From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Master's Theses, 2022-2026

Computer Science, 2023

2024

Long-Context Language Modeling with Parallel Context Encoding

Yen, Howard

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01zs25xc84q

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Long-Context Language Modeling with Parallel Context Encoding

Howard Yen

Advisor: Professor Danqi Chen

Submitted in partial fulfillment

of the requirements for the degree of

Master of Science in Engineering

Department of Computer Science

Princeton University

May 2024

I hereby declare that I am the sole author of this thesis.

I authorize Princeton University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

_____

Howard Yen

I further authorize Princeton University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

_____

Howard Yen

Abstract

Extending large language models (LLMs) to process longer inputs is crucial for numerous applications. However, the considerable computational cost of transformers, coupled with limited generalization of positional encoding, restricts the size of their context window. We introduce **C**ontext **E**xpansion with **P**arallel **E**ncoding (CEPE), a framework that can be applied to any existing decoder-only LLMs to extend their context window. CEPE adopts a small encoder to process long inputs chunk by chunk and enables the frozen decoder to leverage additional contexts via cross-attention. CEPE is efficient, generalizable, and versatile: trained with 8K-token documents, CEPE extends the context window of LLAMA-2 to 128K tokens, offering 10× the throughput with only 1/6 of the memory. CEPE yields strong performance on language modeling and in-context learning. CEPE also excels in retrieval-augmented applications, while existing long-context models degenerate with retrieved contexts. We further introduce a CEPE variant that can extend the context window of instruction-tuned models with only unlabeled data, and showcase its effectiveness on LLAMA-2-CHAT, leading to a strong instruction-following model that can leverage very long context on downstream tasks.[1]

[1] `https://github.com/princeton-nlp/CEPE`.

Acknowledgements

First, I want to thank Tianyu Gao for inspiring me to work on CEPE. This work would not have been possible with all of his help along the way, in both technical details as well as high-level direction. I also want to acknowledge Prof. Danqi Chen for her advising. I often felt lost with the high-level vision throughout the past year, and it was her guidance that kept me on track to ultimately complete this work. My two co-authors taught me so much about how to do research in the past year and I'm eternally grateful. They inspire me pursue research in long-context language modeling and natural language processing, and I'm excited for our future collaborations.

I also want to thank Prof. Karthik Narasimhan for being the second reader. In addition, Sadhika Malladi, Alexander Wettig, Mengzhou Xia, Dan Friedman, Amanda Bertsch, Adithya Bhaskar, Victoria Graf, and other members of the Princeton NLP group gave useful feedback that helpful me along the way. Finally, huge thanks to all of my friends and family for their support throughout the past year, thank you for reminding me to take breaks and have fun outside of work. In particular, I want to thank the other students in the MSE cohort, Austin, Victor, and others, who made my daily lives joyful and listened to me when I needed an ear.

To my friends and family

Contents

Abstract iii

Acknowledgements iv

List of Tables viii

List of Figures xi

1 Introduction **1**

2 Background **4**

2.1 Long-context Language Models 4

2.2 Retrieval-augmented Language Models 5

2.3 Instruction-tuned Models 6

3 Methodology: CEPE **7**

3.1 Architecture 7

3.2 Data 9

3.3 Training 10

3.4 CEPED for Instruction-Tuned Models 12

4 Experiments **13**

4.1 Long-context Language Modeling 13

4.2 Retrieval-augmented Applications 15

4.2.1 Retrieval-augmented Language Modeling 16

4.2.2 Open-domain Question Answering 17
4.3 In-Context Learning 18
4.4 Instruction-tuned Models for Long Text Understanding 19

5 Ablations **22**
5.1 Ablation Studies 22
5.1.1 Training Settings 22
5.1.2 KL Divergence 24

6 Conclusion **25**

A Appendix **27**
A.1 Training Details 27
A.1.1 Pre-training Encoder 27
A.1.2 Training CEPE 28
A.1.3 Training CEPED 29
A.2 Baseline Implementations 30
A.3 Evaluation Settings 31
A.3.1 Open-domain Question Answering 31
A.3.2 In-context Learning 32
A.3.3 ZeroScrolls 33
A.4 Ablations 34
A.4.1 Training Settings 34

List of Tables

4.1 Long-context language modeling text perplexity on ArXiv and Book from RedPajama, PG19, ProofPile, and CodeParrot. Throughput compares the speed (number of tokens/second) of each model with that of LLaMA-2. All experiments are conducted on one A100 80GB GPU, except for LLaMA-2-32K and YaRN with 128K tokens, which require model parallelism, and are conducted on four A100 GPUs. . 14

4.2 Retrieval-augmented language modeling. We report test perplexity on RedPajama across all domains. We calculate perplexity on the last 1792 tokens of the decoder input (to exclude the query tokens). k is the number of retrieved contexts used, and T is the total number of tokens. For LLaMA-2, YaRN-64K, and LLaMA-2-32K, we concatenate the contexts and prepend to the input. Avg. is the macro average across all domains. 16

4.3 ICL results averaged across 3 seeds. k is the number of demonstrations. All models use 2 demonstrations in the decoder, and we add $+k'$ demonstrations to the encoder for CEPE. † denotes the oracle setting with $k = 40$ demonstration in the decoder. 19

4.4 ZeroScrolls validation results. The total number of tokens includes both the input and generated tokens. For NarrativeQA (NQA) and Qspr (QASPER), we report the F1 scores. For QALT (QuALITY), we report accuracy. For GovReport (GvRp), SummScreenFD (SSFD), and QMSum, we report the ROUGE-L scores. CEPED uses 2K tokens in the decoder, and additional tokens are inputted through the encoder. 20

5.1 Test perplexity in long-context and retrieval-augmented language modeling, averaged over all datasets. Full results are in **??** A.6 and **??** A.7. 22

A.1 Performance of StreamingLLM with different stride lengths. . . . 31

A.2 Open-domain QA results. We report exact match scores for the Natural Questions(NQ) test set, TriviaQA(TQA) validation set, and PopQA test set. All models use two-shot in-context learning. k is the number of retrieved passages, and CEPE uses the top 10 passages in the decoder and all passages in the encoder. 32

A.3 Metrics used for each dataset. For normalization, we divide the log-likelihood by the length of the prompt. For calibration, we use Domain Conditional PMI [22] with the empty string as the domain string for all datasets for simplicity. 33

A.4 ZeroScrolls generation example on the Qasper dataset. CEPE sees the entire article through the decoder and the encoder, whereas Llama-2-Chat only sees a 2K token window. For brevity, we omit part of the decoder input and only show 2 out of $k = 17$ encoder inputs for CEPE. 35

A.5 ZeroScrolls generation example on the SummScreenFD dataset. CEPE sees the entire TV script through the decoder and the encoder, whereas LLaMA-2-Chat only sees a 2K token window. For brevity, we omit part of the decoder input and only show 2 out of $k = 44$ encoder inputs for CEPE. 36

A.6 Test perplexity for all ablation settings in the long-context language modeling evaluation setting. 37

A.7 Test perplexity on RedPajama across all domains in the retrieval-augmented setting for all ablation experiments. k is the number of additional contexts used. Avg. is the macro average across all domains. 38

A.8 ZeroScrolls results using different losses during training, where c_{KL} is the coefficient of the KL Divergence loss. $c_{CE} = 1$ is the coefficient of the Cross-Entropy loss for all experiments. 39

List of Figures

1.1 A comparison between CEPE and other techniques of extending LLMs' context window, including YaRN [48], StreamingLLM [71], and RePlug [57]. CEPE trained on 8K tokens can generalize to 128K tokens with minimal computational and memory costs. Fine-tuned methods, such as YaRN trained on 64K sequence lengths, explodes in perplexity past the training length, while inference methods such as StreamingLLM and RePlug incurs significant computational costs with higher perplexity. 2

3.1 The CEPE architecture. The encoder model encodes the additional 3 chunks ($k = 3$) of context $\mathcal{C}$ in parallel, and the final hidden representations from the encoder model are concatenated and used as inputs to the cross-attention layers in the decoder model. The cross-attention layers attend to the encoder representations between the self-attention and feed-forward layers in the decoder model. 8

4.1 Open-domain QA results. We report the exact match (EM) scores. LLaMA-2 is limited to 20 passages, and RePlug is limited to 30 passages due to memory constraints. For the complete results, refer to Table A.2. 17

Chapter 1

Introduction

Long and extensible context is crucial for large language models (LLMs) to effectively perform complex tasks, such as summarizing a book or answering questions with hundreds of retrieved web pages. However, several challenges limit the ability of LLMs to leverage long context: (1) LLMs and popular positional encodings [51, 59] do not generalize to sequence lengths longer than the lengths seen during training [49], even after additional fine-tuning [9, 10, 48, 61]. (2) Transformers [65]—the predominant architecture of LLMs—incur a quadratic computational cost and a linear memory cost with respect to the input length, making it prohibitive to use for long sequences. (3) High-quality long-context data, such as long instruction-following data, are scarce and difficult to obtain [68, 72].

Besides directly fine-tuning LLMs on longer inputs, a series of inference-time modification methods have been proposed recently to scale up the effective context window, either by modifying the attention mechanism [4, 71, 24] or encoding chunks of context in separate forward passes [57, 52, 38]. While these methods generalize to longer sequences, the model often fails to effectively leverage the extra tokens and can incur greater inference costs.

In this work, we propose an efficient and lightweight solution to extending the con-

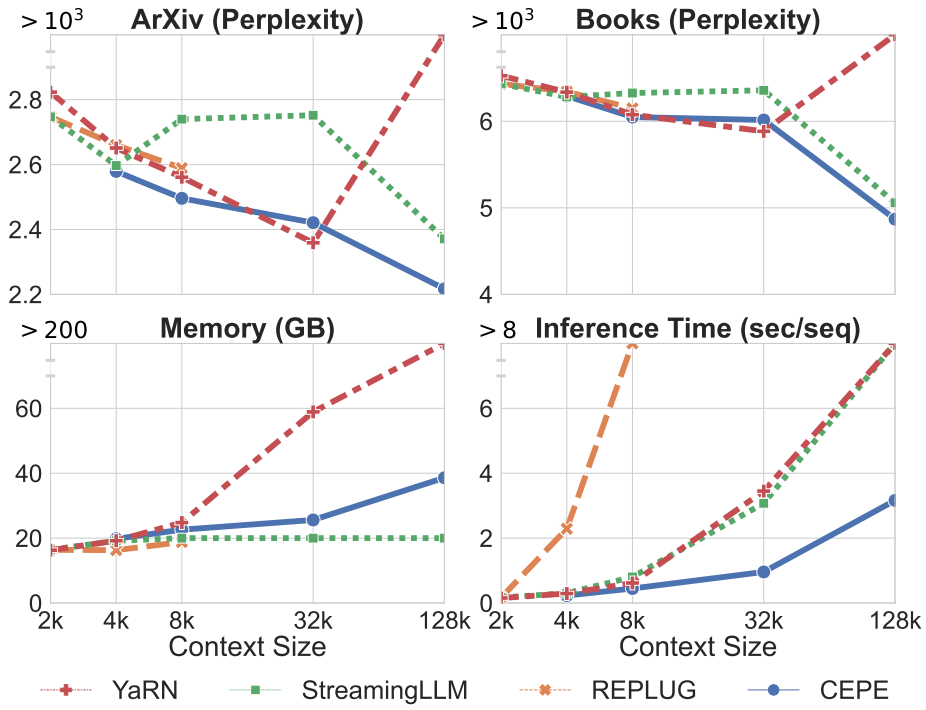


Figure 1.1: A comparison between CEPE and other techniques of extending LLMs' context window, including YaRN [48], StreamingLLM [71], and REPLUG [57]. CEPE trained on 8K tokens can generalize to 128K tokens with minimal computational and memory costs. Fine-tuned methods, such as YaRN trained on 64K sequence lengths, explodes in perplexity past the training length, while inference methods such as StreamingLLM and REPLUG incurs significant computational costs with higher perplexity.

text window of LLMs, called **C**ontext **E**xpansion with **P**arallel **E**ncoding (CEPE). CEPE is applicable to any pre-trained decoder-only LM by adding two components: (1) a small encoder that encodes the long context in chunks, and (2) a cross-attention module that is inserted at each layer of the decoder to attend to the encoder representations (Figure 3.1). With a careful selection of unlabeled training data, CEPE can leverage not only long-context documents, but also retrieved documents effectively.

CEPE offers several benefits: (1) **Length generalization**: CEPE is not limited by positional encoding constraints as the long context is encoded in segments, each with its own positional encoding. (2) **Efficiency**: using a small encoder and processing context in parallel reduce computational cost. Since cross-attention attends

only to the last layer's representations from the encoder, CEPE requires much less memory compared to decoder-only LMs, which cache the key-value pairs of every token in every layer. (3) **Reduced training cost**: unlike full fine-tuning approaches, we only tune the encoder and the cross-attention while keeping the large decoder LM frozen; augmenting a 7B decoder with a 400M encoder and cross-attention layers (1.4B parameters) can be done with one A100 80GB GPU.

We apply CEPE to LLaMA-2 and train it on a filtered version of RedPajama [62] for 20B tokens—only 1% of the pre-training budget of LLaMA-2. We first show that CEPE-LLaMA-2, trained with 8K input length, continues to improve perplexity on longer input up to 128K tokens. Then, we apply CEPE to a retrieval-augmented setting, as our larger context window allows incorporating more retrieved documents. Compared to existing methods, CEPE achieves better performance on both retrieval-augmented language modeling and open-domain question answering. Additionally, we demonstrate that CEPE can effectively leverage more demonstrations for in-context learning [6]. All the above is achieved with a much lower memory and computational cost than most previous solutions.

Finally, we propose CEPE-**D**istilled (CEPED), which extends the context window of *instruction-tuned* models, using only unlabeled data. CEPED distills the behavior of the original instruction-tuned model to the new architecture through an auxiliary KL divergence loss, which eliminates the need to curate expensive long-context instruction-following data [25]. We apply CEPED to LLaMA-2-Chat [64] and show that while preserving their instruction understanding ability, CEPED-LLaMA-2-Chat can incorporate more context and improve performance on long-text understanding tasks [54].

To conclude, CEPE is a lightweight framework that can extend context windows of any pre-trained or instruction-tuned LMs. We hope CEPE can empower future LLM research with cheap and effective long-context abilities.

Chapter 2

Background

2.1 Long-context Language Models

Many recent works on long-context language models aim to solve the problem of positional embedding extrapolation in transformers [9, 48]. Others fine-tune LMs on longer sequences [10, 61, 72] or compress the context into shorter forms [11, 12, 73]. Notably, several recent papers propose to extend the context window of LMs by modifying the attention mechanism: Xiao et al. [71] discover the use of "sink tokens" in sliding windows and reduces the number of tokens the model attends by only keeping a limited number of sink and recent tokens; Bertsch et al. [4] retrieve relevant tokens from a cache instead of attending to all tokens. This results in memory-efficient long-context LMs, but they have diminishing returns with longer contexts, as the same positional embedding may be seen multiple times and they can not fully utilize all tokens. The key advantage of CEPE is that it does not degrade for inputs longer than the training length while achieving great computational and memory efficiency compared to full fine-tuning approaches.

Novel architectures and pre-training methods, such as S4 [18], RPT [53], and Mamba [17], also extend the context window at greater efficiency. However, pre-

training is extremely expensive at scale and these methods cannot leverage existing powerful pre-trained LLMs. It is also unclear if state-space models can achieve comparable performance with transformers [29].

Lastly, existing long-context language models often degrade in performance in long-context settings due to the lack of robustness to noise [39]. In particular, they over-attend to the texts at the start and end of the inputs while ignoring those in between. CEPE avoids this by encoding each chunk in the chunk independently, resulting in a permutation-invariant decoding process.

2.2 Retrieval-augmented Language Models

Augmenting LMs with retrieval has been useful in many applications, such as open-domain question answering. It's often applied to question answering tasks, where retrieval can enrich LMs with external knowledge [36, 31]. Recently, combining LMs with retrieval systems for more generalized purposes, such as language modeling, has been explored: some researchers pre-train LMs with retrieval [19, 5, 28, 45], and others use logits interpolation from separate forward passes to incorporate retrieval information [57, 38].

Our architecture is similar to Atlas [28], RETRO, and RETRO-Fitting [5] , which encode retrieved passages independently and uses cross-attention to incorporate them into the decoder. , but they use retrieval-augmented data for pre-training, which can be expensive to acquire at a large scale. CEPE only requires fine-tuning on long document data, which are much more efficient to obtain; CEPE is also applicable to any decoder-only LM, allowing us to extend context windows for pre-existing strong models. RETRO-Fitting [5] also fine-tunes pre-trained models but requires retrieval-augmented data. Finally, these works do not consider long-context language modeling settings in addition to retrieval-augmented settings.

2.3 Instruction-tuned Models

Instruction-following models are prized for their generalization capabilities [46]. However, it can be difficult to obtain instruction-following data, especially for long-context models. Previous work use self-instruct [68] and a blend of existing data (e.g., QA datasets) to generate long-context instruction-following data [72, 38, 67]. This type of data can be synthetic, or unnatural, especially in the long-context setting. On the other hand, naturally-occurring long texts, such as books and code bases, do not contain instructions, limiting the context length of instruction-following models. While it is possible to concatenate short instruction-following training data together, it is unclear if this result in strong long-context models capable of modeling long-range dependencies. Therefore, we propose CEPED, a self-supervised method that can be easily applied to instruction-following models without additional data collection.

Chapter 3

Methodology: CEPE

We design CEPE to adapt pre-trained LLMs to perform long-context language modeling on sequences with many tokens (e.g., books). For retrieval augmentation, these long contexts may contain a set of retrieved passages instead. We first describe how CEPE modifies the architecture of LLMs to attend to representations encoded by a small encoder, and then describe how the CEPE modules are trained. Finally, we extend CEPE to CEPED, which extends the context window of instruction-tuned models using only unlabeled data.

3.1 Architecture

CEPE augments off-the-shelf decoder-only LMs by (1) adding a small, bidirectional pre-trained encoder model and (2) inserting cross-attention layers between the self-attention and feed-forward layers in every transformer block of the decoder model.

Notation. Given an input context with T tokens $x_1, ..., x_T$, we consider the first m tokens $x_1, ..., x_m$ as the additional context $\mathcal{C}$ and the last $n = T - m$ tokens $x_{m+1}, ..., x_T$ as the main input X. The additional context is split into chunks $\mathcal{C} = C_1, ..., C_k$, which can contain either segments of long documents or a set of retrieved passages. We use

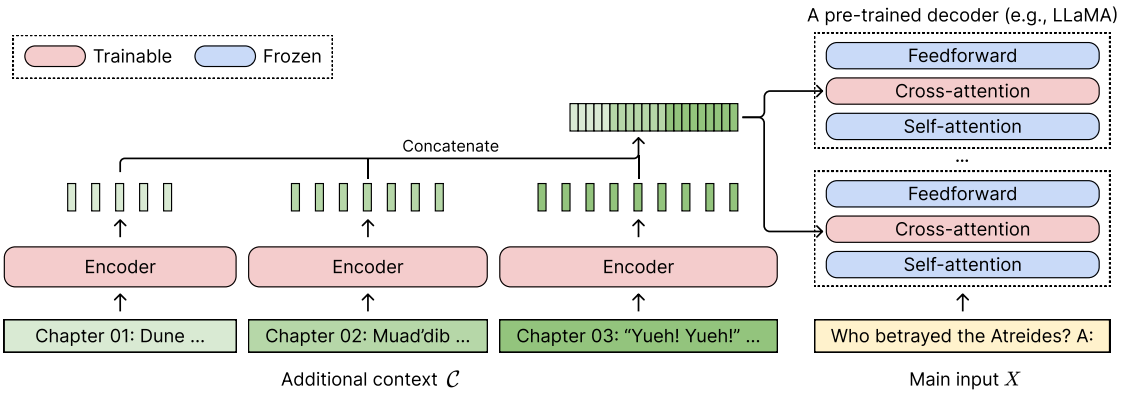


Figure 3.1: The CEPE architecture. The encoder model encodes the additional 3 chunks ($k = 3$) of context $\mathcal{C}$ in parallel, and the final hidden representations from the encoder model are concatenated and used as inputs to the cross-attention layers in the decoder model. The cross-attention layers attend to the encoder representations between the self-attention and feed-forward layers in the decoder model.

$\mathcal{M}_{\text{enc}}$ to denote the encoder model with hidden dimension d_{enc} and $\mathcal{M}_{\text{dec}}$ to denote the decoder-only LLM with hidden dimension d_{dec}.

Encoding chunks. We first encode $C_1, ..., C_k$ chunk by chunk using the trainable encoder $\mathcal{M}_{\text{enc}}$:

$$\phi_i = \mathcal{M}_{\text{enc}}(C_i) \qquad \Phi = \text{CONCAT}(\{\phi_i\}_{i=1}^{k})$$

where $\phi_i \in \mathbb{R}^{|C_i| \times d_{\text{dec}}}$ is the token-wise last layer hidden state from $\mathcal{M}_{\text{enc}}$ and $\Phi \in \mathbb{R}^{m \times d_{\text{dec}}}$. Note that $\mathcal{M}_{\text{enc}}$ is bidirectional; leveraging a bidirectional encoder results in more information-rich representations compared to unidirectional ones.

Cross-attention modules. In every decoder layer of the transformer [65], we insert a cross-attention module between the self-attention and feed-forward layers. To construct the cross-attention module, we provide Φ as keys and values, and the hidden states of X as queries. Note that in order for $\mathcal{M}_{\text{dec}}$ to attend to Φ, the key and value projection matrices in the cross-attention module also serve as an up-projection that transforms the d_{enc}-dimensional Φ into a d_{dec}-dimensional embedding. Figure 3.1 illustrates the architecture of CEPE.

Efficiency. $\mathcal{M}_{\text{enc}}$ is much smaller and encodes contexts in parallel to avoid the quadratic complexity of full attention. This enables CEPE to exhibit a substantially higher training and inference speed than if we were to use $\mathcal{M}_{\text{dec}}$ to process all $T = m+n$ tokens. Moreover, CEPE drastically reduces memory consumption by avoiding caching $(m+n)L$ key-value pairs (L is the number of layers of $\mathcal{M}_{\text{dec}}$) and instead caching only Φ and nL key-value pairs. In particular, using a standard decoder-only model requires $\mathcal{O}((m+n)Ld_{\text{dec}})$ memory whereas CEPE requires $\mathcal{O}(md_{\text{enc}}+nLd_{\text{dec}})$. In our setting, $m \gg n$ and $d_{\text{dec}} \gg d_{\text{enc}}$, so in practice, we observe a substantial gap: for each additional token, CEPE requires only 1/256 of the memory compared to encoding them in $\mathcal{M}_{\text{dec}}$.

Comparison with existing retrieval-augmented LMs. CEPE takes inspiration from retrieval-augmented models such as FiD [27], Atlas [28], and RETRO [5] for attending to parallel encodings. Our method differs in the following key aspects: FiD, Atlas, and RETRO are all full-parameter training methods for encoder-decoder models; Atlas and RETRO are pre-training methods that rely on expensive large-scale retrieval-augmented data; FiD is a fine-tuning method for open-domain QA tasks, which requires task-specific data. In contrast, CEPE is a lightweight framework that can extend the context window of any existing decoder-only model; we only fine-tune an added small encoder and the cross-attention layers; instead of retrieval-augmented data, CEPE only requires a small amount of filtered unlabeled data, as described in the following section; because CEPE is built on top of powerful pre-trained models, it can be easily adapted to a wide range of tasks without explicit fine-tuning.

3.2 Data

We use RedPajama (RP) [62] as our training corpus, which is an open-source reproduction of the LLAMA-1 [63] training data. It contains about 1 trillion tokens from

seven domains: ArXiv, Books, C4-RP, CC, Github, StackExchange, and Wiki. We first partition the corpus into three sets of documents: training, test, and retrieval; we leverage the retrieval corpus for retrieval-augmented language modeling.

We design a data mixture with strong long-range dependencies and a diversity of domains. The training data is preprocessed into two subsets, each of which is a collection of 8192-token sequences: (1) In $\text{RP}_{\text{train-cat}}$, we concatenate documents together to form training sequences; (2) In $\text{RP}_{\text{train-filter}}$, we keep documents from the Arxiv and Books domains that have at least $8,192$ tokens and sample sequences within document boundaries.

Our qualitative analysis found that data from the ArXiv and Books domains naturally contain long documents that are especially useful when training long-context models. It is also important to use a mixture of data from all domains to ensure better generalization. Thus, we only keep the filtered documents from the ArXiv and Books domains, and we use a mixture ratio of 2:1 between $\text{RP}_{\text{train-filter}}$ and $\text{RP}_{\text{train-cat}}$ for training.[1] We alternatively consider using retrieval supervision when generating the training data to improve the model's ability to leverage retrieved documents. However, our ablations found that our unlabeled data mixture transfers well to retrieval-augmented settings at a much cheaper cost (§5.1.1).

3.3 Training

We use LLAMA-2-7B [64] as $\mathcal{M}_{\text{dec}}$ (originally trained on 4K length), and insert the new cross-attention layers described in Section 3.1 into it. We add a bidirectional encoder $\mathcal{M}_{\text{enc}}$ with 435M parameters, yielding 1.8B additional parameters in the CEPE model.

[1] Due to computational constraint, we keep this ratio fixed. We leave explorations on different data mixtures for future works.

Encoder. We first pre-train a bidirectional masked-language model (MLM) on the RedPajama dataset. [2] $\mathcal{M}_{\text{enc}}$ follows the configuration of RoBERTa-large [41] but shares the vocabulary with $\mathcal{M}_{\text{dec}}$, LLAMA-2, which ultimately allows for a shared vocabulary between the encoder and decoder and enable recalling and copying texts from the encoder input during generation. Future works applying CEPE to LLAMA-2 models can simply use our pre-trained encoder. We train $\mathcal{M}_{\text{enc}}$ for 100K steps with a batch size of $2,048$ and a sequence length of 512 tokens. For more details, please refer to §A.1.1.

Cross-attention. We freeze the original weights of the decoder model and train only the added cross-attention layers as well as fine-tuning $\mathcal{M}_{\text{enc}}$ using the cross-entropy loss. We initialize the weights of the cross-attention layers to be the same as the self-attention layers except for the output matrix, which is initialized with all 0s. We first adopt a warmup training stage designed to teach $\mathcal{M}_{\text{dec}}$ to copy from $\mathcal{M}_{\text{enc}}$ through cross-attention—for each position $i \leq T$, the objective is to generate x_{i+1}, conditioned on $\mathcal{M}_{\text{enc}}(x_1, \ldots, x_T)$ and $\mathcal{M}_{\text{dec}}(x_1, \ldots, x_i)$—i.e., $\mathcal{M}_{\text{enc}}$ and $\mathcal{M}_{\text{dec}}$ share the same input. The warmup stage uses 131M tokens sampled from the training set of RP.[3]

After the warmup stage, we move to the standard training, where each sequence has $T = 8,192$ tokens. We use the last $4,096$ tokens as the decoder input X, and chunk the first $4,096$ tokens into $k = 16$ contexts of $|C_i| = 256$ tokens each as the encoder input $\mathcal{C}$, and train for 20B tokens. Freezing the decoder allows CEPE to be trained on a single A100 GPU, which is a significant reduction in computational cost compared to training $\mathcal{M}_{\text{dec}}$ with sequence length T. In practice, we use eight A100 GPUs to

[2] We pre-train our own encoder because we want the cross-attention module to be able to leverage inductive bias from the encoder. However, there are no pre-trained bidirectional encoder with LLAMA-2 vocabulary. Future works applying CEPE to LLAMA-2 models can simply use the MLM that we release.

[3] Preliminary experiments suggest this stage can stabilize training; ablations can be found in §5.1.1.

train the model for 20B tokens. For more training details and hyperparameters, please refer to §A.1.2.

3.4 CEPED for Instruction-Tuned Models

We extend our method to CEPE-DISTILLED (CEPED) to augment instruction-tuned models with longer context. Instruction-tuned models [46, 60, 64] excel in many downstream applications, but their limited context window restricts their performance in tasks that require long documents [54] or a large number of retrieved passages [16]. It is challenging to extend these models to longer context windows directly through fine-tuning due to the scarcity of high-quality instruction data.

To this end, we propose CEPED, which uses an auxiliary distillation loss to encourage $\mathcal{M}_{\text{enc}}$ and the cross-attention layers to learn the capabilities of the already instruction-tuned $\mathcal{M}_{\text{dec}}$. This can be especially useful for settings where the fine-tuning data are not open-sourced, which is the case for LLAMA-2-CHAT [64].

Distillation loss. We design a distillation objective, where the original $\mathcal{M}_{\text{dec}}$ acts as the "teacher" and the CEPED model acts as the "student". We use the input context $\text{CONCAT}(\mathcal{C}, X)$ with $m + n = 4,096$ tokens, which can fit in the context window of LLAMA-2-CHAT, our choice of $\mathcal{M}_{\text{dec}}$. First, we input $\text{CONCAT}(\mathcal{C}, X)$ to $\mathcal{M}_{\text{dec}}$ and save the logits of X as the teacher logits. During training, $\mathcal{C}$ and X are used as inputs to $\mathcal{M}_{\text{enc}}$ and $\mathcal{M}_{\text{dec}}$, respectively, and we minimize the KL divergence between the output logits of X and the corresponding teacher logits as well as the cross-entropy loss. We train the model for 10B tokens from our data mixture. For more details, see §A.1.3.

Chapter 4

Experiments

4.1 Long-context Language Modeling

We start by evaluating CEPE on long-context language modeling benchmarks to assess basic LM abilities, and compare with existing approaches in terms of perplexity, memory, and throughput.

Datasets. We evaluate on ArXiv and Books from our RedPajama test split, as well as three long-context datasets: PG19 [50], Proof-Pile [2], and CodeParrot [69]. We filter all documents to have at least $32,768$ tokens, and sample $5,000$ sequences for each dataset. We calculate the perplexity on the last 256 tokens of each sequence. Following [48], for the experiments with 128K tokens, we filter documents to have at least $131,072$ tokens and evaluate on only 10 sequences.[1]

Models. We focus on 7B parameters model. Our baseline includes LLaMA-2 and its long-sequence fine-tuned versions: LLaMA-2-32K [61], YaRN-64K, and YaRN-128K [48]. LLaMA-2-32K was trained on RP with upsampled ArXiv and Books domains, and a mixture of other data , and YaRN was trained on PG19. While

[1]This is due to computational cost and scarcity of long documents. All 10 sequences are from different documents.

	ArXiv	Book	PG19	ProofPile	CodeParrot	Throughput	Mem. (GB)
Total Tokens $= 4,096$							
LLaMA-2	2.597	**6.282**	7.614	2.409	**1.735**	1.00×	**19.2**
LLaMA-2-32K	2.601	6.621	7.945	2.414	1.785	1.00×	**19.2**
YaRN-64K	2.651	6.337	**7.326**	2.457	1.764	1.04×	**19.2**
CEPE	**2.579**	6.292	7.536	**2.396**	1.763	**1.31**×	19.8
Total Tokens $= 8,192$							
LLaMA-2	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	-	-
LLaMA-2-32K	2.505	6.339	7.744	2.221	1.729	1.00×	24.9
YaRN-64K	2.561	6.077	**7.146**	2.267	**1.714**	2.52×	24.8
RePlug	2.589	6.149	7.554	2.307	1.728	0.17×	**18.8**
StreamingLLM	2.740	6.327	7.783	2.437	1.806	1.94×	20.0
CEPE	**2.496**	**6.049**	7.372	**2.219**	1.715	**3.48**×	22.6
Total Tokens $= 32,768$							
LLaMA-2-32K	**2.322**	6.178	7.420	**2.158**	1.664	1.00×	59.1
YaRN-64K	2.359	**5.884**	**6.809**	2.193	**1.640**	1.03×	58.9
StreamingLLM	2.752	6.358	7.627	2.503	1.853	1.16×	**20.0**
CEPE	2.421	6.015	7.204	2.218	1.702	**3.72**×	25.6
Total Tokens $= 131,072$							
LLaMA-2-32K	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	-	-
YaRN-64K	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	$> 10^3$	-	-
YaRN-128K	2.359	5.270	6.306	2.242	**1.264**	1.00×	235.6
StreamingLLM	2.371	5.058	6.681	2.270	1.280	2.56×	**20.0**
CEPE	**2.217**	**4.869**	**6.305**	**2.099**	1.266	**9.90**×	38.6

Table 4.1: Long-context language modeling text perplexity on ArXiv and Book from RedPajama, PG19, ProofPile, and CodeParrot. Throughput compares the speed (number of tokens/second) of each model with that of LLaMA-2. All experiments are conducted on one A100 80GB GPU, except for LLaMA-2-32K and YaRN with 128K tokens, which require model parallelism, and are conducted on four A100 GPUs.

YaRN-64K and YaRN-128K were fine-tuned for only 1.7B and 3.4B tokens, respectively, fine-tuning all parameters requires much more memory relative to CEPE. We also evaluate on training-free long-context methods: StreamingLLM [71] and RePlug [56] with LLaMA-2-7B. Although RePlug was originally evaluated in a retrieval-augmented setting, we found that it also works well in long-context modeling, when viewing the long context as retrieved context. See §A.2 for more implementation details.

For CEPE, we put 2K tokens in the decoder when $T = 4,096$, and 4K tokens in the decoder in other settings. Additional tokens are split into chunks of 256 tokens and fed into the encoder.

Results. We show the results in Table 4.1. Compared to the two fully fine-tuned models, LLaMA-2-32K and YaRN-64K, CEPE achieves either lower or comparable perplexity across all datasets with lower memory usage and higher throughput. Furthermore, CEPE continues to improve on perplexity while maintaining low memory use at 128K tokens, well beyond its training lengths (8K); on the other hand, LLaMA-2-32K and YaRN-64K cannot generalize beyond its training length and the memory cost increases significantly. At 128K tokens, we outperform or achieve comparable perplexity with all applicable baselines on all domains.

We also outperform REPLUG across all domains while achieving much higher throughput—REPLUG encodes additional context in chunks but each chunk requires a forward pass of the main input, incurring slow speed, due to which we omit REPLUG at longer length. STREAMINGLLM maintains a low memory usage and a reasonable throughput, but the perplexity does not always decrease as the sequence length increases, likely due to the model only using a limited number of cached key-value pairs. STREAMINGLLM also suffers from a slow throughput in the case of long input sequences, as it processes only one token at a time beyond the original context window, and therefore requires a large number of forward pass to build the key-value cache. Compared to STREAMINGLLM, CEPE can leverage all input tokens and achieve better perplexity with better throughput.

4.2 Retrieval-augmented Applications

Retrieval-augmented settings naturally benefit from long-context LMs, as models can leverage the additional context to include more retrieval results. Thus, we test if CEPE trained on long-context data can transfer to retrieval-augmented settings.

	ArXiv	Book	C4-RP	CC	Github	StackEx	Wiki	Avg.
$k = 0$ ($T = 2,048$)								
LLaMA-2	**3.541**	**6.524**	**6.916**	**5.564**	**1.865**	**4.043**	**4.816**	**4.753**
LLaMA-2-32K	3.561	6.892	7.798	5.931	1.932	4.262	4.958	5.048
YaRN-64K	3.633	6.631	7.164	5.701	1.930	4.164	4.837	4.866
$k = 8$ ($T = 4,096$)								
LLaMA-2	3.602	6.581	6.963	5.348	1.829	4.044	4.815	4.740
LLaMA-2-32K	3.642	6.985	7.767	5.645	1.893	4.270	4.988	5.027
YaRN-64K	3.752	6.718	7.218	5.466	1.894	4.178	4.847	4.868
RePlug	3.535	6.494	6.895	5.395	1.833	4.029	4.798	4.711
CEPE	**3.486**	**6.481**	**6.884**	**5.319**	**1.793**	**3.709**	**4.302**	**4.568**
$k = 20$ ($T = 7,168$)								
RePlug	3.531	6.490	6.894	5.386	1.830	4.028	4.795	4.708
CEPE	**3.475**	**6.463**	**6.875**	**5.266**	**1.782**	**3.703**	**4.296**	**4.551**
$k = 50$ ($T = 14,848$)								
RePlug	3.530	6.491	6.899	5.392	1.830	4.028	4.794	4.709
CEPE	**3.467**	**6.457**	**6.881**	**5.273**	**1.777**	**3.701**	**4.292**	**4.550**

Table 4.2: Retrieval-augmented language modeling. We report test perplexity on RedPajama across all domains. We calculate perplexity on the last 1792 tokens of the decoder input (to exclude the query tokens). k is the number of retrieved contexts used, and T is the total number of tokens. For LLaMA-2, YaRN-64K, and LLaMA-2-32K, we concatenate the contexts and prepend to the input. Avg. is the macro average across all domains.

4.2.1 Retrieval-augmented Language Modeling

Datasets. We use the test and retrieval split of RedPajama described in §3.2 for retrieval-augmented LM evaluation. Each sequence contains $2,048$ tokens, and the first 256 tokens are used as the query to retrieve passages from the retrieval split. The retrieval corpus contains 200M documents of 256 tokens each, and we use Contriever [26] to retrieve k passages for each sequence.

Models. We evaluate full-context baselines, LLaMA-2, LLaMA-2-32K, and YaRN-64K, by simply prepending the retrieved passages to the input sequence. We also evaluate RePlug, which runs one forward pass for each retrieved passage and aggregates the results. CEPE uses $2,048$ tokens in the decoder and retrieved passages

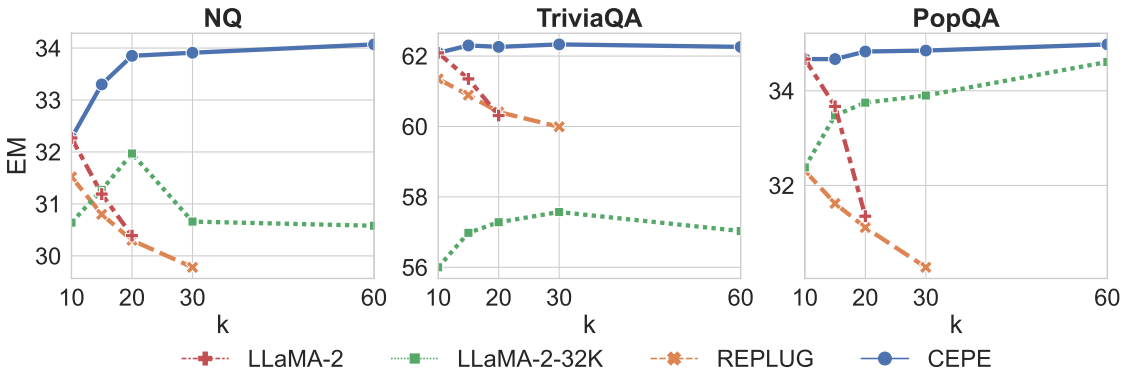


Figure 4.1: Open-domain QA results. We report the exact match (EM) scores. LLaMA-2 is limited to 20 passages, and REPLUG is limited to 30 passages due to memory constraints. For the complete results, refer to Table A.2.

are fed through the encoder in parallel.

Results. The results are shown in Table 4.2. CEPE can effectively improve perplexity by using the retrieved contexts, outperforming REPLUG. Notably, CEPE continues to improve perplexity even with $k = 50$ (trained with $k = 16$). CEPE transfers well to the retrieval-augmented setting whereas the full-context decoder models degrade.

4.2.2 Open-domain Question Answering

Given a question and a large corpus of documents, open-domain question answering (QA) requires the model to retrieve relevant passages and generate the answer. A model that can leverage a large number of retrieved passages without being distracted by irrelevant ones is desirable for this task.

Datasets. We adopt three open-domain QA datasets: Natural Questions (NQ) [34, 37], TriviaQA [30], and PopQA [44]. For each question, we use Contriever to retrieve k passages from Wikipedia.[2]

[2]Snapshot from 2018-12-20, and each passage is 100 words [31].

Models. We compare CEPE with LLaMA-2, LLaMA-2-32K, and REPLUG. For each model, we use two in-context demonstrations. For CEPE, we use 10 passages in the decoder, and all other passages are encoded separately by the encoder. Refer to §A.3.1 for more details.

Results. The results are shown in Figure 4.1 CEPE consistently outperforms all models across all datasets and k. Notably, CEPE outperforms LLaMA-2-32K on NQ and TriviaQA by over 3 and 4 points in exact match, respectively. Furthermore, CEPE does not degrade in performance as the k increases, while other models often perform worse at larger k, as they are sensitive to the large amount of redundant or irrelevant passages.

4.3 In-Context Learning

In-context learning (ICL)[6] is one of the most important emerging qualities of LLMs. In this experiment, we examine whether CEPE can effectively utilize demonstrations from the encoder context to improve performance. Task performance is usually further improved with longer input length and hence more demonstrations. Note that ICL relies on both the interactions between the demonstrations and the query, and the interactions among the demonstrations themselves; hence our architecture has a deficiency as the encoder context does not attend to each other. This experiment is solely to examine whether CEPE can effectively utilize the demonstrations from the encoder context and improve the performance compared to only using the decoder demonstrations, similar to leveraging retrieved passages; we also want to test whether more demonstrations from the encoder side can further boost the results. Specifically, we use a range of classification tasks that contains a large number (up to 150) of categories [52], where the model can benefit from additional demonstrations. Following [52], we use a test set size of 250 examples for each dataset.

	k	SST2	MR	AGNews	SST5	TREC	TREC-F	DBPedia	NLU-S	NLU-I	BANKING	CLINIC
LLaMA-2	2	89.1	96.7	72.7	3.9	**48.0**	16.7	**94.0**	42.3	22.3	38.4	59.1
+ CEPE	2 + 18	90.7	**98.4**	71.9	**46.7**	47.1	22.8	**94.0**	**48.9**	30.4	42.5	62.4
	2 + 38	**92.9**	98.0	**73.2**	45.5	47.5	**25.1**	93.3	48.8	**31.6**	**46.0**	**62.8**
LLaMA-2†	40	94.3	98.7	74.7	52.3	87.7	54.8	95.1	76.7	62.1	50.4	72.0

Table 4.3: ICL results averaged across 3 seeds. k is the number of demonstrations. All models use 2 demonstrations in the decoder, and we add $+k'$ demonstrations to the encoder for CEPE. † denotes the oracle setting with $k = 40$ demonstration in the decoder.

Models. Our baseline is LLaMA-2 with 2 demonstrations. For CEPE, we add additional demonstrations in the encoder. We also compare with an "oracle", where the LLaMA-2 decoder takes 40 demonstrations. Note that the oracle is significantly more expensive. More details are in §A.3.2.

Results. The results are shown in Table 4.3. We first observe that compared to the decoder-only baseline, CEPE can effectively use the additional demonstrations from the encoder context; the performance further increases or remains consistent with more demonstrations in the encoder. However, there is still a large gap to the 40-demonstration oracle. Our hypothesis is that in-context learning requires both query-demonstration interactions and demonstration-demonstration interactions, which CEPE cannot provide. Regardless, CEPE can be always applied on top of the decoder-only model to add additional demonstrations, with little extra computational and memory cost.

4.4 Instruction-tuned Models for Long Text Understanding

In this section, we show that applying CEPED on LLaMA-2-Chat produces instruction-following models that can leverage long inputs.

		Question Answering			Summarization		
	Total tokens	NQA	Qspr	QALT	GvRp	SSFD	QMSum
LLaMA-2-Chat	2K	17.1	14.6	28.6	16.0	16.4	19.3
+ CEPED	2K + 2K	19.5	**20.5**	**30.2**	**16.5**	16.4	**19.6**
	2K + 30K	21.6	19.9	29.6	15.8	**16.7**	19.5
	2K + All	**21.9**	19.9	29.6	15.9	**16.7**	19.5
LLaMA-2-32K Instruct	32K	12.2	18.1	41.6	19.9	10.0	10.3

Table 4.4: ZeroScrolls validation results. The total number of tokens includes both the input and generated tokens. For NarrativeQA (NQA) and Qspr (QASPER), we report the F1 scores. For QALT (QuALITY), we report accuracy. For GovReport (GvRp), SummScreenFD (SSFD), and QMSum, we report the ROUGE-L scores. CEPED uses 2K tokens in the decoder, and additional tokens are inputted through the encoder.

Datasets. ZeroScrolls [54] is a collection of zero-shot long-context understanding tasks that require instruction-following abilities. Specifically, we test on NarrativeQA, QASPER, QuALITY, GovReport, SummScreenFD, and QMSum, which all have large validation sets made available. We follow the formats and instructions of [54] for each dataset, except the long text is placed before the instructions.

Models. For CEPED, we use $2,048$ tokens in the decoder and put the remaining tokens to the encoder as chunks of 256 tokens. We compare with LLaMA-2-Chat and LLaMA-2-32K Instruct,[3] which was fine-tuned on multi-round conversational data as well as long-context summarization and QA data. We allow the model to generate $1,024$ tokens for the summarization tasks and 50 tokens for the question answering tasks.

Results. Table 4.4 shows that CEPED improves upon LLaMA-2-Chat with 2K tokens across all tasks. The performance of CEPED improves or remains consistent as we scale up the number of tokens in the context window. Notably, CEPED improves upon LLaMA-2-Chat by 3 points in F1 scores on NQA, which has the longest

[3] `https://huggingface.co/togethercomputer/Llama-2-7B-32K-Instruct`

average number of words among all ZeroScrolls tasks. Furthermore, CEPED outperforms LLaMA-2-32K Instruct on 4 out of the 6 tasks, despite being trained on unlabeled data. We provide qualitative examples and analysis in §A.3.3.

Chapter 5

Ablations

5.1 Ablation Studies

We conduct comprehensive ablations to show the effectiveness of our training data mixture, pre-training and fine-tuning the encoder, and the warmup training stage. We also ablate to verify the effectiveness of the KL divergence loss in CEPED.

5.1.1 Training Settings

	Long-context (total Tokens)		**Retrieval-augmented (k)**	
	8K	**32K**	**8**	**50**
CEPE	3.97	3.91	4.57	4.55
w/ RetDoc	4.01	3.99	4.53	4.50
w/ $\text{RP}_{\text{train-cat}}$ only	4.01	3.96	4.56	4.54
w/ $\text{RP}_{\text{train-filter}}$ only	3.96	3.89	4.75	4.72
w/ Frozen encoder	4.01	3.99	4.62	4.61
w/ Random encoder	4.03	4.02	4.60	4.60
w/o Warmup	4.03	4.02	4.61	4.61

Table 5.1: Test perplexity in long-context and retrieval-augmented language modeling, averaged over all datasets. Full results are in Table A.6 and Table A.7.

Training with retrieved documents. Even though CEPE was trained with long documents, it also achieves strong performance on retrieval-augmented applications. In this subsection, we study a different data strategy: training CEPE with retrieved documents. Specifically, we pair the training sequences described in §3.2 with retrieved passages from the retrieval split of RP using Contriever. We use the first 256 tokens of the decoder input X as the query and retrieve $k = 16$ passages to form the additional contexts $\mathcal{C}$. Then, we train CEPE with retrieved passages (RetDoc) using the same training settings as in §3.3.

As shown in Table 5.1, we find that training on RetDoc results in slightly stronger performances in the retrieval-augmentation setting, but the results on long documents are worse. Augmenting a pre-training corpus with retrieval contexts can be extremely computationally and storage expensive. CEPE's simple long-document data strategy achieves a good balance between efficient training and strong performance on both long-context and retrieval-augmented applications.

Choices of unlabeled data. In Table 5.1, we show the results of training CEPE with only the filtered documents from the ArXiv and Books domains ($\mathbf{RP}_{\textbf{train-filter}}$), and only the concatenated RP documents ($\mathbf{RP}_{\textbf{train-cat}}$). We find that training on $\text{RP}_{\text{train-filter}}$ is more beneficial for the long-document setting and training on $\text{RP}_{\text{train-cat}}$ is better for the retrieval setting, but using a mixture of both leads to a more balanced and generalizable model. Our findings corroborate with the recent work on long-context data engineering [15].

Encoder training. Finally, we investigate how to best train the encoder. To this end, we train CEPE with (1) freezing the encoder after pre-training and the warmup stage, (2) training with a randomly initialized encoder, and (3) using the pre-trained model without the warmup stage. As shown in Table 5.1, we find that the pre-training the encoder, copying warmup stage, and fine-tuning the encoder during training are

both crucial for strong performance.

5.1.2 KL Divergence

The key component of CEPED is the KL Divergence loss. To understand the importance of this auxiliary loss, we explore the performance of CEPED when trained without the KL Divergence loss as well as with difference coefficients for each loss. Let $\mathcal{L}_{CE}, \mathcal{L}_{KL}$ be the cross entropy loss and the KL Divergence loss, respectively. Then, the total loss is $\mathcal{L} = c_{CE}\mathcal{L}_{CE} + c_{KL}\mathcal{L}_{KL}$. The results are shown in Table A.8. We find that the KL Divergence loss is especially crucial for summarization tasks and QALT.

Chapter 6

Conclusion

We propose CEPE to extend the context window of existing language models. The key idea is to leverage a small encoder and cross-attention modules to process long inputs and achieve low memory and computational complexity. Compared to existing methods, CEPE extrapolates to input lengths well beyond the training length, while remaining efficient and effective. Consequently, CEPE augments pre-trained models to be performant on both long-context and retrieval-augmented applications. We also show that CEPED can be applied to instruction-tuned models with additional contexts using an auxiliary loss with only unlabeled data.

We believe that there is still room for improvement through better data to train flexible and robust models. Given the recent advances in long-context language models, there are still many important research questions to answer. For example, what applications do we want to use long-context language models for and how do we construct the evaluation suite to reflect such use cases? In our experiments, we primarily focused on perplexity on long documents as a sign of long-context abilities, as is done in many previous works, but differences in perplexity are not always reflected in downstream tasks. Other existing benchmarks such as ZeroSCROLLS are no longer long enough to evaluate recent models, as its datasets are mostly under 10K tokens,

whereas recent models are pushing 100K or even 1M tokens as input length. Thus, there is a need for new evaluation suites for long-context models.

Furthermore, the best approach to training a long-context language model in terms of its architecture and training data mixture is still not solved. We hope our work can be a useful and accessible tool for the community to study long-context models in diverse applications.

Appendix A

Appendix

A.1 Training Details

A.1.1 Pre-training Encoder

The encoder follows the configuration of RoBERTa-large [41] – it has 24 layers, a hidden size of 1024, and 16 attention heads. However, we use the architecture of LLaMA-2, which means that the vocabulary size is different and the attention module contains an additional output projection. We refer readers to [41] and [63] for more details.

We pre-trained the encoder for 100K steps on RP using the masked language modeling objective [14]. We used a batch size of 2048 sequences, where each sequence consisted of 512 tokens. The learning rate was set to 10^{-3} with a warm-up of the first 4% of the steps. We used eight A6000 GPUs with a gradient accumulation of 16. Furthermore, we employed a masking rate of 30% and disabled the next sentence prediction objective. We always replace the token with the `[MASK]` token if it is masked instead of replacing it with a random token or the original token. Finally, we used the AdamW optimizer [42] with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$, as implemented by the HuggingFace Transformers library [70].

A.1.2 Training CEPE

The attention module in LLAMA-2 consists of four projection matrices: key, value, query, and output. In contrast to original transformers [65], the output projection matrix is used as an additional attention output projection. When we first insert the cross-attention layers into the decoder, we initialize the weights of the key, value, and query projection matrices with the respective weights from the decoder's self-attention layer in the same transformer block. Furthermore, since the hidden dimension of the encoder is smaller than the hidden dimension of the decoder, $d < D$, we use only copy the first d rows of the key and value projection matrices from the self-attention module to the cross-attention module. Lastly, the output projection matrix is initialized with all zeros. While we did not investigate this initialization in detail, the intuition is that we want the model should use the tokens from the encoder using a similar mechanism as the decoder uses for its own tokens. However, we want the model to learn the output projection from scratch, as it may be too disruptive to have doubled the number of attention modules.

Then, we employ a warmup initialization method that simply trains the model to copy the input tokens from the encoder to the decoder. Specifically, we use the same inputs $X = \mathcal{C}$ for both $\mathcal{M}_{\text{enc}}$ and $\mathcal{M}_{\text{dec}}$, and X consists of $n = 256$ tokens. However, for the encoder, we chunk X into $k = 4$ sequences of 64 tokens to construct $\mathcal{C}$. This step was trained for 4K steps with a batch size of 128 and peak learning rate of 5×10^{-4}, which totals to 131M tokens. We noticed that the model quickly learned to copy the input tokens from the encoder to the decoder, and the loss was close to zero after just 1K steps. The intuition behind this initialization strategy was to instill a strong inductive bias between the encoder input and decoder outputs. From our early experiments, we found that this initialization strategy helped stabilize the later training stage.

Finally, we train CEPE for 20K steps with a batch size of 128. We use eight A100

80 GB GPUs with a per-device batch size of 2 and gradient accumulation of 8, which took approximately 750 GPU hours. We also use a peak learning rate of 3×10^{-4} with a warm-up of 4% of the steps and a cosine learning rate schedule. We use the AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$.

During the standard training, we also use masking to inject noise into the parallel encodings. Intuitively, additional contexts may not always be in chunks of exactly 256 tokens in practice; for example, in open-domain question answering, the retrieved passages may vary in length. Furthermore, we may not always use exactly $k = 16$ contexts during inference, so we may want to train CEPE with instances with different k. To this end, for each encoder context $C_i \in \mathcal{C}$, we apply masking with a probability of 0.3. When masking, we mask out the entire context C_i with probability 0.1. With probability 0.9, we mask out the suffix tokens $x_{|C_i|-t+1}, \ldots, x_{|C_i|}$, where $x_1, \ldots, x_{|C_i|}$ are the tokens in C_i and $t \sim U(1, |C_i|)$. We choose to only mask out suffixes to maintain a natural distribution for encoder inputs. While we did not study the masking rate extensively, we found in preliminary experiments that masking did not hurt perplexity while improving performance on downstream tasks. We leave further explorations on masking in the encoder input for future work.

A.1.3 Training CEPED

We leverage an additional distillation loss for CEPED. For encoder input $\mathcal{C}$ and decoder input X, we first calculate the logits $\mathcal{M}_{\text{dec}}(\text{concat}(\mathcal{C}, X))$ by running forward passes with the original model parameterized by $\mathcal{M}_{\text{dec}}$. Due to storage constraints, we only store the top 50 likelihoods and their indices in the vocabulary for each token in X, following [1, 3].

Then, during training, we define the distillation loss as the KL Divergence between the teacher model's probability distribution and the student model's probability distribution for the previously stored top 50 tokens. Concretely, our distillation loss is

defined as follows:

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(\mathcal{M}_{\text{dec}}(S)||\mathcal{M}_{CEPE}(\mathcal{C}, X))$$

where $S = \text{concat}(\mathcal{C}, X)$, $\mathcal{M}_{\text{dec}}(S)$ is the probability distribution of the top 50 tokens for X, and $\mathcal{M}_{CEPE}(\mathcal{C}, X)$ takes $\mathcal{C}$ as the encoder input and X as the decoder input and outputs the probability distribution of the same 50 tokens on X.

Although [3] also uses an additional category that represents the sum of all other tokens' probabilities, we found that this may cause the KL Divergence to be undefined when the sum of other probabilities is 0. For our main model, we use a coefficient of 2 in front of $\mathcal{L}_{\text{KL}}$ when adding to the cross-entropy loss to calculate the total loss. We experiment with this coefficient in §5.1.1.

A.2 Baseline Implementations

REPLUG. Although REPLUG [57] was introduced as a method to augment language models with retrieval, we found that the technique of interpolating logits from separate forward passes can also transfer well to the long context setting. Among the methods that we compare to, REPLUG uniquely improves performance upon the base model in both the long-context and the retrieval-augmented LM settings. This gives us an additional point of comparison across the two settings.

Following the original authors, we use Contriever [26] to calculate the scores for each previous context by using the first 256 tokens following the previous context as the query in the long-context setting. We did not include the additional memory and inference time costs of calculating the Contriever scores in our evaluation.

	Stride	ArXiv	Book	PG19	ProofPile	CodeParrot
Total Tokens = 8192						
StreamingLLM	1	2.823	6.381	7.817	2.522	1.848
	2048	2.740	6.327	7.783	2.437	1.806

Table A.1: Performance of StreamingLLM with different stride lengths.

StreamingLLM. We follow the implementation of StreamingLLM from the original authors[1] [71]. Specifically, we use their best settings, where we enable the positional shifts and cache 4 sink tokens and 2044 recent tokens.

The original code evaluates the model using a stride of 1 token at a time, where the cache is updated after every token, but this is not feasible for our large-scale evaluation. Therefore, we use a stride of 2048 tokens, and we update the cache after each stride. We show the difference in performance between the two settings in Table A.1, and we found that StreamingLLM benefits from using a larger stride. We leave future exploration in this direction to future work.

A.3 Evaluation Settings

A.3.1 Open-domain Question Answering

The full results for the open-domain question answering experiments are shown in Table A.2. REPLUG only uses up to $k = 30$ passages due to memory constraints, and LLaMA-2 has a window size of 4096, which limits k to 20. While LLaMA-2-32K can use more than $k = 60$ passages with a context size of 32K, we only use up to 60 passages due to the cost of generation. For each demonstration, we only show the top 1 retrieved passages instead of the top k passages.

[1] https://github.com/mit-han-lab/streaming-llm

	k	NQ	TQA	PopQA
LLaMA-2	1	28.37	56.44	27.17
	5	31.91	61.08	33.83
	10	32.27	62.09	34.67
	15	31.19	61.35	33.67
	20	30.39	60.31	31.35
LLaMA-2-32K	10	30.64	56.00	32.38
	15	31.27	56.98	33.48
	20	31.97	57.28	33.75
	30	30.66	57.57	33.90
	60	30.58	57.03	34.61
REPLUG	5	31.27	61.21	32.40
	10	31.52	61.35	32.31
	15	30.80	60.89	31.62
	20	30.30	60.41	31.11
	30	29.78	59.99	30.27
CEPE	10	32.27	62.09	34.67
	15	33.30	62.30	34.67
	20	33.85	62.26	34.83
	30	33.91	62.33	34.85
	60	34.07	62.26	34.98

Table A.2: Open-domain QA results. We report exact match scores for the Natural Questions(NQ) test set, TriviaQA(TQA) validation set, and PopQA test set. All models use two-shot in-context learning. k is the number of retrieved passages, and CEPE uses the top 10 passages in the decoder and all passages in the encoder.

A.3.2 In-context Learning

For our in-context learning experiments, we use the datasets commonly used in previous works [75, 43, 20, 52]: SST-2 [58], MovieReview [**?**, MR]]pang-lee-2005-seeing, AGNews [74], SST-5 [58], TREC [66], DBPedia [74], NLU [40], BANKING77 [7], CLINIC150 [35]. We follow the prompts used in [52] for all datasets. During evaluation, we first calculate the log-likelihood of each option and select the option with the highest likelihood. We sample the in-context learning demonstrations from the training set such that each label has an equal number of demonstrations (except for possible remainders).

Furthermore, we first calculate the accuracy for each dataset using four different metrics: likelihood, likelihood normalized for length, calibrated likelihood, and cali-

brated likelihood normalized for length. We calibrate using Domain Conditional PMI [22], but use the empty string as the domain string for all datasets for simplicity. We then choose the metrics that yield the highest score for the LLaMA-2 model in the two-shot setting and apply the same metrics to all other models. The metrics used for each dataset are shown in Table A.3. In this work, we did not investigate how to best calibrate CEPE in ICL settings. We leave these explorations for future work.

	Normalized	Calibrated
SST2	No	Yes
MR	No	No
AGNews	No	No
SST5	No	Yes
TREC	No	No
TREC-F	No	No
DBPedia	Yes	Yes
NLU-S	Yes	Yes
NLU-I	No	No
BANKING	No	No
CLINIC	No	No

Table A.3: Metrics used for each dataset. For normalization, we divide the log-likelihood by the length of the prompt. For calibration, we use Domain Conditional PMI [22] with the empty string as the domain string for all datasets for simplicity.

A.3.3 ZeroScrolls

We use a subset of the ZeroScrolls [54] with large validation sets: NarrativeQA [32], Qasper [13], QuALITY [47], GovReport [23], SummScreenFD [8], and QMSum [76]. Specifically, we use the validation sets of these datasets made available by Scrolls [55], as they are preprocessed in similar fashion.

However, we follow the same evaluation setup as ZeroScrolls, where models are evaluated in the zero-shot setting. We also use the same evaluation metrics as ZeroScrolls for each dataset. For the question answering datasets (NarrativeQA, Qasper, and QuALITY), we allow the model to generate up to 50 tokens, and we use

greedy decoding.

For the summarization datasets (GovReport, SummScreenFD, and QMSum), we allow the model to generate up to $1,024$ tokens, following the original authors. For SummScreenFD and QMSum, we use greedy decoding, and for GovReport we use nucleus sampling [21] with a temperature of 1.0 and top-p of 0.95 and a minimum generation length of 10 tokens.

This is because GovReport has a much longer gold summary than the other datasets, and sampling methods are typically used in long-generation settings. Furthermore, greedy decoding may degenerate. The minimum generation length helps prevent trivial outputs, such as empty strings. To account for the randomness in the sampling method, we averaged GovReport performance over 3 seeded runs, and we found that the standard deviation is less than 0.20 ROUGE-L scores in all settings.

Furthermore, we hypothesize that LLaMA-2-32K Instruct may overfit to specific domains such as GovReport due to being trained on BookSum, a summarization dataset that also has long gold summaries [33].

We also show some generation examples in Table A.4 and A.5. We find that CEPED can especially benefit from the additional contexts in the encoder in the QA datasets, where the answer may be localized to just one small part of the entire input. On the other hand, summarization tasks do not catastrophically fail when the model only has access to only part of the input, as the model can still generate a coherent summary for the provided context, achieving reasonable ROUGE-L scores.

A.4 Ablations

A.4.1 Training Settings

We show the full results for the ablation studies in Table A.6 and A.7, and A.8.

Encoder Input C_1:
We propose a novel pre-training method called BRLM, which can effectively alleviates the distance between different source language spaces. Our proposed approach significantly improves zero-shot translation performance, consistently surpassing pivoting and multilingual approaches. Meanwhile, the performance on supervised translation direction remains the same level or even better when using our method. Related Work In recent years, zero-shot translation in NMT has attracted widespread attention in academic research. Existing methods are mainly divided into four categories: pivot-based method, transfer learning, multilingual NMT, and unsupervised NMT. Pivot-based Method is a common strategy to obtain a source→target model by introducing a pivot language. This approach is further divided into pivoting and pivot-synthetic. While the former firstly translates a source language into the pivot language which is later translated to the target language BIBREF4, BIBREF5, BIBREF12, the latter trains a source→target model with pseudo

Encoder Inputs $[C_2, \ldots, C_{17}]$ Omitted...

Decoder Input X:
tokens are selected to be masked. Among the selected tokens, 80% of them are replaced with [MASK] token, 10% are replaced with a random BPE token, and 10% unchanged. The prediction accuracy of masked words is used as a stopping criterion in the pre-training stage. Besides, we use fastalign tool BIBREF34 to extract word alignments for BRLM-HA. Additional Decoder Input Omitted...
You are given a scientific article and a question. Answer the question as concisely as you can, using a single phrase or sentence if possible. If the question cannot be answered based on the information in the article, write "unanswerable". If the question is a yes/no question, answer "yes", "no", or "unanswerable".
Question:
what are the pivot-based baselines?
Answer:

Model Outputs:
LLaMA-2-Chat output: unanswerable.
CEPE output with encoder contexts: pivot-based baselines include pivoting and pivot-synthetic.
Gold answers: pivoting, pivoting$_{\text{m}}$

Table A.4: ZeroScrolls generation example on the Qasper dataset. CEPE sees the entire article through the decoder and the encoder, whereas LLaMA-2-Chat only sees a 2K token window. For brevity, we omit part of the decoder input and only show 2 out of $k = 17$ encoder inputs for CEPE.

Encoder Input C_1:

Phoebe: Almost sunrise. Do you think you're ready to try the window again? Prue: Yeah, yeah, but Abraxas will be ready for us here. We have to take him by surprise, go where we're most powerful, where we're most connected. [Cut to the park. Prue, Piper and Phoebe have joined hands around a stone.] Prue, Piper and Phoebe: "Hear now the words of the witches, the secrets we hid in the night, the oldest of Gods are invoked here, the great work of magic is sought." [Cut to Abraxas undoing the spell that gave them their powers.] [Cut back to the girls.] Prue, Piper, Phoebe: "In this night, and in this hour we call upon the ancient power." [Cut back to Abraxas. The pages of the Book of Shadows turn over. He turns back to the spell and continues to read it backwards.] [Cut back to the girls.] Prue, Piper, Phoebe: "Hear

Encoder Input C_2:

hoebe: Honey, we've done a lot of good as witches too. You know that. But that doesn't mean that bad things still aren't gonna happen. But just because we can't help that, doesn't mean that is our fault. Piper: Or yours. Phoebe: Come here. (They all hug.) Grams' Voice: The Power of Three. (They look around.) Phoebe: Okay, did you guys hear that? Piper: Mm-hmm. Prue: Grams? Grams' Voice: The Power of Three. Prue: How is that possible? Phoebe: I don't know but it's gotta to be a good thing if we're all hearing it. Piper: Maybe it means we're connected. Phoebe: To be as one. Maybe that's what the Book meant by the Power of One. Piper: The three of us working as one. Prue: Which is what the triquetra represents. (Phoebe looks at her watch)

Encoder Inputs $[C_3, \ldots, C_{44}]$ Omitted...

Decoder Input X:

the secrets we hid in the night, the oldest of Gods are invoked here, the great work of magic is sought." [Cut back to Abraxas. The Book disappears.] [Cut back to the girls. The Book drops out from the sky and lands on the rock.] Phoebe: Prue, Piper... Prue: Keep chanting. Additional Decoder Input Omitted... You are given a script of a TV episode. Summarize the episode in a paragraph.

Model Outputs:

LLAMA-2-CHAT output (ROUGE-L = 19.1): Prue, Piper, and Phoebe find an old Book of Shadows in the attic of their home, which they believe to be the original Book of Shadows. They start to read it, but Abraxas, the demon who wrote the book, appears to them, warning them to stop reading the book because he has returned to Earth and is seeking the power in the book to take over. The three sisters are able to defeat Abraxas and close the book, but the experience causes them to reflect on their lives as witches and their bond as sisters..

CEPE output with encoder contexts (ROUGE-L = 19.3): In this episode of Charmed, Prue is worried about the future because Andy is not coming back and they don't know what their new life will be like. Prue and Piper have a hard time coming to terms with the fact that their powers are no longer needed and they are losing their sense of purpose and identity. At the same time, they are struggling to cope with their feelings for each other and their relationship. Phoebe, meanwhile, has found a new source of power and is learning to control her abilities in a new way. As the three sisters work together to defeat an evil spirit that is threatening their home, they are able to come to terms with their powers and their relationship and find a new sense of purpose and identity.

Gold answers: The sisters face a demon named Abraxas , who steals the Book of Shadows and undoes some of the spells they've cast (in reverse order), releasing old vanquished enemies of the Charmed Ones and forcing them to remember the vanquishing spells from memory alone. The sisters have one chance to recapture the Book of Shadows or they will lose their powers forever. The sisters also meet their new neighbors, Jenny and her uncle Dan Gordon . Phoebe and Piper learn through the Wiccan community that because their anniversary of activating their inherent powers falls on an equinox , a wiccan holy day , each of their powers will be more developed and greatly magnified but only temporarily.

Table A.5: ZEROSCROLLS generation example on the SummScreenFD dataset. CEPE sees the entire TV script through the decoder and the encoder, whereas LLAMA-2-CHAT only sees a 2K token window. For brevity, we omit part of the decoder input and only show 2 out of $k = 44$ encoder inputs for CEPE.

	ArXiv	Book	PG19	ProofPile	CodeParrot
Total Tokens = 4096					
CEPE	2.579	6.292	7.536	2.396	1.763
w/ RetDoc	2.649	6.340	7.586	2.465	1.775
w/ RP Only	2.633	6.335	7.604	2.446	1.766
w/ AB Only	2.569	6.287	7.525	2.386	1.772
w/ Frozen Encoder	2.631	6.353	7.603	2.446	1.785
w/ Random Encoder	2.680	6.374	7.617	2.488	1.797
w/ No Copy Init	2.678	6.372	7.613	2.487	1.796
Total Tokens = 8192					
CEPE	2.496	6.049	7.372	2.219	1.715
w/ RetDoc	2.553	6.089	7.417	2.278	1.724
w/ RP Only	2.543	6.085	7.434	2.262	1.718
w/ AB Only	2.485	6.040	7.357	2.208	1.720
w/ Frozen Encoder	2.541	6.099	7.430	2.261	1.734
w/ Random Encoder	2.571	6.108	7.439	2.291	1.739
w/ No Copy Init	2.572	6.113	7.439	2.292	1.739
Total Tokens = 32768					
CEPE	2.421	6.015	7.204	2.218	1.702
w/ RetDoc	2.546	6.088	7.280	2.332	1.726
w/ RP Only	2.497	6.059	7.271	2.288	1.709
w/ AB Only	2.396	5.995	7.178	2.195	1.702
w/ Frozen Encoder	2.520	6.091	7.282	2.297	1.739
w/ Random Encoder	2.571	6.108	7.303	2.346	1.752
w/ No Copy Init	2.571	6.110	7.301	2.346	1.752

Table A.6: Test perplexity for all ablation settings in the long-context language modeling evaluation setting.

	ArXiv	Book	C4-RP	CC	Github	StackEx	Wiki	Avg.
Total Tokens = 2048 ($k = 0$)								
CEPE	3.486	6.481	6.884	5.319	1.793	3.709	4.302	4.568
w/ RetDoc	3.413	6.399	6.854	5.263	1.788	3.694	4.287	4.528
w/ RP Only	3.485	6.479	6.901	5.313	1.777	3.700	4.281	4.562
w/ AB Only	3.505	6.504	7.185	5.444	1.859	4.018	4.763	4.754
w/ Frozen Encoder	3.501	6.495	6.933	5.505	1.821	3.734	4.323	4.616
w/ Random Encoder	3.426	6.442	6.904	5.541	1.838	3.728	4.338	4.602
w/ No Copy Init	3.452	6.459	6.914	5.546	1.842	3.732	4.344	4.613
Total Tokens = 7168 ($k = 20$)								
CEPE	3.475	6.463	6.875	5.266	1.782	3.703	4.296	4.551
w/ RetDoc	3.413	6.393	6.839	5.169	1.779	3.693	4.286	4.510
w/ RP Only	3.479	6.467	6.894	5.249	1.767	3.696	4.276	4.547
w/ AB Only	3.491	6.481	7.140	5.401	1.846	4.004	4.738	4.729
w/ Frozen Encoder	3.485	6.482	6.930	5.500	1.815	3.727	4.318	4.608
w/ Random Encoder	3.426	6.442	6.904	5.540	1.837	3.727	4.337	4.602
w/ No Copy Init	3.447	6.457	6.913	5.545	1.841	3.721	4.332	4.608
Total Tokens = 14848 ($k = 50$)								
CEPE	3.467	6.457	6.881	5.273	1.777	3.701	4.292	4.550
w/ RetDoc	3.413	6.392	6.835	5.098	1.776	3.692	4.287	4.499
w/ RP Only	3.472	6.465	6.900	5.243	1.762	3.693	4.274	4.544
w/ AB Only	3.480	6.471	7.114	5.412	1.838	3.994	4.719	4.718
w/ Frozen Encoder	3.474	6.474	6.930	5.509	1.814	3.723	4.316	4.606
w/ Random Encoder	3.426	6.442	6.904	5.540	1.837	3.726	4.337	4.602
w/ No Copy Init	3.445	6.456	6.913	5.545	1.841	3.716	4.329	4.606

Table A.7: Test perplexity on RedPajama across all domains in the retrieval-augmented setting for all ablation experiments. k is the number of additional contexts used. Avg. is the macro average across all domains.

	Question Answering			Summarization		
c_{KL}	NQA	Qspr	QALT	GvRp	SSFD	QMSum
Total Tokens = 4K						
2	19.5	20.5	**30.2**	**16.5**	**16.4**	**19.6**
1	**21.6**	20.7	27.2	16.3	5.3	4.7
0	21.3	21.0	**27.4**	14.6	14.9	15.6
Total Tokens = 16K						
2	20.6	19.9	**29.6**	15.9	**16.8**	**19.4**
1	21.8	**20.6**	26.8	**16.0**	15.2	16.1
0	**22.9**	**20.6**	26.4	14.8	5.2	4.7
Total Tokens = 32K						
2	21.6	19.9	**29.6**	15.8	**16.7**	**19.5**
1	**22.7**	**20.6**	26.8	**16.0**	15.2	15.8
0	22.3	**20.6**	26.4	14.6	5.2	4.7
All Tokens						
2	21.9	19.9	**29.6**	**15.9**	**16.7**	**19.5**
1	22.6	**20.6**	26.8	**15.9**	15.2	15.8
0	**23.0**	**20.6**	26.4	14.6	5.2	4.7

Table A.8: ZeroScrolls results using different losses during training, where c_{KL} is the coefficient of the KL Divergence loss. $c_{CE} = 1$ is the coefficient of the Cross-Entropy loss for all experiments.

Bibliography

[1] A. Askell, Y. Bai, A. Chen, D. Drain, D. Ganguli, T. Henighan, A. Jones, N. Joseph, B. Mann, N. DasSarma, N. Elhage, Z. Hatfield-Dodds, D. Hernandez, J. Kernion, K. Ndousse, C. Olsson, D. Amodei, T. Brown, J. Clark, S. McCandlish, C. Olah, and J. Kaplan. A general language assistant as a laboratory for alignment, 2021.

[2] Z. Azerbayev, E. Ayers, and B. Piotrowski. Proofpile: A pre-training dataset of mathematical texts, 2023.

[3] Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli, T. Henighan, N. Joseph, S. Kadavath, J. Kernion, T. Conerly, S. El-Showk, N. Elhage, Z. Hatfield-Dodds, D. Hernandez, T. Hume, S. Johnston, S. Kravec, L. Lovitt, N. Nanda, C. Olsson, D. Amodei, T. Brown, J. Clark, S. McCandlish, C. Olah, B. Mann, and J. Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022.

[4] A. Bertsch, U. Alon, G. Neubig, and M. R. Gormley. Unlimiformer: Long-range transformers with unlimited length input. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

[5] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. B. Van Den Driessche, J.-B. Lespiau, B. Damoc, A. Clark, D. De Las Casas, A. Guy, J. Menick, R. Ring, T. Hennigan, S. Huang, L. Maggiore, C. Jones, A. Cassirer,

A. Brock, M. Paganini, G. Irving, O. Vinyals, S. Osindero, K. Simonyan, J. Rae, E. Elsen, and L. Sifre. Improving language models by retrieving from trillions of tokens. In *International Conference on Machine Learning (ICML)*, volume 162, pages 2206–2240, 2022.

[6] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[7] I. Casanueva, T. Temčinas, D. Gerz, M. Henderson, and I. Vulić. Efficient intent detection with dual sentence encoders, 2020.

[8] M. Chen, Z. Chu, S. Wiseman, and K. Gimpel. 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, Dublin, Ireland, May 2022. Association for Computational Linguistics.

[9] S. Chen, S. Wong, L. Chen, and Y. Tian. Extending context window of large language models via positional interpolation, 2023.

[10] Y. Chen, S. Qian, H. Tang, X. Lai, Z. Liu, S. Han, and J. Jia. Longlora: Efficient fine-tuning of long-context large language models, 2023.

[11] A. Chevalier, A. Wettig, A. Ajith, and D. Chen. Adapting language models to compress contexts. In H. Bouamor, J. Pino, and K. Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3829–3846, Singapore, Dec. 2023. Association for Computational Linguistics.

[12] K. M. Choromanski, V. Likhosherstov, D. Dohan, X. Song, A. Gane, T. Sarlos, P. Hawkins, J. Q. Davis, A. Mohiuddin, L. Kaiser, D. B. Belanger, L. J. Colwell, and A. Weller. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021.

[13] P. Dasigi, K. Lo, I. Beltagy, A. Cohan, N. A. Smith, and M. Gardner. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610, Online, June 2021. Association for Computational Linguistics.

[14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional Transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.

[15] Y. Fu, R. Panda, X. Niu, X. Yue, H. Hajishirzi, Y. Kim, and H. Peng. Data engineering for scaling language models to 128k context, 2024.

[16] T. Gao, H. Yen, J. Yu, and D. Chen. Enabling large language models to generate text with citations. In H. Bouamor, J. Pino, and K. Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore, Dec. 2023. Association for Computational Linguistics.

[17] A. Gu and T. Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2023.

[18] A. Gu, K. Goel, and C. Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.

[19] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M.-W. Chang. REALM: Retrieval-augmented language model pre-training. In *International Conference on Machine Learning (ICML)*, 2020.

[20] Z. Han, Y. Hao, L. Dong, Y. Sun, and F. Wei. Prototypical calibration for few-shot learning of language models. In *The Eleventh International Conference on Learning Representations*, 2023.

[21] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations (ICLR)*, 2020.

[22] A. Holtzman, P. West, V. Shwartz, Y. Choi, and L. Zettlemoyer. Surface form competition: Why the highest probability answer isn't always right. In M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7038–7051, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics.

[23] L. Huang, S. Cao, N. Parulian, H. Ji, and L. Wang. Efficient attentions for long document summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1419–1436, Online, June 2021. Association for Computational Linguistics.

[24] M. Ivgi, U. Shaham, and J. Berant. Efficient Long-Text Understanding with Short-Text Models. *Transactions of the Association for Computational Linguistics*, 11:284–299, 03 2023.

[25] H. Ivison, Y. Wang, V. Pyatkin, N. Lambert, M. Peters, P. Dasigi, J. Jang, D. Wadden, N. A. Smith, I. Beltagy, and H. Hajishirzi. Camels in a changing climate: Enhancing lm adaptation with tulu 2, 2023.

[26] G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, and E. Grave. Unsupervised dense information retrieval with contrastive learning. *Transactions on Machine Learning Research*, 2022.

[27] G. Izacard and E. Grave. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*, 2020.

[28] G. Izacard, P. Lewis, M. Lomeli, L. Hosseini, F. Petroni, T. Schick, J. Dwivedi-Yu, A. Joulin, S. Riedel, and E. Grave. Atlas: Few-shot learning with retrieval augmented language models. *arXiv preprint arXiv:2208.03299*, 2022.

[29] S. Jelassi, D. Brandfonbrener, S. M. Kakade, and E. Malach. Repeat after me: Transformers are better than state space models at copying, 2024.

[30] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension, 2017.

[31] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, and W.-t. Yih. 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, Online, Nov. 2020. Association for Computational Linguistics.

[32] T. Kočiský, J. Schwarz, P. Blunsom, C. Dyer, K. M. Hermann, G. Melis, and E. Grefenstette. The NarrativeQA reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018.

[33] W. Kryscinski, N. Rajani, D. Agarwal, C. Xiong, and D. Radev. BOOKSUM: A collection of datasets for long-form narrative summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6536–6558, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.

[34] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, et al. Natural questions: a benchmark for question answering research. *TACL*, 2019.

[35] S. Larson, A. Mahendran, J. J. Peper, C. Clarke, A. Lee, P. Hill, J. K. Kummerfeld, K. Leach, M. A. Laurenzano, L. Tang, and J. Mars. An evaluation dataset for intent classification and out-of-scope prediction. In K. Inui, J. Jiang, V. Ng, and X. Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316, Hong Kong, China, Nov. 2019. Association for Computational Linguistics.

[36] K. Lee, M.-W. Chang, and K. Toutanova. Latent retrieval for weakly supervised open domain question answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy, July 2019. Association for Computational Linguistics.

[37] K. Lee, M.-W. Chang, and K. Toutanova. Latent retrieval for weakly supervised open domain question answering. In A. Korhonen, D. Traum, and L. Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy, July 2019. Association for Computational Linguistics (ACL).

[38] X. V. Lin, X. Chen, M. Chen, W. Shi, M. Lomeli, R. James, P. Rodriguez, J. Kahn, G. Szilvasy, M. Lewis, L. Zettlemoyer, and S. Yih. Ra-dit: Retrieval-augmented dual instruction tuning, 2023.

[39] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang. Lost in the middle: How language models use long contexts, 2023.

[40] X. Liu, A. Eshghi, P. Swietojanski, and V. Rieser. Benchmarking natural language understanding services for building conversational agents. *ArXiv*, abs/1903.05566, 2019.

[41] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

[42] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[43] Y. Lu, M. Bartolo, A. Moore, S. Riedel, and P. Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In S. Muresan, P. Nakov, and A. Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland, May 2022. Association for Computational Linguistics.

[44] A. Mallen, A. Asai, V. Zhong, R. Das, D. Khashabi, and H. Hajishirzi. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Association for Computational Linguistics (ACL)*, pages 9802–9822, Toronto, Canada, July 2023. Association for Computational Linguistics.

[45] S. Min, W. Shi, M. Lewis, X. Chen, W.-t. Yih, H. Hajishirzi, and L. Zettlemoyer. Nonparametric masked language modeling. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2097–2118, Toronto, Canada, July 2023. Association for Computational Linguistics.

[46] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:27730–27744, 2022.

[47] R. Y. Pang, A. Parrish, N. Joshi, N. Nangia, J. Phang, A. Chen, V. Padmakumar, J. Ma, J. Thompson, H. He, and S. Bowman. QuALITY: Question answering with long input texts, yes! In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5336–5358, Seattle, United States, July 2022. Association for Computational Linguistics.

[48] B. Peng, J. Quesnelle, H. Fan, and E. Shippole. Yarn: Efficient context window extension of large language models, 2023.

[49] O. Press, N. Smith, and M. Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. In *International Conference on Learning Representations (ICLR)*, 2022.

[50] J. W. Rae, A. Potapenko, S. M. Jayakumar, C. Hillier, and T. P. Lillicrap. Compressive transformers for long-range sequence modelling. In *International Conference on Learning Representations*, 2020.

[51] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified

text-to-text Transformer. *The Journal of Machine Learning Research (JMLR)*, 21(140), 2020.

[52] N. Ratner, Y. Levine, Y. Belinkov, O. Ram, I. Magar, O. Abend, E. Karpas, A. Shashua, K. Leyton-Brown, and Y. Shoham. Parallel context windows for large language models. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6383–6402, Toronto, Canada, July 2023. Association for Computational Linguistics.

[53] O. Rubin and J. Berant. Long-range language modeling with self-retrieval, 2023.

[54] U. Shaham, M. Ivgi, A. Efrat, J. Berant, and O. Levy. ZeroSCROLLS: A zero-shot benchmark for long text understanding. In H. Bouamor, J. Pino, and K. Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7977–7989, Singapore, Dec. 2023. Association for Computational Linguistics.

[55] U. Shaham, E. Segal, M. Ivgi, A. Efrat, O. Yoran, A. Haviv, A. Gupta, W. Xiong, M. Geva, J. Berant, and O. Levy. 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, Dec. 2022. Association for Computational Linguistics.

[56] F. Shi, X. Chen, K. Misra, N. Scales, D. Dohan, E. Chi, N. Schärli, and D. Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning (ICML)*, 2023.

[57] W. Shi, S. Min, M. Yasunaga, M. Seo, R. James, M. Lewis, L. Zettlemoyer, and W. tau Yih. Replug: Retrieval-augmented black-box language models, 2023.

[58] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2013.

[59] J. Su, Y. Lu, S. Pan, A. Murtadha, B. Wen, and Y. Liu. Roformer: Enhanced transformer with rotary position embedding, 2021.

[60] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto. Stanford Alpaca: An Instruction-following LLaMA model, 2023.

[61] Together. Llama-2-7b-32k, July 2023.

[62] Together. Redpajama: An open source recipe to reproduce llama training dataset, April 2023.

[63] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*, 2023.

[64] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.

[65] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NIPS*, 2017.

[66] E. M. Voorhees and D. M. Tice. Building a question answering test collection. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '00, page 200–207, New York, NY, USA, 2000. Association for Computing Machinery.

[67] B. Wang, W. Ping, L. McAfee, P. Xu, B. Li, M. Shoeybi, and B. Catanzaro. InstructRetro: Instruction Tuning post Retrieval-Augmented Pretraining, 2024.

[68] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada, July 2023. Association for Computational Linguistics.

[69] T. Wolf, L. Ben Allal, L. von Werra, L. Jia, and A. Zebaze. A dataset of python files from github. `https://github.com/huggingface/blog/blob/main/codeparrot.md?version=codeparrot/codeparrot-valid-v2-near-dedup`, 2023.

[70] T. Wolf, J. Chaumond, L. Debut, V. Sanh, C. Delangue, A. Moi, P. Cistac, M. Funtowicz, J. Davison, S. Shleifer, et al. Transformers: State-of-the-art natural language processing. In *Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, 2020.

[71] G. Xiao, Y. Tian, B. Chen, S. Han, and M. Lewis. Efficient streaming language models with attention sinks, 2023.

[72] W. Xiong, J. Liu, I. Molybog, H. Zhang, P. Bhargava, R. Hou, L. Martin, R. Rungta, K. A. Sankararaman, B. Oguz, M. Khabsa, H. Fang, Y. Mehdad, S. Narang, K. Malik, A. Fan, S. Bhosale, S. Edunov, M. Lewis, S. Wang, and H. Ma. Effective long-context scaling of foundation models, 2023.

[73] D. Yoshida, A. Ettinger, and K. Gimpel. Adding recurrence to pretrained transformers for improved efficiency and context size, 2020.

[74] X. Zhang, J. J. Zhao, and Y. LeCun. Character-level convolutional networks for text classification. In *Neural Information Processing Systems*, 2015.

[75] Z. Zhao, E. Wallace, S. Feng, D. Klein, and S. Singh. Calibrate before use: Improving few-shot performance of language models. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12697–12706. PMLR, 18–24 Jul 2021.

[76] M. Zhong, D. Yin, T. Yu, A. Zaidi, M. Mutuma, R. Jha, A. H. Awadallah, A. Celikyilmaz, Y. Liu, X. Qiu, and D. Radev. QMSum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921, Online, June 2021. Association for Computational Linguistics.