

# LONGEMBED: Extending Embedding Models for Long Context Retrieval

Dawei Zhu \*<sup>♥♠</sup> Liang Wang <sup>◇</sup> Nan Yang <sup>◇</sup> Yifan Song <sup>♥♠</sup> Wenhao Wu <sup>♥♠</sup>  
Furu Wei <sup>◇</sup> Sujian Li <sup>♥♠♠</sup>

<sup>♥</sup> School of Computer Science, Peking University

<sup>♠</sup> National Key Laboratory for Multimedia Information Processing, Peking University

<sup>♣</sup> Jiangsu Collaborative Innovation Center for Language Ability, Jiangsu Normal University

<sup>◇</sup> Microsoft Corporation

{dwzhu, lisujian}@pku.edu.cn wangliang@microsoft.com

<https://github.com/dwzhu-pku/LongEmbed>

## Abstract

Embedding models play a pivotal role in modern NLP applications such as document retrieval. However, existing embedding models are limited to encoding short documents of typically 512 tokens, restrained from application scenarios requiring long inputs. This paper explores context window extension of existing embedding models, pushing their input length to a maximum of 32,768. We begin by evaluating the performance of existing embedding models using our newly constructed LONGEMBED benchmark, which includes two synthetic and four real-world tasks, featuring documents of varying lengths and dispersed target information. The benchmarking results highlight huge opportunities for enhancement in current models. Via comprehensive experiments, we demonstrate that training-free context window extension strategies can effectively increase the input length of these models by several folds. Moreover, comparison of models using Absolute Position Encoding (APE) and Rotary Position Encoding (RoPE) reveals the superiority of RoPE-based embedding models in context window extension, offering empirical guidance for future models. Our benchmark, code and trained models will be released to advance the research in long context embedding models.

## 1 Introduction

Text embeddings are vector representations of natural language that encode its semantic information. They play a pivotal role in various natural language processing (NLP) tasks, including information retrieval (IR) and retrieval-augmented generation (RAG). However, embedding models for producing these vector representations still operates within a very narrow context window, many supporting only 512 input tokens (Wang et al., 2022; Xiao et al., 2023; Ni et al., 2022). This narrow

context window has greatly hindered their application in scenarios requiring long inputs, such as long Wikipedia articles and meeting scripts (Saad-Falcon et al., 2024).

Previous efforts that train a long context embedding model *from scratch* suffer significant computational overhead, due to the combined demand for large batch sizes and long sequences. For example, Chen et al. (2024) utilized 96 A100 GPUs to train BGE-M3 which supports 8k context. Meanwhile, there have been many successes in extending context window of *existing* LLMs in a plug-and-play way or via efficient fine-tuning, pushing their context from 4k to 128k (Xiong et al., 2023) and even 2 million tokens (Ding et al., 2024). Motivated by this, instead of training long context embedding models from scratch, this paper explores context window extension of *existing* embedding models.

First, we examine the capability of existing embedding models in processing long context. Retrieval is selected as the proxy task, as it closely mirrors real-world application scenarios. While there have been some retrieval benchmarks such as BEIR (Thakur et al., 2021) and LoCo (Saad-Falcon et al., 2024), we identify two major limitations with these existing benchmarks: 1) limited document length, 2) biased distribution of target information. To overcome this, we introduce the LONGEMBED benchmark that integrates two synthetic tasks to enable flexible control over document length, and four real tasks featuring dispersed target information. Results on LONGEMBED indicates huge room for improvement in current embedding models.

Based on this, we explore plug-and-play strategies to extend embedding models, including parallel context windows, reorganizing position ids, and position interpolation. Comprehensive experiments show that these strategies can effectively extend the context window of existing embedding models by several folds, regardless of their original context being 512 or beyond 4k. Furthermore,

\* Contribution during Dawei’s internship at MSR Asia. Sujian Li is the corresponding author.




Figure 1: (a) Overview of the LONGEMBED benchmark. (b) Performance of current embedding models on passkey retrieval, with evaluation length ranging from 256 to 32,768<sup>1</sup>.  $\blacktriangle$  /  $\blacklozenge$  denotes embedding models with 512 /  $\geq 4k$  context. The greener  $\blacksquare$  a cell is, the higher retrieval accuracy this model achieves on the corresponding evaluation length. (c) Effects of context window extension methods on E5, E5-RoPE, E5-Mistral, measured by improvements of Avg. Scores on LONGEMBED. SE / NTK is short for SelfExtend / NTK-Aware Interpolation.

for models employing absolute position encoding (APE), we show the possibility of harvesting further improvements via fine-tuning while strictly preserving original behavior within the short context. In this way, we have extended  $E5_{\text{Base}}$  (Wang et al., 2022) from 512 to 4k (See Figure 1c).

For models utilizing RoPE (Su et al., 2021), substantial enhancements on LONGEMBED are observed when employing methods that fully leverage RoPE’s advantages, such as NTK (Peng and Quesnelle, 2023) and SelfExtend (Jin et al., 2024). As illustrated in Figure 1b and 1c, leveraging NTK extends the context window of E5-Mistral to 32k, achieving close-to-perfect accuracy on passkey retrieval and state-of-the-art performance on LONGEMBED. Further, for fair comparison of APE / RoPE-based embedding models, we pre-train E5-RoPE following the training procedure and data of E5. Thorough comparison of E5 and E5-RoPE reveals the superiority of RoPE-based embedding models in context window extension. To sum up, our contributions are as follows:

- We construct LONGEMBED to benchmark long context retrieval, which includes two synthetic and four real-world tasks, featuring documents of varying lengths and dispersed target information.
- We have conducted comprehensive experiments on training-free context window extension, extending the input length of existing embedding models by several folds.
- We reveal the superiority of RoPE-based embedding models in context window extension via

thorough comparison of models adopting APE and RoPE, offering empirical guidance for future embedding models.

- Our benchmark and trained models ( $E5_{\text{Base}}-4k$ ,  $E5\text{-RoPE}_{\text{Base}}$ ) will be released to advance the research in long context embedding models.

## 2 Related Work

**Text Embedding Models.** Text embeddings encode semantic information of text as low-dimensional vectors, enabling numerous NLP applications. Early attempts on embeddings models include latent semantic indexing (Deerwester et al., 1990) and weighted average of word embeddings (Mikolov et al., 2013). Modern embedding models (Wang et al., 2022; Xiao et al., 2023; Nee-lakantan et al., 2022) exploit supervision from labeled query-document pairs, adopting a multi-stage training paradigm that pre-trained on large-scale raw text pairs using contrastive loss, then fine-tuned on small scale but high-quality datasets.

Existing efforts in developing long-context embedding models typically involve first obtaining a long-context backbone model, either by pre-training with long inputs from scratch (Günther et al., 2023; Nussbaum et al., 2024; Chen et al., 2024) or using existing ones (Wang et al., 2023b), followed by training the backbone model to produce embeddings. Instead, this paper endows *ex-*

<sup>1</sup>For simplicity, we report results from the *base* versions of the included models by default. The supported context length of each model is presented in Table 2. Inputs exceeding the supported context length are truncated.

isting embedding models with the ability to handle long context through context window extension.

**Context Window Extension for LLMs.** Due to the high cost of pre-training an LLM from scratch, there have been many efforts towards extending the context window of *existing* LLMs in a plug-and-play manner. We categorize these efforts as follows: 1) *Divide-and-conquer*, which involves segmenting long inputs into short chunks, processing each chunk with the model, and aggregating the results, as demonstrated by PCW (Ratner et al., 2023); 2) *Position reorganization*, which reorganizes position ids to accommodate longer inputs, as exemplified by SelfExtend (Jin et al., 2024), DCA (An et al., 2024). 3) *Position interpolation*, which introduces new position embeddings by interpolating existing ones, includes PI (Chen et al., 2023), NTK (Peng and Quesnelle, 2023), YaRN (Peng et al., 2023), and Resonance RoPE (Wang et al., 2024a). Our paper thoroughly investigates these three lines of methods on embedding models. We also acknowledge other efforts in context extension, such as token compression (Jiang et al., 2023; Ge et al., 2023; Zhang et al., 2024a) and memory-based transformers (Wang et al., 2024b; Xiao et al., 2024). However, the former is not applicable for bidirectional attention, and the latter requires complex mechanisms for accessing encoded content, hence we do not experiment with these two categories.

In addition to their plug-and-play usability, further fine-tuning on top of these methods with long training samples has been proven to yield better performance (Xiong et al., 2023; Fu et al., 2024; Zhang et al., 2024b; Yen et al., 2024). Addressing the overhead of training on long inputs and the scarcity of extremely long training data, a line of research investigates simulating long inputs within short context, including Randomized Positions (Rouss et al., 2023), Positional Skip-wise (PoSE) training (Zhu et al., 2023), and SkipAlign (Wu et al., 2024). This paper also leverage these efforts to synthesize long training samples from the original training data, facilitating further fine-tuning on top of plug-and-play methods.

### 3 The LONGEMBED benchmark

In this section, we first identify two limitations of existing retrieval benchmarks for evaluating long-context capabilities (§ 3.1). Then, we introduce the retrieval tasks adopted in our LONGEMBED, including both synthetic ones (§ 3.2) and real ones (§ 3.3).




Figure 2: Results of E5<sub>Base</sub> on 8 LoCo tasks that are publicly available.

#### 3.1 Examining Existing Retrieval Benchmarks

There are two main desiderata for curating a long context retrieval benchmark. First, the candidate documents should be long enough. Second, the target information to answer user query should be as uniformly distributed across the document as possible. This prevents embedding models from solely focusing on specific parts, such as the beginning (Coelho et al., 2024), to achieve unreasonably high scores. Based on these criteria, we examine existing retrieval benchmarks as follows:

**BEIR Benchmark** (Thakur et al., 2021) is a collection of 18 information retrieval datasets, ranging across ad-hoc web search, question answering, fact verification, etc. However, documents in this benchmark contains fewer than 300 words on average (See Table 5 in Appendix), making it unsuitable for measuring long context retrieval that usually involves documents of thousands or tens of thousands of words.

**LoCo Benchmark** (Saad-Falcon et al., 2024) consists 12 retrieval tasks that requires long context reasoning, spanning diverse domains such as law and finance. However, it still suffers from biased distribution of key information, as demonstrated in Figure 2. With only 512 context length, E5<sub>Base</sub> achieves >85% nDCG scores on 3 out of 8 publicly-available LoCo tasks. This severely biased distribution of target information undermines its ability to reflect model performance as context increases.

#### 3.2 Synthetic Tasks in LONGEMBED

First, we introduce the passkey and needle retrieval task for embedding models as follows:

**Personalized Passkey Retrieval.** Passkey retrieval (Mohtashami and Jaggi, 2023) requires LLMs to recover a random passkey hidden within a long document comprising garbage information. For embedding models, we adopt the *personal-*

Dataset	Domain	# Queries	# Docs	Avg. Query Words	Avg. Doc Words
<i>Real Tasks</i>					
NarrativeQA	Literature, Film	10,449	355	9	50,474
QMSum	Meeting	1,527	197	71	10,058
2WikiMultihopQA	Wikipedia	300	300	12	6,132
SummScreenFD	ScreenWriting	336	336	102	5,582
<i>Synthetic Tasks</i>					
Passkey	Synthetic	400	800	11	†
Needle	Synthetic	400	800	7	†

Table 1: Overview of the LONGEMBED benchmark. Average word number is rounded to the nearest integer. † For needle and passkey test, we have 8 groups of queries and candidate documents, with the documents averaging  $\{0.25, 0.5, 1, 2, 4, 8, 16, 32\} \times 0.75k$  words, respectively.

#### Passkey Test Examples:

Query: What is the pass key for **Sky Morrow**?

Doc1: <prefix> **Sky Morrow**'s passkey is **123**. Remember it. **123** is the passkey for **Sky Morrow**. <suffix>

Doc2: <prefix> **Cesar McLean**'s passkey is **456**. Remember it. **456** is the passkey for **Cesar McLean**. <suffix>

...

#### Needle Test Examples:

Query: Who discovered the **law of gravity**?

Doc1: <prefix> **The law of gravity was discovered by Sir Issac Newton**. <suffix>

Doc2: <prefix> **The best thing to do in San Francisco is eat a sandwich and sit in Dolores Park on a sunny day**. <suffix>

...

Figure 3: Example for the passkey and needle test. For the passkey test, the <prefix / suffix> are repeats of "The grass is green. The sky is blue. The sun is yellow. Here we go. There and back again." For the needle test, the <prefix> and <suffix> form a long essay.

ized passkey retrieval (Wang et al., 2023b), where each document contains a unique person name and his/her passkey at random position. The goal is to retrieve the document containing the given person’s passkey from all candidates documents.

**Needle-in-a-haystack Retrieval.** While passkey retrieval surrounds key information with garbage sentences, needle-in-a-haystack retrieval (Kamradt, 2023; Liu et al., 2024) randomly inserts key information into an arbitrary position of a long essay, making the task more challenging. To tailor this task for embedding models, we instruct GPT-4 to generate 100 facts covering a variety of domains including physics, history, geometry, art, etc, and 100 queries correspondingly. The facts are subsequently treated as *needles* and randomly inserted into the PaulGrahamEssay to form 100 candidate

documents. Our task is to correctly retrieve the document that contains corresponding needle given the query.

The advantage of synthetic data is that we can flexibly control context length and distribution of target information. For both tasks, we evaluate a broad context range of  $\{0.25, 0.5, 1, 2, 4, 8, 16, 32\} \times 1,024$  tokens<sup>2</sup>. For each context length, we include 50 test samples, each comprising 1 query and 100 candidate documents.<sup>3</sup> In this way, we can measure the effective context size of embedding models for up to 32k tokens. Examples for both tasks are in Figure 3.

### 3.3 Real Tasks in LONGEMBED

While synthetic tasks offer flexibility in manipulating context length and distributing target information, they still differ from real-world scenarios. To conduct a comprehensive evaluation, we have tailored following long-form QA and summarization tasks for long context retrieval. For QA datasets, we use the questions as queries, the set of all input documents as candidate documents. For summarization datasets, we use the summaries as queries, and the set of all input documents as candidate documents.

**NarrativeQA** (Kočíský et al., 2018) is a QA dataset comprising long stories and corresponding questions about specific content such as characters,

<sup>2</sup>Since token numbers vary w.r.t. tokenizers, we use a rough estimation that 1 token = 0.75 word, and constraint the word numbers to not exceed  $\{0.25, 0.5, 1, 2, 4, 8, 16, 32\} \times 1,024 \times 0.75$ .

<sup>3</sup>The original version of personalized passkey retrieval uses different candidate documents for each query, resulting in 50 queries and 5,000 documents to encode for each context length. To speed up evaluation, we share the candidate documents for different queries within each context length.

events. As these details are dispersed throughout the story, models must process the entire long context to get the correct answers.

**2WikiMultihopQA** (Ho et al., 2020) is a multi-hop QA dataset featuring questions with up to 5 hops, synthesized through manually designed templates to prevent shortcut solutions. This necessitates the ability to process and reason over long context, ensuring that answers cannot be obtained by merely focusing on a short span within the document.

**QMSum** (Zhong et al., 2021) is a query-based meeting summarization dataset that requires selecting and summarizing relevant segments of meetings in response to queries. Due to the involvement of multiple participants and topics in the meeting, summarization regarding specific queries naturally requires aggregating information dispersed throughout the entire text.

**SummScreenFD** (Chen et al., 2022) is a screenplay summarization dataset comprising pairs of TV series transcripts and human-written summaries. Similar to QMSum, its plot details are scattered throughout the transcript and must be integrated to form succinct descriptions in the summary.

Table 1 presents the overall statistics of LONGEMBED. Considering the computational complexity that increases quadratically with input length, we intentionally restrict the number of candidate documents in each task to not exceed  $10^3$ . In this way, we can efficiently evaluate the basic long context capabilities of embedding models. For further elaboration on the source and examples for each dataset, please refer to Appendix C.

## 4 Methodology

### 4.1 Preliminary: APE & RoPE

**Absolute Position Embedding (APE)** stands as the predominant positional encoding strategy for embedding models, as majority of them follows the BERT architecture (Devlin et al., 2019). APE-based models first embed absolute position ids into position vectors and add token embeddings to their corresponding position vectors, before feeding them to a stack of transformer layers.

**Rotary Position Embedding (RoPE)** is the most pervasive position embedding strategy in the era of LLMs, including LLaMA (Touvron et al., 2023), QWen (Bai et al., 2023a), etc. It encodes position information of tokens with a rotation matrix that naturally incorporates explicit relative position dependency. To elucidate, given a hidden vector

$\mathbf{h} = [h_0, h_1, \dots, h_{d-1}]$  of dimension  $d$ , and a position index  $m$ , RoPE operates as follows:

$$f(\mathbf{h}, m) = [(h_0 + ih_1)e^{im\theta_0}, (h_2 + ih_3)e^{im\theta_1}, \dots, (h_{d-2} + ih_{d-1})e^{im\theta_{d/2-1}}]$$

where  $\theta_j = 10000^{-2j/d}$ ,  $j \in \{0, 1, \dots, d/2 - 1\}$ ,  $i = \sqrt{-1}$  is the imaginary unit. Unlike APE that is directly applied to the input vector  $\mathbf{x}$ , RoPE is employed on the query and key vectors at each layer. The attention score  $a(\mathbf{q}, \mathbf{k})$  between a query  $\mathbf{q}$  at position  $m$  and a key  $\mathbf{k}$  at position  $n$  is:

$$\begin{aligned} a(\mathbf{q}, \mathbf{k}) &= \text{Re}\langle f(\mathbf{q}, m), f(\mathbf{k}, n) \rangle \\ &= \text{Re} \left[ \sum_{j=0}^{d/2-1} (q_{2j} + iq_{2j+1})(k_{2j} - ik_{2j+1})e^{i(m-n)\theta_j} \right] \\ &:= g(\mathbf{q}, \mathbf{k}, (m-n)\boldsymbol{\theta}) \end{aligned} \tag{1}$$

where  $g(\cdot)$  is an abstract mapping function exclusively dependent on  $\mathbf{q}, \mathbf{k}$  and  $(m-n)\boldsymbol{\theta}$ .

### 4.2 Extending APE-based Models

As delineated in Section 2, training-free context extension strategies applicable to embedding models can be classified into 3 categories: 1) Divide-and-conquer; 2) Position reorganization; 3) Position interpolation. In this section, we introduce methods from each of these categories to assess their applicability to embedding models. Further fine-tuning on top of these methods is also included. Let  $L_o$  represent the original context length,  $\mathcal{D} = \{x_1, x_2, \dots, x_{L_t}\}$  denote a long document of target context length  $L_t$ , and  $s = \lceil L_t/L_o \rceil$  indicate the context scaling factor. The context extension methods we investigated are described below:

**Parallel Context Windows (PCW).** To process a long document with a short-context model, PCW divides the long document into multiple short chunks, processes each chunk in parallel, and aggregates their results (Ratner et al., 2023; Yen et al., 2024). In practice, we first segment  $\mathcal{D}$  into chunks of  $L_o$  tokens, then average over each chunk’s embeddings to represent  $\mathcal{D}$ . For simplicity, we set the overlap between adjacent chunks to 0, except for the last chunk, to ensure it contains  $L_o$  tokens.

**Grouped & Recurrent Positions (GP & RP).** Dividing inputs into chunks and processing them separately sacrifices their interaction in between. By contrast, position reorganization accommodates longer context by reusing the original position ids. To be specific, we experiment with two simple

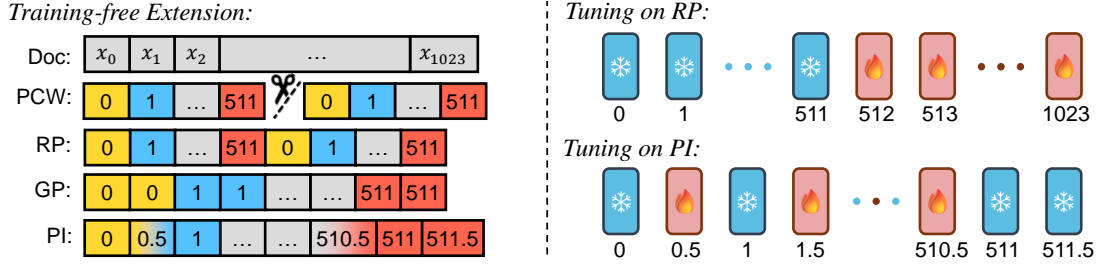


Figure 4: (Left) Arrangement of pids for extending APE-based models from 512 to 1,024. (Right) Illustration of learnable (🔥) and frozen (\*❄️) position vectors when further tuning on RP / PI.

strategies: *Grouped Positions* and *Recurrent Positions*. The former groups the original position ids as such:  $f_{gp}(pid) \rightarrow \lfloor pid/s \rfloor$ , while the latter assigns the position ids recurrently, formulated as:  $f_{rp}(pid) \rightarrow pid \bmod L_o$ .

**Linear Position Interpolation (PI).** Instead of reusing position ids, Chen et al. (2023) introduces new position embeddings via linear interpolation of existing ones. To apply PI on APE-based models, we map the positions ids as such:  $f_{pi}(pid) \rightarrow pid/s$ , and assign embeddings for non-integers as linear interpolation of that of neighboring integers. In practice, we first extend the original position embedding matrix  $E_o \in \mathbb{R}^{L_o \times d}$  into  $E_t \in \mathbb{R}^{L_t \times d}$ , where  $d$  stands for hidden size. Next, we assign  $E_t[i \cdot s] = E_o[i], i \in \{0, 1, \dots, L_o - 1\}$ . For non-integer position id  $j$  between  $i$  and  $i + 1$ , we determine their embeddings as follows:  $E_t[s \cdot j] = ((i + 1 - j)E_t[i \cdot s] + (j - i)E_t[(i + 1) \cdot s])$ .

**Further Tuning.** Except for PCW, which divides long texts into smaller blocks and processes separately, GP, RP, and PI can all be seen as extending the position embedding matrix. Since APE-based models assign an independent vector to each position, we can freeze the original model parameters while updating only the newly added position embeddings. In this way, we can strictly maintain model ability within 512 context, while harvesting further performance gains in handling long context as free lunch. Specifically, further fine-tuning on top of RP and PI is explored in this paper, as illustrated in Figure 4 (Right). Since the traditional training data for embedding models are short queries and passages not exceeding 512 tokens, we manipulate position ids to simulate long training samples, as proposed in Zhu et al. (2023). See Appendix B for details of further fine-tuning.

### 4.3 Extending RoPE-based Models

For RoPE-based models, we further explore Self Extend and NTK, which respectively advances over GP and PI, harnessing the inherent advantages of RoPE. Since there is no simple strategy for further training while exactly maintaining original performance like APE, we leave comprehensive exploration of training-based context window extension for RoPE-based models for future work.

**Self Extend (SE).** Compared with APE, RoPE operates on the query and key vectors at each layer to encode relative positions, offering enhanced flexibility for position reorganization. For each token, instead of assigning grouped relative positions to all other tokens, SelfExtend (Jin et al., 2024) re-introduces normal relative positions within the nearest neighbor window  $w$ , achieving improved performance. For example, consider a document of 10 tokens  $\{x_0, x_1, \dots, x_9\}$  with a neighbor window size  $w = 4$  and a group size  $g = 2$ . The relative positions to  $x_0$  are  $\{0, 1, 2, 3, 4, 4, 5, 5, 6, 6\}$ . For  $x_4$ , the relative positions of the other tokens are  $\{-4, -3, -2, -1, 0, 1, 2, 3, 4, 4\}$ .

**NTK-Aware Interpolation (NTK).** Given a scaling factor  $s$ , PI proportionally down-scales position index  $m$  to  $m/s$ . In this way, the attention score  $a(\mathbf{q}, \mathbf{k})$  defined in Equation 1 becomes  $g(\mathbf{q}, \mathbf{k}, (m - n)\theta/s)$ . This is also equivalent to reducing the frequencies  $\theta$  uniformly, which may prevent the model from learning high-frequency features, as shown by the Neural Tangent Kernel (NTK) theory (Jacot et al., 2018). To remedy this, NTK-Aware interpolation (Peng and Quesnelle, 2023) scales high frequencies less and low frequencies more to spread out the interpolation pressure across multiple dimensions. This is achieved by directly altering the original  $\theta_j = 10000^{-2j/d}$  into  $\theta'_j = (10000\lambda)^{-2j/d}$ , where  $\lambda$  is conventionally chosen to be slightly greater than  $s$ .

Model	Param.	CTX Len.	Synthetic (Acc@1)		Real (nDCG@10)				Avg.
			Passkey	Needle	NQA	QMS	SFD	WQA	
<i>512 Context Models</i>									
E5 <sub>Base</sub> (Wang et al., 2022)	110M	512	38.0	28.5	25.3	23.8	74.7	55.8	<b>41.0</b>
E5-RoPE <sub>Base</sub>	110M	512	38.5	31.5	24.6	23.2	66.6	58.8	40.5
GTE <sub>Base</sub> (Li et al., 2023)	110M	512	31.0	24.5	28.6	21.8	55.8	47.3	34.8
BGE <sub>Base</sub> (Xiao et al., 2023)	110M	512	18.0	25.3	25.6	22.4	60.3	51.7	33.9
Contriever (Izacard et al., 2021)	110M	512	38.5	29.0	26.7	25.5	73.5	47.3	40.1
GTR <sub>Base</sub> (Ni et al., 2022)	110M	512	38.5	26.3	26.5	18.3	63.7	52.2	36.5
$\geq 4k$ Context Models									
E5-Mistral (Wang et al., 2023b)	7B	4,096	71.0	48.3	44.6	43.6	96.8	82.0	<b>64.4</b>
Jina-V2 (Günther et al., 2023)	137M	8,192	50.3	54.5	37.9	38.9	93.5	74.0	58.2
Nomic-V1(Nussbaum et al., 2024)	137M	8,192	60.7	39.5	41.2	36.7	93.0	73.8	57.5
BGE-M3 (Chen et al., 2024)	568M	8,192	59.3	40.5	45.8	35.5	94.0	78.0	58.9
OpenAI-Ada-002	-	-	50.8	36.8	41.1	40.0	91.8	80.1	56.8
<i>Our Extended Models</i>									
E5 <sub>Base</sub> + Tuning (4k)	110M	4,096	67.3	41.5	30.4	35.7	95.2	69.2	56.6
E5-RoPE <sub>Base</sub> + SelfExtend (4k)	110M	4,096	73.5	53.5	32.3	39.1	91.9	74.6	60.8
E5-Mistral + NTK (32k)	7B	32,768	<b>93.8</b>	<b>66.8</b>	<b>49.8</b>	<b>49.2</b>	<b>97.1</b>	<b>95.2</b>	<b>75.3</b>

Table 2: Results (%) of existing and extended embedding models on LONGEMBED. *NQA*, *QMS*, *SFD*, *WQA* is short for *NarrativeQA*, *QMSum*, *SummScreenFD*, *2WikiMultihopQA*, respectively. We show that context window extension can effectively improve existing embedding models in processing long context.

## 5 Experiments

### 5.1 Experimental Setup

**Benchmarked Models.** We evaluate both open-sourced and proprietary models on LONGEMBED, including E5, GTE, BGE, Contriever, GTR, E5-Mistral, Jina-V2, Nomic-V1, BGE-M3, OpenAI-ada-002. M2 (Saad-Falcon et al., 2024) is not included in our evaluation, given its training data partly overlaps with test samples in LONGEMBED.

**Candidate Models for Extension.** From each of the APE-based and RoPE-based category, we select 2 candidate models for comprehensive study. The former includes E5<sub>Base</sub> and GTE<sub>Base</sub>. The latter includes the 4,096-context E5-Mistral, and a newly trained E5-RoPE<sub>Base</sub>, which supports 512 context (See Appendix A for its training details and BEIR results). Note that E5-RoPE<sub>Base</sub> employs the same training procedure and training data as E5<sub>Base</sub>, only with APE substituted with RoPE. This facilitates fair comparison of APE / RoPE-based models in context window extension, as presented in Section 5.4. For implementation details of each context window extension strategies on each model, please refer to Appendix B.

### 5.2 Main Results

Table 2 demonstrates the performance of existing embedding models on our LONGEMBED benchmark. Among the 512-context models, E5<sub>Base</sub> achieves the highest average score of 41.0 points, closely followed by E5-RoPE<sub>Base</sub> and Contriever. As the supported context length increases beyond 4k, exemplified by E5-Mistral and Jina-V2, a discernible increase in scores is observed. This verifies both the efficacy of these long-context models and the validity of LONGEMBED to assess long-context retrieval. Note that even the best performing model attains only 64.4 pts on average, indicating huge room for improvement in current models.

In the last row block of Table 2, we further include the best results achieved by E5<sub>Base</sub>, E5-RoPE<sub>Base</sub> and E5-Mistral after context window extension. For E5<sub>Base</sub> and E5-RoPE<sub>Base</sub>, we extend their contexts from 512 to 4,096. For E5-Mistral, we extend its context from 4,096 to 32,768. Compared to the original versions, the extended models achieve an average score increase of +15.6 / +20.3 / +10.9 points. This indicates the efficacy of these context extension strategies on embedding models, enabling them to handle inputs of several folds longer. Detailed performance comparison of different extension strategies on APE & RoPE-based embedding models is presented in Section 5.3.

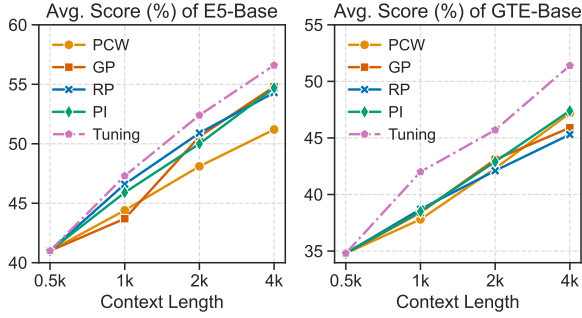


Figure 5: Effects of different context window extension methods on  $E5_{Base}$  and  $GTE_{Base}$ . We show that further tuning yields the best results.

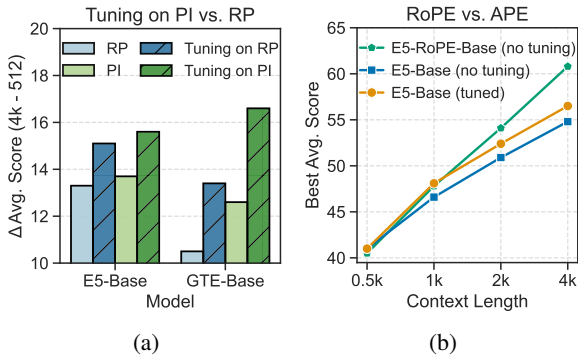


Figure 6: (a) Performance gain after tuning on PI / RP, compared with the original model. (b) Best results achieved by extended versions of  $E5_{Base}$  /  $E5-RoPE_{Base}$ .

### 5.3 Comparison of Extension Methods

**APE-based Models.** Figure 5 illustrates the impact of various context extension strategies on  $E5_{Base}$  and  $GTE_{Base}$  across different target context lengths. We observe that plug-and-play methods including GP, RP, PI and PCW strategies yield comparable results with no significant disparities. On the other hand, further tuning consistently yields additional performance gains for both models, across all target context lengths. Particularly noteworthy is  $GTE_{Base}$ , which showcases a substantial average score increase of approximately 5 points after further tuning. This suggests that freezing the original model weights and fine-tuning exclusively the added position embeddings can effectively extend the model’s context window while strictly maintaining model’s original ability.

**RoPE-based Models.** Table 3 depicts the outcomes of  $E5-RoPE_{Base}$  and  $E5-Mistral$  on each dataset of  $LONGEMBED$  after context window extension via PCW, GP, PI, SE and NTK. It is observed that RoPE-specific methods including NTK and SE yield significant improvements for both

Model	Synthetic		Real				Avg.
	P	N	NQA	QMS	SFD	WQA	
$E5-RoPE_{Base}$	38.5	31.5	24.6	23.2	66.6	58.8	40.5
+PCW (4k)	42.5	50.8	25.1	34.9	<b>94.9</b>	69.3	52.9
+GP (4k)	68.0	38.8	25.9	30.9	85.8	65.8	52.5
+PI (4k)	68.3	36.0	25.9	30.8	84.9	65.3	51.9
+SE (4k)	<b>73.5</b>	<b>53.5</b>	<b>32.3</b>	<b>39.1</b>	91.9	<b>74.6</b>	<b>60.8</b>
+NTK (4k)	66.3	46.5	25.5	35.8	90.8	71.7	56.1
$E5-Mistral$	71.0	48.3	44.6	43.6	96.8	82.0	64.4
+PCW (32k)	63.5	49.5	<b>59.3</b>	51.3	<b>97.3</b>	91.2	68.7
+GP (32k)	81.0	48.8	37.0	42.9	90.6	88.1	64.7
+PI (32k)	89.8	48.5	37.8	40.4	76.8	63.0	59.4
+SE (32k)	90.8	52	49.3	48.7	97.2	<b>96.4</b>	72.4
+NTK (32k)	<b>93.8</b>	<b>66.8</b>	49.8	<b>49.2</b>	97.1	95.2	<b>75.3</b>

Table 3: Results (%) of context window extension methods on  $E5-RoPE_{Base}$  and  $E5-Mistral$ . For datasets,  $P$ ,  $N$ ,  $NQA$ ,  $QMS$ ,  $SFD$ ,  $WQA$  is short for *Passkey*, *Needle*, *NarrativeQA*, *QMSum*, *SummScreenFD*, *2WikiMulti-hopQA*. For extension methods,  $PCW$ ,  $GP$ ,  $PI$ ,  $SE$ ,  $NTK$  are short for *Parallel Context Windows*, *Grouped Positions*, *Linear Position Interpolation*, *SelfExtend*, and *NTK-Aware Interpolation*, respectively.

models across all datasets, surpassing PCW, PI and GP by a large margin.

### 5.4 Analysis

**Tuning on PI vs. RP.** Figure 6a compares further tuning on top of RP vs. PI. In the former approach, the initial 512 position embeddings are frozen while the remaining embeddings are tuned, whereas for the latter, the frozen / learnable embedding vectors are arranged in an interleaved manner. We observe that tuning on PI consistently produces superior results on both  $GTE_{Base}$  and  $E5_{Base}$ . A possible explanation is that fixed vectors in PI serve intrinsically as anchors, preventing the learnable vectors from converging to suboptimal values.

**RoPE vs. APE.** We further discuss the potential of APE / RoPE-based models for context window extension.  $E5_{Base}$  and  $E5-RoPE_{Base}$  are selected as the comparison subjects thanks to their shared training process, training data, and comparable performance on BEIR and  $LONGEMBED$  benchmarks. At each target context length ( $\{1k, 2k, 4k\}$ ), we report the best scores achieved by each model on  $LONGEMBED$ , as illustrated in Figure 6b. Without requiring further training,  $E5-RoPE_{Base}$  consistently demonstrates superior performance compared to  $E5_{Base}$  across all target lengths. Furthermore, as the target window length increases, this



superiority becomes more pronounced, even surpassing the fine-tuned version of  $E5_{\text{Base}}$  by a large margin. This suggests that RoPE-based models can better extrapolate to longer context. Consequently, we advocate for the use of RoPE in future embedding models.

## 6 Conclusion

This paper explores context window extension of existing embedding models. Through extensive experiments on our LONGEMBED benchmark, we show that training-free context window extension strategies can effectively increase the input length of these models by several folds. Further, our analysis reveals the superiority of RoPE-based embedding models over APE-based ones in context window extension. Hence, we advocate for the use of RoPE for future embedding models.

## Limitations

As a pioneering work in applying context window extension on embedding models, this paper is still limited in several aspects, particularly in that most of the context extension strategies explored in this paper are training-free. As evidenced by previous findings (Xiong et al., 2023; Fu et al., 2024; Zhang et al., 2024b; Yen et al., 2024), and the additional performance gain achieved via tuning on  $E5_{\text{Base}}$  and  $GTE_{\text{Base}}$ , we believe further fine-tuning on top of plug-and-play methods can bring even better extension results. In the future, we will make comprehensive exploration of training-based context window extension for embedding models, especially for RoPE-based ones.

## Ethics Statement

This work fully complies with the ACL Ethics Policy. We declare that there are no ethical issues in this paper, to the best of our knowledge.

## Acknowledgement

We thank the anonymous reviewers for their helpful comments on this paper. We thank Xueguang Ma, Niklas Muennighoff, and Kenneth Enevoldsen for their thoughtful discussion and assistance in integrating LongEmbed into MTEB. This work was partially supported by National Natural Science Foundation of China (No. 62476010).

## References

- Chenxin An, Fei Huang, Jun Zhang, Shansan Gong, Xipeng Qiu, Chang Zhou, and Lingpeng Kong. 2024. Training-free long-context scaling of large language models. *arXiv preprint arXiv:2402.17463*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. 2023b. Longbench: A bilingual, multitask benchmark for long context understanding. *arXiv preprint arXiv:2308.14508*.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216*.
- Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2022. Summscreen: A dataset for abstractive screenplay summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8602–8615.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.
- David Chiang and Peter Cholak. 2022. [Overcoming a theoretical limitation of self-attention](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7654–7664, Dublin, Ireland. Association for Computational Linguistics.
- João Coelho, Bruno Martins, João Magalhães, Jamie Callan, and Chenyan Xiong. 2024. Dwell in the beginning: How language models embed long documents for dense retrieval. *arXiv preprint arXiv:2404.04163*.
- Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391–407.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Yiran Ding, Li Lina 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 preprint arXiv:2402.13753*.
- Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hananeh Hajishirzi, Yoon Kim, and Hao Peng. 2024. Data engineering for scaling language models to 128k context. *arXiv preprint arXiv:2402.10171*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910.
- Tao Ge, Jing Hu, Xun Wang, Si-Qing Chen, and Furu Wei. 2023. In-context autoencoder for context compression in a large language model. *arXiv preprint arXiv:2307.06945*.
- Michael Günther, Jackmin Ong, Isabelle Mohr, Alaeddine Abdessalem, Tanguy Abel, Mohammad Kalim Akram, Susana Guzman, Georgios Mastrapas, Saba Sturua, Bo Wang, et al. 2023. Jina embeddings 2: 8192-token general-purpose text embeddings for long documents. *arXiv preprint arXiv:2310.19923*.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. **Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Towards unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*, 2(3).
- Arthur Jacot, Franck Gabriel, and Clément Hongler. 2018. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. LlmLingua: Compressing prompts for accelerated inference of large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13358–13376.
- 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 preprint arXiv:2401.01325*.
- Greg Kamradt. 2023. Needle in a haystack - pressure testing llms. [https://github.com/gkamradt/LLMTest\\_NeedleInAHaystack](https://github.com/gkamradt/LLMTest_NeedleInAHaystack).
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781.
- Tomáš Kočický, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. **The NarrativeQA reading comprehension challenge**. *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Benjamin Lefaudeux, Francisco Massa, Diana Liskovich, Wenhan Xiong, Vittorio Caggiano, Sean Naren, Min Xu, Jieru Hu, Marta Tintore, Susan Zhang, Patrick Labatut, Daniel Haziza, Luca Wehrstedt, Jeremy Reizenstein, and Grigory Sizov. 2022. xformers: A modular and hackable transformer modelling library. <https://github.com/facebookresearch/xformers>.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Amirkeivan Mohtashami and Martin Jaggi. 2023. Landmark attention: Random-access infinite context length for transformers. *arXiv preprint arXiv:2305.16300*.
- Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, et al. 2022. Text and code embeddings by contrastive pre-training. *arXiv preprint arXiv:2201.10005*.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human-generated machine reading comprehension dataset.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, et al. 2022. Large dual encoders are generalizable retrievers. In *Proceedings*

- of the 2022 Conference on Empirical Methods in Natural Language Processing, pages 9844–9855.
- Zach Nussbaum, John X Morris, Brandon Duderstadt, and Andriy Mulyar. 2024. Nomic embed: Training a reproducible long context text embedder. *arXiv preprint arXiv:2402.01613*.
- Bowen Peng and Jeffrey Quesnelle. 2023. Ntk-aware scaled rope allows llama models to have extended (8k+) context size without any fine-tuning and minimal perplexity degradation. [https://www.reddit.com/r/LocalLLaMA/comments/14lz7j5/ntkaware\\_scaled\\_rope\\_allows\\_llama\\_models\\_to\\_have](https://www.reddit.com/r/LocalLLaMA/comments/14lz7j5/ntkaware_scaled_rope_allows_llama_models_to_have).
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2023. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. Parallel context windows for large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6383–6402.
- Anian Ruoss, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Róbert Csordás, Mehdi Bannani, Shane Legg, and Joel Veness. 2023. Randomized positional encodings boost length generalization of transformers. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1889–1903.
- Jon Saad-Falcon, Daniel Y Fu, Simran Arora, Neel Guha, and Christopher Ré. 2024. Benchmarking and building long-context retrieval models with loco and m2-bert. *arXiv preprint arXiv:2402.07440*.
- 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 Conference on Empirical Methods in Natural Language Processing*, pages 12007–12021, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jianlin Su. 2021. Understanding attention scaling from the perspective of entropy invariance. <https://spaces.ac.cn/archives/8823>.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2021. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. **BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models**. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2023a. Simlm: Pre-training with representation bottleneck for dense passage retrieval. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2244–2258.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2023b. Improving text embeddings with large language models. *arXiv preprint arXiv:2401.00368*.
- Suyuchen Wang, Ivan Kobzyev, Peng Lu, Mehdi Rezagholizadeh, and Bang Liu. 2024a. Resonance rope: Improving context length generalization of large language models. *arXiv preprint arXiv:2403.00071*.
- Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. 2024b. Augmenting language models with long-term memory. *Advances in Neural Information Processing Systems*, 36.
- Wenhao Wu, Yizhong Wang, Yao Fu, Xiang Yue, Dawei Zhu, and Sujian Li. 2024. Long context alignment with short instructions and synthesized positions. *arXiv preprint arXiv:2405.03939*.
- Chaojun Xiao, Pengl Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, Song Han, and Maosong Sun. 2024. Inllm: Unveiling the intrinsic capacity of llms for understanding extremely long sequences with training-free memory. *arXiv preprint arXiv:2402.04617*.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighof. 2023. C-pack: Packaged resources to advance general chinese embedding. *arXiv preprint arXiv:2309.07597*.
- Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajjwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, et al. 2023. Effective long-context scaling of foundation models. *arXiv preprint arXiv:2309.16039*.
- Howard Yen, Tianyu Gao, and Danqi Chen. 2024. **Long-context language modeling with parallel context encoding**. *Preprint*, arXiv:2402.16617.

- Peitian Zhang, Zheng Liu, Shitao Xiao, Ninglu Shao, Qiwei Ye, and Zhicheng Dou. 2024a. Soaring from 4k to 400k: Extending llm’s context with activation beacon. *arXiv preprint arXiv:2401.03462*.
- Yikai Zhang, Junlong Li, and Pengfei Liu. 2024b. Extending llms’ context window with 100 samples. *arXiv preprint arXiv:2401.07004*.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 2021. QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization. In *North American Association for Computational Linguistics (NAACL)*.
- 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. In *The Twelfth International Conference on Learning Representations*.

## A Training Details for E5-RoPE<sub>Base</sub>

Params	Pre-training		Fine-tuning	
	E5 <sub>Base</sub>	E5-RoPE <sub>Base</sub>	E5 <sub>Base</sub>	E5-RoPE <sub>Base</sub>
learning rate	$2 \times 10^{-4}$	$2 \times 10^{-4}$	$2 \times 10^{-5}$	$2 \times 10^{-5}$
GPUs (V100)	32	32	8	8
warmup steps	1000	1000	400	400
max length	128	512	192	192
batch size	32k	16k	256	256
max steps	20k	20k	n.a.	n.a.
epochs	n.a.	n.a.	3	3
$\tau$	0.01	0.01	0.01	0.01
$\alpha$	n.a.	n.a.	0.2	0.2
weight decay	0.01	0.01	0.01	0.01
hard negatives	0	0	7	7
pos embedding	APE	RoPE	APE	RoPE

Table 4: Hyperparameters for contrastive pre-training and fine-tuning of E5<sub>Base</sub> and E5-RoPE<sub>Base</sub>.

In this section, we describe the training details of E5-RoPE<sub>Base</sub>. Our training procedure and data exactly follows that of E5 (Wang et al., 2022), where we first perform contrastive pre-training on their collected CCPairs, then perform fine-tuning on the concatenation of 3 datasets: MS-MARCO passage ranking (Nguyen et al., 2016), NQ (Karpukhin et al., 2020; Kwiatkowski et al., 2019), and NLI (Gao et al., 2021). Each example is paired with 7 hard negatives. We leverage the mined hard negatives and re-ranker scores from SimLM (Wang et al., 2023a) for the first two datasets. As the NLI dataset only provides 1 hard negative per example, we randomly sample 6 sentences from the entire corpus. xFormers (Lefaudeux et al., 2022) is used for memory efficient training. As presented in Table 4, training hyperparameters for E5<sub>Base</sub> and E5-RoPE<sub>Base</sub> are identical, except in two aspects:

- **Initialization.** Before contrastive pre-training, E5<sub>Base</sub> is initialized on BERT<sub>Base</sub> (Devlin et al., 2019), which employs absolute position embeddings (APE). For the initialization of E5-RoPE<sub>Base</sub>, we simply replace the APE part of BERT<sub>Base</sub> with RoPE. It’s worth noting that the BERT<sub>Base</sub> model after this replacement cannot function properly. We count on the subsequent pre-training phase to adapt the model to RoPE.
- **Pre-training length and batch size.** E5<sub>Base</sub> does not update its position embedding matrix during the training phase, i.e., it utilizes the same position embedding matrix as BERT<sub>Base</sub>. This

Tasks	# W/Q.	# W/D.	E5 <sub>Base</sub>	E5-RoPE <sub>Base</sub>
MS MARCO	6.0	56.0	41.8	42.4
Trec-Covid	10.6	160.8	69.6	73.3
NFCorpus	3.3	232.3	35.4	34.9
NQ	9.2	78.9	58.2	60.1
HotpotQA	17.6	46.3	69.1	61.0
FiQA	10.8	132.3	39.8	36.4
ArguAna	193.0	166.8	44.6	54.2
Touche-2020	6.6	292.4	26.4	26.6
CQADupStack	8.6	129.1	37.4	36.5
Quora	9.5	11.4	86.6	87.7
DBPedia	5.4	49.7	42.2	40.0
Scidocs	9.4	176.2	18.7	18.1
Fever	8.1	84.8	85.0	68.0
Climate-Fever	20.1	84.8	26.6	19.0
Scifact	12.4	213.6	72.0	71.0
Average	< 200	< 300	50.23	48.61

Table 5: Statistics and performance comparison of E5<sub>Base</sub> and E5-RoPE<sub>Base</sub> on 15 publicly available BEIR tasks. # W/Q. and # W/D. stands for word number per query and per document, respectively.

allows it to generalize to input sequences of up to 512 tokens, while being trained with a max training length of 192. As for E5-RoPE, replacing APE with RoPE during initialization prevents us from directly inheriting the original model’s capability in handling 512 tokens. Consequently, in the pre-training phase of E5-RoPE, we set the maximum training length to 512, and reduce the batch size to 16k according to memory constraints.

Table 5 demonstrates results of E5<sub>Base</sub> and E5-RoPE<sub>Base</sub> on 15 publicly available BEIR tasks. We observe comparable overall scores between both models. This comparable performance, along with their shared training process and training data, facilitates fair comparison of APE and RoPE-based models’s capabilities in length extrapolation. Note that the slight performance loss of E5-RoPE<sub>Base</sub> could possibly be attributed to the replacement of position embedding in the initialization phase, or the reduced batch size in the pre-training phase, as mentioned before.

## B Implementation Details for Context Extension Strategies

This section describes implementation details for the explored context extension strategies. For plug-and-play methods including PCW, RP, GP, PI, NTK and SE, Table 6 summarizes their hyperparameters under each condition.

Extension	PCW & GP & RP & PI	NTK	SE
<i>GTE<sub>Base</sub> &amp; E5<sub>Base</sub></i>			
512 -> 1,024	$L_o = 512, L_t = 1,024, s = 2$	-	-
512 -> 2,048	$L_o = 512, L_t = 2,048, s = 4$	-	-
512 -> 4,096	$L_o = 512, L_t = 4,096, s = 8$	-	-
<i>E5-RoPE<sub>Base</sub></i>			
512 -> 1,024	$L_o = 512, L_t = 1,024, s = 2$	$\lambda = 3$ (10,000 -> 30,000)	$g = 3, w = 256$
512 -> 2,048	$L_o = 512, L_t = 2,048, s = 4$	$\lambda = 5$ (10,000 -> 50,000)	$g = 5, w = 128$
512 -> 4,096	$L_o = 512, L_t = 4,096, s = 8$	$\lambda = 10$ (10,000 -> 100,000)	$g = 9, w = 64$
<i>E5-Mistral</i>			
4,096 -> 8,192	$L_o = 4,096, L_t = 8,192, s = 2$	$\lambda = 3$ (10,000 -> 30,000)	$g = 3, w = 2,048$
4,096 -> 16,384	$L_o = 4,096, L_t = 16,384, s = 4$	$\lambda = 5$ (10,000 -> 50,000)	$g = 5, w = 1,024$
4,096 -> 32,768	$L_o = 4,096, L_t = 32,768, s = 8$	$\lambda = 10$ (10,000 -> 100,000)	$g = 9, w = 512$

Table 6: Hyperparameters for plug-and-play context extension strategies.

**Further Tuning.** On top of PI and RP, we perform further tuning on both E5<sub>Base</sub> and GTE<sub>Base</sub>, utilizing the fine-tuning dataset mentioned in Appendix A. Following the practice of PoSE (Zhu et al., 2023), we manipulate position ids to simulate long training samples. Concretely, given an input document  $\mathcal{D} = \{x_0, x_1, \dots, x_{L_o-1}\}$  of original context length  $L_o$ , we introduce a skipping bias term  $u$  at the beginning of  $\mathcal{D}$ , transferring the original position ids  $\mathcal{D}$  into  $\{0, 1, \dots, L_o - 1\}$  into  $\{u, u+1, \dots, u+L_o-1\}$ .<sup>4</sup> For every piece of training data,  $u$  is re-sampled from the discrete uniform distribution  $\mathcal{U}(\{0, 1, \dots, L_t - L_o\})$ . In this way, we ensure comprehensive coverage of target context window. The training procedure spans 3 epochs on 2 A100 GPUs, with a learning rate of  $5e^{-4}$ , a batch size of 512, and 100 steps for warmup. Other hyperparameters are same as Table 4.

**Inference.** In inference time, attention scaling (Su, 2021; Chiang and Cholak, 2022) is used by default for all tested models for better length extrapolation ability. Especially for GTE<sub>Base</sub> and E5<sub>Base</sub> tuned on PI, we use the original position ids when input length not exceeds 512. This is achieved by mapping the position ids  $\{0, 1, \dots, l\}$  into  $\{0, s, \dots, l \times s\}$ , where  $s$  is the scaling factor,  $l < 512$ .

## C Further details on LONGEMBED

Figure 7 presents source and examples for each dataset included in LONGEMBED. For QA datasets including NarrativeQA and 2WikiMultihopQA, we

<sup>4</sup>The original practice of PoSE focuses on relative position, hence introduces bias terms at the middle of document  $\mathcal{D}$ . For APE-based models, we simply skips from the beginning.

Method	Synthetic		Real				Avg.
	P	N	NQA	QMS	SFD	WQA	
BM25	<b>100</b>	<b>95.3</b>	<b>71.5</b>	<b>81.3</b>	<b>97.6</b>	<b>96.5</b>	<b>90.4</b>
E5-Mistral	71.0	48.3	44.6	43.6	96.8	82.0	64.4
+NTK (32k)	93.8	66.8	49.8	49.2	97.1	95.2	75.3

Table 7: BM25 Results on LONGEMBED.  $P, N, NQA, QMS, SFD, WQA$  is short for *Passkey, Needle, NarrativeQA, QMSum, SummScreenFD, 2WikiMultihopQA*.

adopt their test splits. Note that for 2WikiMultihopQA, we adopt the length-uniformly sampled version from Bai et al. (2023b) to better assess the model’s capabilities across various context lengths. For summarization datasets including QMSum and SummScreenFD, we adopt the version processed by SCROLLS (Shaham et al., 2022). Since SCROLLS does not include ground truth summarization in its test sets, we switch to validation set for these two datasets. Particularly for QMSum, as its validation set only have 60 documents, which is too small for document retrieval, we included the train set as well.

## D BM25 Results on LONGEMBED

Table 7 shows the scores of BM25 on LONGEMBED, along with those of the best-performing long context embedding model, E5-Mistral. The significant gap between BM25 and E5-Mistral highlights substantial room for improvement in current long context embedding models.

Dataset Name	Source / Split	Query Example	Document Example
Narrative QA	- / test	Why is Bobolink eventually eager to help Martin?	The Project Gutenberg EBook of The Purple Cloud, by M.P. Shiel\n [...] Title: The Purple Cloud\n\nAuthor: M.P. Shiel\n\nRelease Date: February 22, 2004, [...]
QMSum	Scrolls / train + valid	The team wanted to understand how they could combine different linguistic features to make a more robust recognition model. They were [...]	Project Manager: Can I close this ?\nUser Interface: Uh we don't have any changes , do we ?\nProject Manager: Oh , okay .\nUser Interface: So no . {vocalsound}\nProject Manager: {vocalsound} There we go . Okay , here we are again . Detailed design {disfmarker} oh , come on . Well {disfmarker} Ah {gap} s Forgot to insert the minutes [...]
2WikiMultihop QA	LongBench / test	Where was the director of film The Central Park Five born	Passage 1:\nMargaret, Countess of Brienne\nMarguerite d'Engbien (born 1365 - d. after 1394), was the ruling suo jure Countess of Brienne and of Conversano, suo jure Lady of Engbien, and Lady of Beauvois from 1394 until an unknown date. [...] Passage 2:\nNocher II, Count of Soissons\nNocher II (died 1019), Count of Bar-sur-Aube, Count of Soissons. He was the son of Nocher I, Count of Bar-sur-Aube. Nocher's brother Beraud (d. 1052) was Bishop of Soissons.Nocher became Count of Soissons, jure uxoris, upon his marriage to Adelise, Countess of Soissons. [...]
SummScreenFD	Scrolls / valid	Penny gets a new chair, which Sheldon enjoys until he finds out that she picked it up from the street. He constantly pesters Penny to dispose of it, to no avail. Note: Melissa Rauch is absent in this episode.	[PREVIOUSLY_ON]\nYou make jumps you can't explain, Will. The evidence explains. Then help me find some evidence. I wouldn't put him out there! Should he get too close, I need you to make sure he's not out there alone. I don't think the Shrike killed that girl in the field. This girl's killer thought that she was a pig. You think this was a copycat? I think I can help good Will, see his face. Hello? They know.\n(gunshots)\nYou said he wouldn't get too close. See?\n(gunshots)\n(knocking)\nJack: We're here!\n(police radio chatter)\nWill: Could be a permanent installation in your Evil Minds Museum. [...]
Passkey	- / -	what is the passkey for Kyree Mays?	[...] The grass is green. The sky is blue. The sun is yellow. Here we go. There and back again. The grass is green. The sky is blue.\nMalayah Graves's pass key is 41906. Remember it. 41906 is the pass key for Malayah Graves.\n\nThe sun is yellow. Here we go. There and back again. The grass is green. The sky is blue. The sun is yellow. Here we go. There and back again. [...]
Needle	- / -	What is the best thing to do in San Francisco?	Aaron Swartz created a scraped feed of the essays page. November 2021(This essay is derived from a talk at the Cambridge Union. ) [...] The best thing to do in San Francisco is eat a sandwich and sit in Dolores Park on a sunny day.\n\nThere's a narrow sense in which it refers to aesthetic judgements and a broader one in which it refers to preferences of any kind. [...]

Figure 7: Source and examples for each dataset in LONGEMBED.