ben Allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, Mar'ia Grandury, Mario vSavsko, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad Ali Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gilbert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla A. Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto L'opez, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, S. Longpre, So-maieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesh Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Sai-ful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal V. Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Févry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiang Tang, Zheng Xin Yong, Zhiqing Sun, Shaked Brody, Y Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre Francois Lavall'ee, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Re-vena, Suraj Patil, Tim Dettmers, Ahmed Baruwaa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aur'elie N'ev'eol, Charles Lovering, Daniel H Garrette, Deepak R. Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Xiangru Tang, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian

Gehrmann, Shachar Mirkin, S. Osher Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdenvek Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ananda Santa Rosa Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behrooz, Benjamin Olusola Ajibade, Bharat Kumar Saxena, Carlos Muñoz Fer-randis, Danish Contractor, David M. Lansky, Davis David, Douwe Kiela, Duong Anh Nguyen, Edward Tan, Emily Baylor, Ezinwanne Ozoani, Fatim T Mirza, Frankline Ononiwu, Habib Rezanejad, H.A. Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jan Passmore, Joshua Seltzer, Julio Bonis Sanz, Karen Fort, Livia Macedo Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, M. K. K. Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nourhan Fahmy, Olanrewaju Samuel, Ran An, R. P. Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas L. Wang, Sourav Roy, Sylvain Viguier, Thanh-Cong Le, Tobi Oyebade, Trieu Nguyen Hai Le, Yoyo Yang, Zachary Kyle Nguyen, Abhinav Ramesh Kashyap, A. Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Kumar Singh, Benjamin Beilharz, Bo Wang, Caio Matheus Fonseca de Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel Le'on Perin'an, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabc, Iman I.B. Bello, Isha Dash, Ji Soo Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthi Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, María Andrea Castillo, Marianna Nezhurina, Mario Sanger, Matthias Samwald, Michael Cullan, Michael Weinberg, M Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myung-sun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patricia Haller, R. Chandrasekhar, R. Eisenberg, Robert Martin, Rodrigo L. Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aaroonsiri, Srishti Kumar, Stefan Schweter, Sushil Pratap Bharati, T. A. Laud, Th'eo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yashasvi Bajaj, Y. Venkatraman, Yifan Xu, Ying Xu, Yun chao Xu, Zhee Xiao Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2022. Bloom: A 176b-parameter open-access multilingual language model. *ArXiv*, abs/2211.05100.

Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, and Omer Levy. 2022. [SCROLLS: Standardized CompaRison over long language sequences](#). In *Proceedings of the 2022 Con-*

- ference on Empirical Methods in Natural Language Processing*, pages 12007–12021, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Eva Sharma, Chen Li, and Lu Wang. 2019. Bigpatent: A large-scale dataset for abstractive and coherent summarization. *ArXiv*, abs/1906.03741.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *North American Chapter of the Association for Computational Linguistics*.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *ArXiv*, abs/2104.09864.
- Simeng Sun, Y. Liu, Shuo Wang, Chenguang Zhu, and Mohit Iyyer. 2023. Pearl: Prompting large language models to plan and execute actions over long documents. *ArXiv*, abs/2305.14564.
- Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shao-han Huang, Alon Benhaim, Vishrav Chaudhary, Xia Song, and Furu Wei. 2022. [A length-extrapolatable transformer](#). *ArXiv*, abs/2212.10554.
- Mingxu Tao, Yansong Feng, and Dongyan Zhao. 2023. A frustratingly easy improvement for position embeddings via random padding. *ArXiv*, abs/2305.04859.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971.
- Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Ulme Wennberg and Gustav Eje Henter. 2021. The case for translation-invariant self-attention in transformer-based language models. *ArXiv*, abs/2106.01950.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oğuz, Madian Khabsa, Han Fang, Yashar Mehdad, Sharan Narang, Kshitiz Malik, Angela Fan, Shruti Bhosale, Sergey Edunov, Mike Lewis, Sinong Wang, and Hao Ma. 2023. [Effective long-context scaling of foundation models](#). *ArXiv*, abs/2309.16039.
- Wenhan Xiong, Barlas Ouguz, Anchit Gupta, Xilun Chen, Diana Liskovich, Omer Levy, Wen tau Yih, and Yashar Mehdad. 2021a. Simple local attentions remain competitive for long-context tasks. In *North American Chapter of the Association for Computational Linguistics*.
- Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Moo Fung, Yin Li, and Vikas Singh. 2021b. Nyströmformer: A nyström-based algorithm for approximating self-attention. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 35 16:14138–14148.
- Dawei Zhu, Nan Yang, Liang Wang, Yifan Song, Wenhao Wu, Furu Wei, and Sujian Li. 2023. [Pose: Efficient context window extension of llms via positional skip-wise training](#).

Table 6: Perplexity of "out of the box" extrapolation for models with APE, RoPE, and ALiBi positional embeddings.

(ppl.)	1×	2×	3×	4×	5×
APE	6.675	6.326	6.394	7.099	8.438
RoPE	6.677	17.348	45.797	69.288	-
ALiBi	7.217	7.295	7.653	7.773	-

A Full Results