

Anchor-based Large Language Models

Anonymous ACL submission

Abstract

Large language models (LLMs) predominantly employ decoder-only transformer architectures, necessitating the retention of keys/values information for historical tokens to provide contextual information and avoid redundant computation. However, the substantial size and parameter volume of these LLMs require massive GPU memory. This memory demand increases with the length of the input text, leading to an urgent need for more efficient methods of information storage and processing. This study introduces Anchor-based LLMs (AnLLMs), which utilize an innovative anchor-based self-attention network (AnSAN) and also an anchor-based inference strategy. This approach enables LLMs to compress sequence information into an anchor token, reducing the keys/values cache and enhancing inference efficiency. Experiments on question-answering benchmarks reveal that AnLLMs maintain similar accuracy levels while achieving up to 99% keys/values cache reduction and up to 3.5 times faster inference. Despite a minor compromise in accuracy, the substantial enhancements of AnLLMs employing the AnSAN technique in resource utilization and computational efficiency underscore their potential for practical LLM applications.

1 Introduction

Large language models (LLMs) primarily utilize decoder-only transformer architectures, which necessitate caching keys/values information for historical tokens during the auto-regressive inference to supply contextual information and avoid redundant computation (Wei et al., 2022; Touvron et al., 2023a; OpenAI, 2023; Touvron et al., 2023b). However, due to their immense size and high parameter count, a considerable amount of GPU memory is required for loading. Furthermore, as the length of input text grows, storing keys/values caches requires more and more GPU memory, as evidenced in in-context learning, complex instructions, and

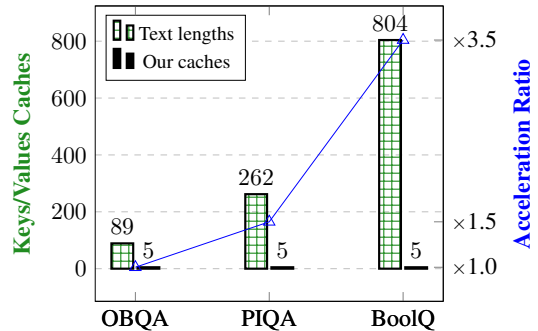


Figure 1: Keys/Values Caches and Inference Acceleration Ratio of Ours in OBQA, PIQA, and BoolQ Tasks with Five-Shot Demonstrations. The bars denote the Keys/Values Cache and text length, while the curve represents the Inference Acceleration Ratio. As text length increases, our method demonstrates an impressive reduction of up to 99% in Keys/Values Caches compared to conventional methods. Moreover, by caching prefix texts, we have enhanced inference efficiency by a factor of 3.5 compared to non-caching inference.

extended conversations (Dong et al., 2022; Jiang et al., 2023; Wang et al., 2023), which is not conducive to scenarios with limited computational resources. An alternative approach entails recalculating these extensive inputs, which, however, results in increased time overhead. Therefore, this study aims to *reduce the storage demand for keys/values caches during the inference phase of LLMs, improving the memory efficiency and, consequently, accelerating the inference speed as well.*

In a recent study, Wang et al. (2023) demonstrate that label words in prefix demonstrations can act as anchors during inference, providing an effective context compression approach for improving inference efficiency in in-context learning. However, in practical applications, not all prefix inputs or demonstrations contain label words suitable for compressing information, making the reliance on label words a less universal approach for text information compression. Additionally, Pang et al.

(2024) observe that LLMs tend to attend to only a few, yet consistent, prefix tokens during inference. However, the specific tokens utilized are often unpredictable and uncontrollable. These observations raise an intriguing question: *do natural language texts contain anchor points that compress the overall semantic information of sequences?* In this context, previous studies on sequence embeddings have shown that the hidden state of a special token in neural network models can encapsulate semantic information (Baudiš et al., 2016; Devlin et al., 2018). Furthermore, contemporary LLMs typically utilize the causal self-attention mechanism during both training and inference phases (Touvron et al., 2023a,b), attending on each prior token. This suggests that the final token in a sequence may be better suited to serve as a natural information compression point compared to other tokens, as they cannot observe future tokens. Therefore, a methodical approach that identifies and exploits sequence anchor tokens in a dependable and controllable manner is essential for compressing sequence information, effectively reducing keys/values caches, and improving inference efficiency for LLMs.

To this end, we propose novel **Anchor-based Large Language Models (AnLLMs)**, equipped with an innovative anchor-based self-attention network (AnSAN) and an anchor-based inference strategy. The AnSAN is designed to compel the models to compress sequence information into the anchor token (the last token in our implementation) during the training process, with the aid of anchor-based attention masks. During inference, the anchor-based inference strategy retains the keys/values caches of anchor tokens, which have aggregated the entire sequence information, and discards those of non-anchor tokens, thereby reducing memory demands. Specifically, the anchor-based attention masks for AnSAN serve two objectives: 1) to ensure anchor tokens attend exclusively to tokens within the same sequence, preventing attention to other sequences, and 2) to direct non-anchor tokens’ attention to previous sequence anchors, blocking the other non-anchor tokens from previous sequences. It is noteworthy that the technique of anchor-based attention bears similarities to the principles underlying sparse attention (Child et al., 2019). However, unlike the existing research that employs sparse attention to extend the context length of LLMs (Chen et al., 2023; Ratner et al., 2023), our method focuses on continually pre-training the model to compress sequence information into the anchor token.

In our implementation, we utilize the publicly available RedPajama datasets (Computer, 2023) to continuously pre-train the open-source Llama2 models (Touvron et al., 2023b), resulting in AnLLMs that incorporate our proposed anchor-based attention mechanism. Experimental results on question answering benchmarks, as depicted in Figure 1, reveal that our method achieves up to a 99% reduction in keys/values caches and up to a 3.5-fold increase in inference acceleration ratios, while maintaining comparable accuracy to the original model. Despite a minor decrease in accuracy (within 1.5%), these findings underscore the significant improvements in computational efficiency and memory utilization offered by our method.

2 Related Work

Our research is inspired by the recent investigation into the understanding of in-context learning (ICL) within LLMs by Wang et al. (2023). In their study, the authors delve into the underlying mechanisms of ICL, emphasizing the influence of label words in demonstration examples on information flow. They reveal that these label words serve as anchors, wherein semantic information converges into these anchors during inference, subsequently directing the LLMs’ final predictions. Motivated by their findings, our objective is to extend this feature to natural language modeling by guiding sequence information compression into manually designed anchor tokens, rather than solely relying on label words. This is crucial because natural language texts may not always contain an explicit label.

The most relevant method to our approach in the existing literature is the learning to compress prompts with gist tokens (Mu et al., 2023). Their approach centers around compressing task-specific prompts by fine-tuning the model using the proposed gist masking, thereby enforcing prompt compression. However, there are several crucial divergences between our study and theirs. Unlike their focus on compressing a task prompt, our objective lies in training the LLM to condense sequence information into the anchor tokens. Consequently, our approach can be universally applied to a range of tasks without requiring task-specific training, a feature not shared by gist tokens, as the anchor tokens are seamlessly incorporated into the model’s language modeling. Furthermore, our anchor-based attention masks account for information compression within a sequence and information interaction be-

tween sequences, thus extending beyond the mere compression of task prompts.

On the other hand, FlashAttention (Dao et al., 2022) and PagedAttention (Kwon et al., 2023) both present memory-efficient attention mechanisms for LLMs. While they focus on optimizing attention computation and subdividing attention processing, our proposed method offers a distinct approach that specifically targets the compression of sequence information into anchor tokens, making it orthogonal to these existing works.

3 Anchor-based Large Language Models

3.1 Background

Transformers. LLMs are predominantly implemented as transformer decoders (Vaswani et al., 2017; Touvron et al., 2023a,b), consisting of an input embedding layer followed by multiple decoder layers. Each of these decoder layers is composed of a self-attention network and a feed-forward network, succeeded by respective normalization modules. It is important to note that causal attention masks are employed within the self-attention mechanisms, ensuring that current tokens solely attend to preceding tokens.

Self-Attention Networks. Typically for decoder-only LLMs like Llama2 (Touvron et al., 2023b), self-attention networks (SANs) map queries Q , keys K , and values V into an output, as delineated in the following equations,

$$\text{SAN}(Q, K, V) = \text{Softmax}(Q, K)V, \quad (1)$$

$$\text{Softmax}(Q, K)_{i,j} = \frac{M_{i,j} \exp(Q_i K_j^T)}{\sum_k M_{i,k} \exp(Q_i K_k^T)}, \quad (2)$$

$$M_{i,j} = \begin{cases} 1, & \text{if } i \geq j \\ 0, & \text{else} \end{cases}, \quad (3)$$

where M denotes an $L \times L$ masking matrix, facilitating the current i -th token to attend to only preceding tokens whilst disregarding subsequent tokens during the training and inference phases.

Keys/Values Caches. In the application of LLMs, the keys/values caches increase with lengthy prefix texts and continuously generated tokens during the inference phase, such as in question-answering (Saad-Falcon et al., 2023), text summarization (Basyal and Sanghvi, 2023), and machine translation (Pang et al., 2024). The key and value matrices associated with tokens of prefix inputs are cached to avoid recomputation and expedite subsequent

token prediction (Radford et al., 2019). Additionally, the model generates the output token-by-token in the real-time inference process, which requires more cache memory to store the newly generated sequence. Therefore, addressing the challenges arising from the ever-expanding texts is crucial for enhancing the efficiency of LLM inference.

3.2 Anchor-based Self-Attention Networks

Given an input text with n ordered sequences, $P = \{S_1, S_2, \dots, S_n\}$, their associated anchor tokens are the last tokens that represented as $A = \{a_1, a_2, \dots, a_n\}$. The primary objective of AnSAN is to encapsulate the information of a sequence into its anchor token, with the anchor hidden states representing the comprehensive semantic information. In this manner, an AnLLM equipped with AnSAN generates subsequent tokens based on the keys/values caches of preceding tokens within the current sequence and the keys/values caches of anchor tokens from previous sequences.

Anchor-based Attention Masks. To accomplish this, we devise anchor-based attention masks, as illustrated in Figure 2. Assuming that the current token in the sequence is a non-anchor token, we allow attention towards previous non-anchor tokens within the same sequence and anchor tokens from preceding sequences, while blocking attention towards non-anchor tokens from previous sequences. This approach ensures that non-anchor tokens can only access information from anchor tokens in previous sequences and the current sequence’s information. Conversely, when the current token is an anchor token, which is the last token in the sequence, we exclusively permit its attention towards previous non-anchor tokens within the same sequence, blocking all other attention. This constraint forces the anchor token to aggregate information solely from its current sequence. Consequently, we replace Eq. (3) with anchor-based attention masks in Eq. (4) to determine the mask of the i -th token in the input text concerning the j -th token (assuming that the i -th token belongs to the k -th sequence).

$$M_{i,j} = \begin{cases} 0, & \text{if } ((w_i, w_j) \notin A) \wedge (w_j \in S_{\leq k-1}) \\ 0, & \text{else if } (w_i \in A) \wedge (w_j \in S_{\leq k-1}) \\ 1, & \text{else if } i \geq j \\ 0, & \text{else} \end{cases}, \quad (4)$$

where $S_{\leq k-1}$ represents previous $k-1$ sequences. The number 0 denotes blocking attention, whereas the number 1 indicates the opposite.

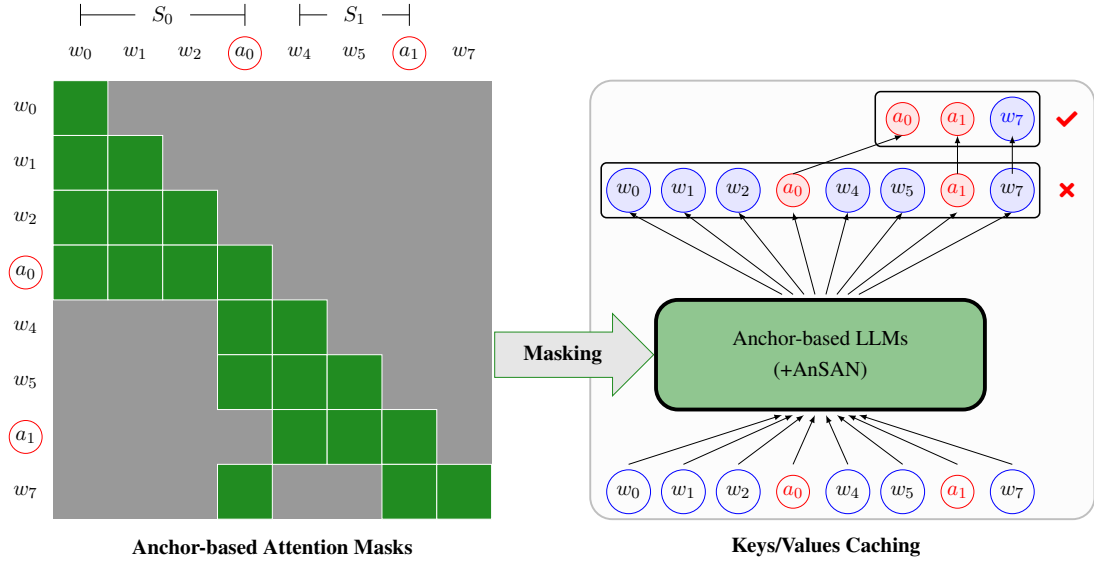


Figure 2: Anchor-based Attention Masking and Efficient Caching in Anchor-based LLMs. On the left, the gray and green squares represent the masking and unmasking operations respectively, with the circled “a” symbols denoting the anchor tokens. On the right, the shaded circles depict keys/values caches. By employing anchor-based attention masking during training, we compel the model to compress sequence information into the anchor tokens. On the right, during inference, with the AnSAN technique, AnLLMs compress information into the anchor tokens and discard the previous remaining keys/values caches, thereby facilitating an efficient caching mechanism.

Anchor Token Selection. By implementing the AnSAN mechanism for training LLMs, we can compel the model to compress sequence information into the anchor token and generate new tokens based on the anchor token information from previous sequences and non-anchor token information from the current sequence.

The challenge now lies in selecting an appropriate anchor token. In our experiment, we propose two implementation methods: one using the endpoint as the anchor token, and the other appending a new token specifically as the anchor token.

3.3 Anchor-based Inference

By training the model to compress information into the anchor token of a natural language sequence, we can optimize the inference process by modifying the keys/values caching mechanism. Specifically, during inference, upon encountering an anchor token that condenses the comprehensive semantic information of preceding tokens in the current sequence, the model can reduce the keys/values caches by deleting the caches of non-anchor tokens within that sequence.

We introduce the inference method in Algorithm 1. The function “REDUCTION” in Line 1 is utilized to remove keys/values caches when the model processes prefix texts in Line 10 or generates an anchor token during the prediction of the next

Algorithm 1 Anchor-based Inference

Require: Anchor-based LLM Θ , prefix text P with anchor tokens, keys/values cache list \mathcal{C} , predicted token w_{new} ;

Output: Generated text \mathcal{T} ;

```

1: function REDUCTION( $\mathcal{C}$ )
2:    $j \leftarrow$  last anchor index in  $\mathcal{C}$ ;
3:    $\mathcal{C} \leftarrow \{c \in \mathcal{C} \mid \text{index}(c) \geq j \text{ or } c \text{ is anchor}\}$ ;
4:   return  $\mathcal{C}$ .
5: end function
6: Initialize  $\mathcal{T}, \mathcal{C}$  as empty lists;
7:  $\mathcal{M} \leftarrow$  GetMasks( $P, \mathcal{C}$ ) using Eq. (4);
8: Update  $w_{new}, \mathcal{C}$  using Forward( $P; \mathcal{M}, \mathcal{C}, \Theta$ );
9: Append  $w_{new}$  to  $\mathcal{T}$ ;
10:  $\mathcal{C} \leftarrow$  Reduction( $\mathcal{C}$ );
11: while  $w_{new}$  is not [eos] do
12:    $\mathcal{M} \leftarrow$  GetMasks( $w_{new}, \mathcal{C}$ ) using Eq. (4);
13:   Update  $w_{new}, \mathcal{C}$  using Forward( $w_{new}; \mathcal{M}, \mathcal{C}, \Theta$ );
14:   Append  $w_{new}$  to  $\mathcal{T}$ ;
15:   if  $w_{new}$  is the anchor token then
16:      $\mathcal{C} \leftarrow$  Reduction( $\mathcal{C}$ );
17:   end if
18: end while
19: return  $\mathcal{T}$ .

```

token in Line 16. This approach aims to reduce the keys/values caches for both prefix tokens and generated outputs during real-time inference.

4 Experiments

In this section, we first detail AnLLM’s implementation, then outline the training procedure and model perplexity. Finally, we introduce the evaluation datasets and metrics.

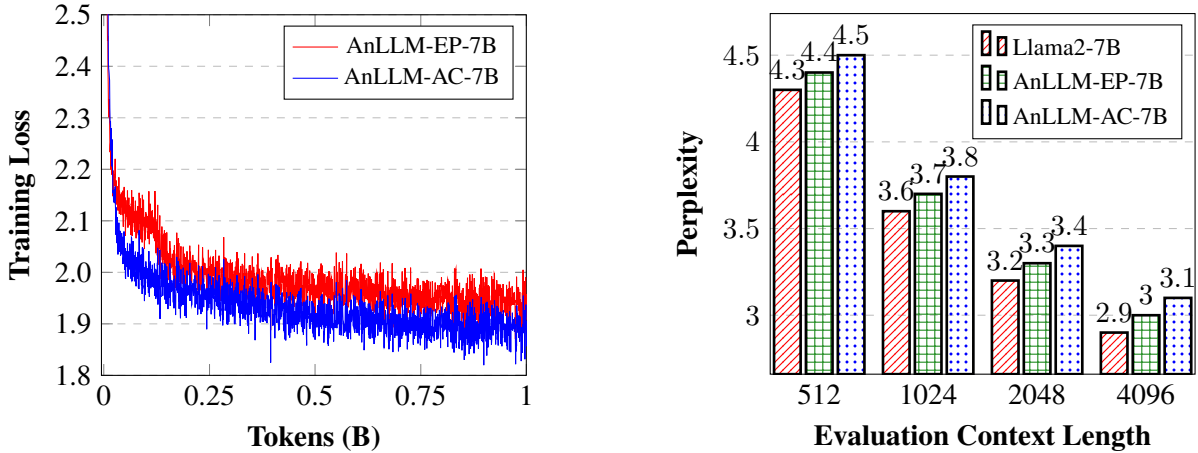


Figure 3: Training Process and Perplexity Evaluation of the Anchor-based Large Language Model.

4.1 Our Implementation

Llama2-7b (Touvron et al., 2023b) is adopted as the base model in our experiments, which is an open-source and English-centric LLM. In accordance with the principles outlined in Section 3, we present our implementations here. The crux is to identify which tokens in a sequence can be considered anchor tokens. In light of this, we describe two implementation strategies: one employs the endpoints directly, and the other involves appending a new token at the end of a sequence to serve as the anchor token. The details are as follows:

- **AnLLM-EP.** This approach uses punctuation marks within the sequence as anchor tokens. Punctuation marks, such as commas, periods, and question marks, are viewed as semantic boundaries within a sequence. As such, they can serve as anchor tokens in AnLLM. In our experiments of AnLLM-EP, we use the endpoint in English as the anchor tokens.
- **AnLLM-AC.** This strategy entails the introduction of a new token to act as the sequence anchor. In our implementation, we designate <AC> as the new token and initialize its embedding using the mean value of the embedding matrix. For training data, we use the sentence tokenizer from the NLTK package to split texts into sentences, appending <AC> at the end of each sentence as the anchor token.¹ During inference, <AC> tokens can be flexibly added to the text based on user requirements, such as adding one anchor for each demonstration, allowing for flexible and controllable sequence compression.

¹<https://www.nltk.org/api/nltk.tokenize.punkt.html>

4.2 Data and Training Procedure

Considering that AnLLMs are expected to predict subsequent tokens within the context of keys/values hidden states of anchor tokens, this presents a significant challenge for existing open-source LLMs. To this end, by substituting the self-attention networks with anchor-based self-attention networks as detailed in Section 3.2, we continually pre-train the Llama2 model using a publicly available corpus.

Data. We employ the RedPajama-Data-1T-Sample dataset (Computer, 2023) for the continuous pre-training purpose.² This dataset comprises 850,000 samples with approximately 1 billion tokens, which have been subjected to right truncation to fit the model context of 4,096.

Training Procedure. We train each model via the next token prediction objective on the dataset for one epoch, with a batch size of 512. The learning rate is set to 0.00002 and constant after a linear warmup with 20 update steps. The AdamW (Loshchilov and Hutter, 2019) with $\beta_1 = 0.9$ and $\beta_2 = 0.95$ is adopted as the gradient backtrack propagation optimizer. All the training procedures are conducted with four $8 \times$ A100 GPU machines with 40G GPU Memory.

Training Loss and Perplexity. The left-hand side of Figure 3 depicts the training loss associated with our models. The loss curves for AnLLM-EP and AnLLM-AC consistently decline to approximately 1.9, with AnLLM-AC achieving a lower loss. This observation suggests that continually pre-training an LLM using anchor-based attention

²<https://huggingface.co/datasets/togethercomputer/RedPajama-Data-1T-Sample>

masks is indeed viable, enabling the LLM to effectively learn the process of compressing sequence information into anchor tokens.

The right-hand side of Figure 3 displays the perplexity evaluation of the models with varying context lengths. Full attention is utilized to assess the language modeling capabilities of all models. Following the settings of Chen et al. (2023), the perplexity is evaluated on the test samples of the Proof-Pile datasets (Rae et al., 2020). The results demonstrate that both AnLLM-EP and AnLLM-AC models maintain a promising performance, exhibiting language modeling capacity comparable to the base model, Llama2-7B. Moreover, this finding suggests that AnLLMs are compatible with full attention, as indicated by minimal perplexity decline.

4.3 Evaluation

In our investigation, we employ a diverse collection of benchmarks with varying text lengths to evaluate our outcomes, including OpenBookQA (OBQA) (Mihaylov et al., 2018), WinoGrande (WG) (Sakaguchi et al., 2021), ARC-easy (ARC-e) and ARC-challenge (ARC-c) (Clark et al., 2018), PIQA (Bisk et al., 2020), HellaSwag (HS) (Zellers et al., 2019), SCIQ (Welbl et al., 2017), and BoolQ (Clark et al., 2019). These benchmarks provide a comprehensive evaluation of various aspects, including reasoning, comprehension, understanding of the physical world, and predicting future events. Importantly, they cover texts of varying lengths, facilitating a thorough assessment of our model’s performance across diverse tasks and text complexities, ranging from shorter input contexts in OBQA to longer texts in BoolQ. To measure the precision and efficiency of our models, we evaluate them across three dimensions using three distinct metrics for both zero-shot and five-shot settings. For AnLLM-AC in the five-shot setting, we incorporate the anchor token <AC> at the end of each demonstration.

- **Accuracy (Acc).** This conventional metric is utilized to gauge the prediction accuracy of models. In accordance with previous studies (Gao et al., 2023), we choose the options with the highest probabilities as predictions and calculate accuracy using the gold-standard labels.
- **Keys/Values Caches Reduction (C_{\downarrow}).** In the context of the five-shot evaluation, the demonstrations can be cached in GPU memory for subsequent reuse. Nevertheless, extended demonstrations may require increased memory consump-

tion. This metric is designed to assess the memory efficiency of the AnSAN technique.

- **Inference Acceleration Ratio (T_{\uparrow}).** Similar to Wang et al. (2023), capitalizing on the cached keys/values, we present the Inference acceleration ratio, which serves as an indicator of the inference efficiency of the AnSAN technique.

Note that we first report full attention inference results for all models, then present results with the AnSAN method (+AnSAN) applied, compressing sequence information into anchor tokens.

5 Results

As evident from the results presented in Table 1, both the AnLLM-AC and AnLLM-EP models demonstrate promising accuracy, comparable to that of the base model, while simultaneously improving memory and inference efficiency.

Accuracy (Acc). The proposed AnLLM-EP and AnLLM-AC models exhibit commendable accuracy across various benchmarks.

In the zero-shot setting, with full attention, AnLLM-EP and AnLLM-AC achieve average accuracies of 64.6% and 65.1%, respectively, comparable to Llama2-7B’s 65.8% accuracy. This suggests that training with integrated anchor tokens barely affects the model capacity, emphasizing the robustness of LLMs. Furthermore, our models excel in OBQA, PIQA, and SCIQ tasks.

In the five-shot setting, with five prior examples, AnLLM-EP and AnLLM-AC maintain dependable performance using full attention. When implementing the AnSAN technique, a slight accuracy decline across all models is observed. This is expected, as AnSAN, designed for memory efficiency, necessitates token removal, potentially leading to information loss. The degradation in BoolQ is most pronounced, which contains the longest demonstration tasks, indicating that the longer the text, the greater the information loss after compression. However, the average accuracy reduction is minimal, approximately 1.5%, suggesting that AnSAN effectively balances memory-saving and model performance.

Keys/Values Cache Reduction (C_{\downarrow}). The size of the keys/values cache is a critical factor in the practical implementation of LLMs, particularly concerning memory efficiency and computational resources. In this respect, the AnLLM-EP and AnLLM-AC models offer significant advantages.

	OBQA	WG	ARC-e	ARC-c	PIQA	HS	SCIQ	BoolQ	AVG.
Llama2-7B	31.4	69.1	76.3	43.4	78.1	57.1	93.7	77.7	65.8
AnLLM-EP	33.2	68.0	73.4	40.8	77.8	55.0	94.4	74.4	64.6
AnLLM-AC	31.6	68.5	74.4	42.5	78.3	54.7	93.8	77.0	65.1

(a) The Zero-Shot Performance.

		OBQA	WG	ARC-e	ARC-c	PIQA	HS	SCIQ	BoolQ	AVG.
	L_d	89	133	145	209	262	426	603	804	334
	L_x	18	26	36	42	42	90	130	169	69
Llama2-7B	<i>Acc</i>	37.2	73.7	79.8	50.0	78.7	58.3	96.8	78.4	69.1
+AnSAN	<i>Acc</i>	34.6	68.6	62.6	35.8	68.3	30.8	65.7	50.8	52.1
AnLLM-EP	<i>Acc</i>	36.8	71.0	79.4	49.4	78.1	55.3	96.6	75.6	67.8
+AnSAN	<i>Acc</i>	36.2	68.0	76.7	45.6	78.2	52.6	93.1	74.0	65.6
	L_{kv}	89	8	5	30	9	25	50	43	32
	\mathcal{C}_{\downarrow}	-0%	-94%	-97%	-86%	-97%	-94%	-92%	-95%	-90%
	\mathcal{T}_{\uparrow}	$\times 1.0$	$\times 1.0$	$\times 1.0$	$\times 1.2$	$\times 1.4$	$\times 2.1$	$\times 2.6$	$\times 3.5$	$\times 1.7$
AnLLM-AC	<i>Acc</i>	37.2	72.3	79.8	49.0	78.6	56.9	96.8	77.5	68.5
+AnSAN	<i>Acc</i>	35.6	70.6	79.2	47.9	78.7	55.6	95.7	76.6	67.5
	L_{kv}	5	5	5	5	5	5	5	5	5
	\mathcal{C}_{\downarrow}	-94%	-96%	-97%	-98%	-98%	-99%	-99%	-99%	-99%
	\mathcal{T}_{\uparrow}	$\times 1.0$	$\times 1.0$	$\times 1.1$	$\times 1.2$	$\times 1.5$	$\times 2.0$	$\times 2.6$	$\times 3.5$	$\times 1.7$

(b) The Five-Shot Performance.

Table 1: Accuracy and Efficiency of LLMs on Question Answering Benchmarks. \mathcal{C}_{\downarrow} represents the reduction in keys/values cache size, while \mathcal{T}_{\uparrow} denotes the inference acceleration ratio during testing. *Acc* stands for Accuracy. L_{kv} represents the length of the keys/values cache. L_d and L_x denote the lengths of in-context learning demonstrations and input queries, respectively. Our methods effectively reduce cache sizes and boost inference efficiency.

By adopting the AnSAN, these models are designed to dramatically reduce the keys/values cache size during inference. As shown in Table 1, these models achieve remarkable reductions in cache size. Specifically, the average reduction percentages are around 90% for AnLLM-EP and an impressive 99% for AnLLM-AC. This is a substantial improvement compared to conventional approaches, which typically necessitate large cache sizes to store keys/values. These reductions in cache size translate to considerable savings in memory and computational resources, rendering these models highly efficient for practical applications.

Inference Acceleration Ratio (\mathcal{T}_{\uparrow}). The inference acceleration ratio serves as a crucial metric reflecting the model’s efficiency during the testing phase. By incorporating anchor tokens into natural language texts, we can repurpose the hidden states of anchor tokens as keys/values caches in the demonstrations, and then adopt an inference strategy as suggested by Wang et al. (2023). In this

scenario, both the AnLLM-EP and AnLLM-AC models demonstrate significant improvements.

Specifically, in the five-shot testing, both AnLLM-EP and AnLLM-AC models attain an average inference acceleration ratio of approximately 1.7 times. This represents a considerable advancement over the conventional non-caching method, which typically necessitates prolonged processing times due to the large number of tokens involved. As L_d increases, reaching up to 3.5 times in the BoolQ task, the acceleration ratios also escalate, corroborating the findings of Wang et al. (2023). This enhancement in processing speed leads to increased efficiency, making these models particularly apt for scenarios with limited resources.

The AnLLM-EP and AnLLM-AC models exhibit remarkable performance in natural language understanding benchmarks, effectively balancing accuracy, memory efficiency, and time acceleration. The incorporation of anchor tokens into AnLLMs, along with the utilization

of the AnSAN technique for reducing keys/values cache size, allows these models to maintain performance on par while significantly improving memory efficiency and inference speed. The equilibrium achieved between model performance and computational efficiency is noteworthy and opens up new possibilities for the advancement of LLMs.

6 Analysis

To further elucidate our method’s insights, we conduct a natural language generation experiment with the German-to-English (De2En) translation task. We evaluate the models using COMET-DA (Rei et al., 2022), indicating translation quality, and the Keys/Values Cache Reduction C_{\downarrow} metric, denoting memory efficiency as previously described. In line with previous findings, AnLLMs accept a minor accuracy trade-off (about 3 COMET-DA points) for enhanced memory efficiency. All LLMs are fine-tuned on the Alpaca dataset, combined with the newstest2017-2020 datasets, following Jiao et al. (2023). Results are presented in Table 2.

6.1 Compatibility and Flexibility of Full Attention and Anchor-based Attention

The results offer significant insight into the interplay between anchor-based attention and full attention mechanisms in the De2En translation task. Since source sentences are vital in translation tasks, applying full attention to them is crucial for maintaining model performance. Thus, retaining the source sentence keys/values caches is expected to enhance AnLLM performance when implementing the AnSAN technique. Specifically, when combining full attention with the AnSAN method, both AnLLM-EP and AnLLM-AC achieve approximately 80.0 COMET-DAE scores, comparable to other models using full attention exclusively. This indicates that the AnSAN technique is compatible with the full attention mechanism. Consequently, our proposed models allow users to choose between full attention and anchor-based attention for input texts based on their needs, emphasizing the compatibility and flexibility of our models.

6.2 Effective Cache Reduction for Real-Time Inference with the AnSAN Technique

The results in Table 2 show that our reduction strategy effectively minimizes keys/values caches during real-time inference. In Specific, as indicated in Line 15 of Algorithm 1, when generating an anchor token (i.e., the endpoint or <AC> tokens),

Model	Src Cache	De2En	MaxKV	C_{\downarrow}
Llama2-7b	✓	83.1	220	0%
AnLLM-EP	✓	81.6	220	0%
+AnSAN	✗	78.5	50	77%
	✓	80.3	124	44%
AnLLM-AC	✓	82.4	220	0%
+AnSAN	✗	78.0	35	84%
	✓	80.0	125	43%

Table 2: COMET-DA Scores and Keys/Values Cahces for the WMT23 German-to-English (De2En) Translation Task. The term “Src Cach” denotes retaining source sentence hidden states in Keys/Values Caches, while “MaxKV” refers to the average maximum keys/values length during inference.

our AnSAN-equipped models execute the reduction function to minimize the current keys/values caches. When discarding source sentence caches, we achieve approximately 77% and 84% reduction for the AnLLM-EP and AnLLM-AC models, respectively, albeit with a low COMET-DA score. However, when retaining source sentence caches, we still reduce around 44% of caches for both models, achieving a COMET-DA score of approximately 80.0. These results confirm the effectiveness of our anchor-based inference strategy for practical real-time inference applications.

7 Conclusion

LLMs have emerged as a significant research area in the field of artificial intelligence. However, despite their exceptional performance across various natural language tasks, the practical application of these models is limited by their significant memory overhead and time efficiency. Implementing LLMs on resource-constrained devices, such as smartphones, poses a unique challenge. To address this issue, we propose anchor-based LLMs with the AnSAN technique. Our experiments demonstrate that by sacrificing a marginal 1.5% in precision, our approach saves 99% of keys/values cache memory while simultaneously improving inference speed by up to 3.5 times. Our methods’ application in machine translation showcases their compatibility and flexibility, effectively enhancing memory efficiency for practical use. Our novel approach is practical, straightforward, flexible, and compatible with existing methods, paving the way for further adoption of LLMs in real-world applications.

578 Limitations

579 While our proposed AnLLMs demonstrate signif- 627
580 icant improvements in memory efficiency and infer- 628
581 ence acceleration, there are several limitations 629
582 to consider: 630

- 583 1. **Accuracy trade-off:** As observed in the ex- 631
584 perimental results, our method incurs a minor 632
585 decrease in accuracy (within 1.5%) compared 633
586 to the original model. This limitation 634
587 stems from the information compression pro- 635
588 cess, which may lead to information loss. Al- 636
589 though the degradation in accuracy is rela- 637
590 tively small, it is crucial to acknowledge this 638
591 trade-off when deploying our method in prac- 639
592 tical applications. 640
- 593 2. **Applicability to various tasks:** Our experi- 642
594 ments primarily focus on question-answering 643
595 benchmarks and machine translation tasks. 644
596 The effectiveness of our method in other natu- 645
597 ral language processing tasks, such as summa- 646
598 rization, sentiment analysis, and entity recog- 647
599 nition, remains to be thoroughly investigated. 648
600 Future work should explore the applicabil- 649
601 ity and performance of our method across a 650
602 broader range of tasks. 651
- 603 3. **Optimal anchor token selection:** In our im- 652
604 plementation, we chose the last token in a se- 653
605 quence as the anchor token. However, the op- 654
606 timal anchor token selection may vary across 655
607 different tasks and domains. Further research 656
608 is needed to develop more sophisticated strate- 657
609 gies for identifying and leveraging the most 658
610 suitable anchor tokens. 659
- 611 4. **Scalability to other LLMs:** We have applied 660
612 our method to the open-source Llama2 mod- 661
613 els. It remains to be seen how our approach 662
614 would perform when applied to other open- 663
615 source LLMs, such as Falcon and Qwen (Al- 664
616 mazrouei et al., 2023; Bai et al., 2023). Evalu- 665
617 ating the effectiveness and scalability of our 666
618 method on more extensive language models 667
619 is an essential direction for future research. 668

620 Despite these limitations, our work presents a 669
621 novel approach to enhance memory efficiency and 670
622 inference acceleration in LLMs. Future research 671
623 efforts should address these limitations, refining 672
624 our method and extending its applicability to a 673
625 wider range of tasks and model architectures. 674

Ethics Statement

626 In conducting this research, we have adhered to 627
628 the highest ethical standards and principles of aca- 629
630 demic integrity. The development and implemen- 630
631 tation of the AnLLMs and the AnSAN have been 631
632 carried out with the primary aim of improving the 632
633 memory efficiency and inference speed of large 633
634 language models, without any intention to cause 634
635 harm or promote malicious applications. 635

636 Our methodology and experimental design have 636
637 been thoroughly reviewed to ensure that the 637
638 datasets and models employed are used respon- 638
639 sibly and appropriately. The RedPajama datasets 639
640 and the open-source Llama2 models, which we 640
641 utilized in our study, are publicly available and 641
642 widely recognized as reliable resources in the re- 642
643 search community. All data used in this study have 643
644 been processed and analyzed in compliance with 644
645 relevant guidelines and best practices. 645

646 We acknowledge that the advancements in large 646
647 language models and their applications may have 647
648 potential implications for privacy, security, and fair- 648
649 ness. In light of these concerns, we emphasize the 649
650 importance of responsible usage and deployment 650
651 of our proposed AnLLMs and AnSAN techniques. 651
652 Researchers and practitioners adopting our meth- 652
653 ods should be aware of the potential risks and take 653
654 necessary precautions to mitigate any unintended 654
655 consequences. 655

656 Throughout this study, we have strived for trans- 656
657 parency and reproducibility. Our results and find- 657
658 ings are reported honestly and accurately, without 658
659 any manipulation or misrepresentation. We are 659
660 committed to sharing our knowledge and insights 660
661 with the broader research community, and we en- 661
662 courage open discussion and constructive feedback 662
663 to further advance the understanding and develop- 663
664 ment of efficient and ethical large language models. 664

665 In conclusion, this research has been conducted 665
666 in accordance with the highest ethical standards, 666
667 and we are dedicated to fostering a responsible and 667
668 collaborative research environment in the field of 668
669 large language models and artificial intelligence. 669

References

- 670 Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Al- 670
671 shamsi, Alessandro Cappelli, Ruxandra Cojocaru, 671
672 Mérouane Debbah, Étienne Goffinet, Daniel Hess- 672
673 low, Julien Launay, Quentin Malartic, et al. 2023. 673
674 The falcon series of open language models. *arXiv* 674
675 *preprint arXiv:2311.16867*. 675

676	Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,	Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong	729
677	Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei	Wu, Baobao Chang, Xu Sun, Jingjing Xu, and	730
678	Huang, et al. 2023. Qwen technical report. <i>arXiv</i>	Zhifang Sui. 2022. A survey for in-context learning.	731
679	<i>preprint arXiv:2309.16609</i> .	<i>arXiv preprint arXiv:2301.00234</i> .	732
680	Lochan Basyal and Mihir Sanghvi. 2023. Text	Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman,	733
681	summarization using large language models: A	Sid Black, Anthony DiPofi, Charles Foster, Laurence	734
682	comparative study of mpt-7b-instruct, falcon-7b-	Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li,	735
683	instruct, and openai chat-gpt models. <i>arXiv preprint</i>	Kyle McDonell, Niklas Muennighoff, Chris Ociepa,	736
684	<i>arXiv:2310.10449</i> .	Jason Phang, Laria Reynolds, Hailey Schoelkopf,	737
685	Petr Baudiš, Silvestr Stanko, and Jan Šedivý. 2016.	Aviya Skowron, Lintang Sutawika, Eric Tang, Anish	738
686	Joint learning of sentence embeddings for relevance	Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023.	739
687	and entailment. In <i>Proceedings of the 1st Workshop</i>	A framework for few-shot language model	740
688	<i>on Representation Learning for NLP</i> , pages 8–17,	evaluation.	741
689	Berlin, Germany. Association for Computational Lin-	Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng	742
690	guistics.	Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023.	743
691	Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi,	Longllmlingua: Accelerating and enhancing llms	744
692	et al. 2020. Piqa: Reasoning about physical commonsense	in long context scenarios via prompt compression.	745
693	in natural language. In <i>Proceedings of the</i>	<i>arXiv preprint arXiv:2310.06839</i> .	746
694	<i>AAAI conference on artificial intelligence</i> , volume 34,	Wenxiang Jiao, Jen-tse Huang, Wenxuan Wang, Zhiwei	747
695	pages 7432–7439.	He, Tian Liang, Xing Wang, Shuming Shi, and	748
696	Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai,	Zhaopeng Tu. 2023. Parrot: Translating during chat	749
697	Zhijian Liu, Song Han, and Jiaya Jia. 2023. Longlora:	using large language models tuned with human trans-	750
698	Efficient fine-tuning of long-context large language	lation and feedback. In <i>Findings of the Association</i>	751
699	models. <i>arXiv:2309.12307</i> .	<i>for Computational Linguistics: EMNLP 2023</i> , pages	752
700	Rewon Child, Scott Gray, Alec Radford, and	15009–15020, Singapore. Association for Computa-	753
701	Ilya Sutskever. 2019. Generating long sequences	tional Linguistics.	754
702	with sparse transformers. <i>arXiv preprint</i>	Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying	755
703	<i>arXiv:1904.10509</i> .	Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.	756
704	Christopher Clark, Kenton Lee, Ming-Wei Chang,	Gonzalez, Hao Zhang, and Ion Stoica. 2023. Effi-	757
705	Tom Kwiatkowski, Michael Collins, and Kristina	cient memory management for large language model	758
706	Toutanova. 2019. BoolQ: Exploring the surprising	serving with pagedattention. In <i>Proceedings of the</i>	759
707	difficulty of natural yes/no questions. In <i>Proceedings</i>	<i>ACM SIGOPS 29th Symposium on Operating Systems</i>	760
708	<i>of the 2019 Conference of the North American Chap-</i>	<i>Principles</i> .	761
709	<i>ter of the Association for Computational Linguistics:</i>	Ilya Loshchilov and Frank Hutter. 2019. Decoupled	762
710	<i>Human Language Technologies, Volume 1 (Long and</i>	weight decay regularization. In <i>International Confer-</i>	763
711	<i>Short Papers)</i> , pages 2924–2936, Minneapolis, Min-	ence on Learning Representations.	764
712	nesota. Association for Computational Linguistics.	Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish	765
713	Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot,	Sabharwal. 2018. Can a suit of armor conduct elec-	766
714	Ashish Sabharwal, Carissa Schoenick, and Oyvind	tricity? a new dataset for open book question answer-	767
715	Tafford. 2018. Think you have solved question	ing. <i>arXiv preprint arXiv:1809.02789</i> .	768
716	answering? try arc, the ai2 reasoning challenge. <i>arXiv</i>	Jesse Mu, Xiang Lisa Li, and Noah Goodman. 2023.	769
717	<i>preprint arXiv:1803.05457</i> .	Learning to compress prompts with gist tokens.	770
718	Together Computer. 2023. Redpajama: an open dataset	<i>arXiv preprint arXiv:2304.08467</i> .	771
719	for training large language models.	OpenAI. 2023. Gpt-4 technical report.	772
720	Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra,	Jianhui Pang, Fanghua Ye, Longyue Wang, Dian Yu,	773
721	and Christopher Ré. 2022. FlashAttention: Fast and	Derek F Wong, Shuming Shi, and Zhaopeng Tu. 2024.	774
722	memory-efficient exact attention with IO-awareness.	Salute the classic: Revisiting challenges of machine	775
723	In <i>Advances in Neural Information Processing Sys-</i>	translation in the age of large language models. <i>arXiv</i>	776
724	<i>tems</i> .	<i>preprint arXiv:2401.08350</i> .	777
725	Jacob Devlin, Ming-Wei Chang, Kenton Lee, and	Alec Radford, Jeffrey Wu, Rewon Child, David Luan,	778
726	Kristina Toutanova. 2018. Bert: Pre-training of deep	Dario Amodei, Ilya Sutskever, et al. 2019. Language	779
727	bidirectional transformers for language understand-	models are unsupervised multitask learners. <i>OpenAI</i>	780
728	ing. <i>arXiv preprint arXiv:1810.04805</i> .	<i>blog</i> , 1(8):9.	781

782	Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. 2020. Compressive transformers for long-range sequence modelling . In <i>International Conference on Learning Representations</i> .	2022. Emergent abilities of large language models. <i>arXiv preprint arXiv:2206.07682</i> .	839
783			840
784		Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. Crowdsourcing multiple choice science questions . In <i>Proceedings of the 3rd Workshop on Noisy User-generated Text</i> , pages 94–106, Copenhagen, Denmark. Association for Computational Linguistics.	841
785			842
786			843
787	Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. Parallel context windows for large language models . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 6383–6402, Toronto, Canada. Association for Computational Linguistics.		844
788			845
789		Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? <i>arXiv preprint arXiv:1905.07830</i> .	846
790			847
791			848
792			849
793			
794			
795	Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. COMET-22: Unbabel-IST 2022 submission for the metrics shared task . In <i>Proceedings of the Seventh Conference on Machine Translation (WMT)</i> , pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.		
796			
797			
798			
799			
800			
801			
802			
803	Jon Saad-Falcon, Joe Barrow, Alexa Siu, Ani Nenkova, Ryan A Rossi, and Franck Dernoncourt. 2023. Pdf-triage: Question answering over long, structured documents. <i>arXiv preprint arXiv:2309.08872</i> .		
804			
805			
806			
807	Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. <i>Communications of the ACM</i> , 64(9):99–106.		
808			
809			
810			
811	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models . <i>arXiv preprint arXiv:2302.13971</i> .		
812			
813			
814			
815			
816			
817	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models . <i>arXiv preprint arXiv:2307.09288</i> .		
818			
819			
820			
821			
822			
823	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need . <i>Advances in neural information processing systems</i> , 30.		
824			
825			
826			
827			
828	Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. Label words are anchors: An information flow perspective for understanding in-context learning . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing</i> , pages 9840–9855, Singapore. Association for Computational Linguistics.		
829			
830			
831			
832			
833			
834			
835			
836	Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al.		
837			
838			

A More Experimental Results

A.1 Testing Acceleration Ratio to Full-Caching Method

In Section 5, we report the testing acceleration ratio following the setting of Wang et al. (2023), comparing the time difference between caching and non-caching inference. Although our method reduces the keys/values caches, enabling smaller space for prefix information storage and improving testing time up to $\times 3.5$, we are still curious about whether it would enhance inference efficiency if conventional methods use full caches that save all keys/values of prefix tokens. As a supplement to Table 1, we present the testing acceleration ratio between anchor-caching and full-caching inference in Table 3. The acceleration ratios for AnLLM-EP-AnSAN and AnLLM-AC-AnSAN achieve the highest improvements observed in tasks such as HS, SCIQ, and BoolQ. The average acceleration ratios for AnLLM-EP-AnSAN and AnLLM-AC-AnSAN are 1.03. These results demonstrate that our anchor-based caching method can enhance inference efficiency even when compared to conventional methods that save all keys/values of prefix tokens. These results suggest that our anchor-based caching approach, which saves only the keys/values caches of anchor tokens, can effectively accelerate the inference process for the lengthy prefix texts.

A.2 Model Scalability Assessment

To examine the scalability of our approach, we extend the AnLLM-AC model to 13B and assess its performance on eight question-answering benchmarks using the same evaluation strategy as previously mentioned. In comparison to the 7B AnLLM models in Table 1, Results in Table 4 indicate that as the model size expands, the AnLLM-AC model achieves accuracies of 67.5% and 70.0% for 0-shot and 5-shot testing, respectively, resulting in up to a 2.4% improvement. Moreover, by incorporating anchor-based attention, the AnLLM-AC-AnSAN model achieves an average accuracy of 69.5%, signifying a 2.0% increase. The performance enhancement underscores the effectiveness of our methods in accommodating larger model capacities. The consistent improvements observed in the AnLLM-AC model across various scenarios highlight its robustness and adaptability. Furthermore, the increased performance of the AnLLM-AC-AnSAN model, facilitated by anchor-based attention, emphasizes the potential of our approaches in opti-

mizing LLMs. Collectively, these findings point to promising avenues for future research aimed at maximizing the utility and efficiency of AnLLM.

A.3 Case Study in Real-Time Inference

To elaborate on the optimization of keys/values caches by AnLLM-EP and AnLLM-AC during real-time inference, we reference examples from the translation task in Section 6.2. As per Table 5, AnLLM-EP and AnLLM-AC use "endpoints" (".") and "<AC>" tokens as anchor tokens, respectively. During inference, both models employ auto-regressive generation, creating outputs token-by-token. Upon generating an anchor token (as per Line 16, Algorithm 1), the Reduction function (defined in Line 1) is activated, preserving relevant caches and eliminating others. As a result, the Keys/Values Cache lengths are reduced to roughly the sequence length, averaging around 50 for AnLLM-EP and 35 for AnLLM-AC, as shown in Table 2.

B Data Settings

To provide a thorough insight into how we continually pre-train the model into AnLLM and carry out evaluations, we showcase some data examples in this section for both training and testing data.

B.1 Training Data Examples

In this section, we provide examples to illustrate the specific data format used in training the AnLLM models. For the AnLLM-EP model, the endpoints act as anchor tokens, allowing us to directly utilize natural language texts. For the AnLLM-AC model, we append a new token <AC> at the end of each sequence in the input texts, which are initially split into sentences using the NLTK toolkits.³ Some examples are presented in Table 6. All the trainig data are downloaded from HuggingFace⁴, an open-source community.

B.2 Testing Data Examples

For the testing outlined in the results section (Section 5), we employ the same evaluation method as in previous work (Gao et al., 2023), which treats each choice as text generation and computes the corresponding probabilities, respectively. Table 7 presents some evaluation examples.

³<https://www.nltk.org/api/nltk.tokenize.punkt.html>

⁴<https://huggingface.co/datasets/togethercomputer/RedPajama-Data-1T-Sample>

	OBQA	WG	ARC-e	ARC-c	PIQA	HS	SCIQ	BoolQ	AVG.
AnLLM-EP-AnSAN	×1.00	×1.00	×1.00	×1.00	×1.00	×1.06	×1.14	×1.13	×1.03
AnLLM-AC-AnSAN	×1.00	×1.02	×1.00	×1.00	×1.00	×1.01	×1.10	×1.13	×1.03

Table 3: Testing Acceleration Ratio on Question-Answering Tasks between Anchor-Caching and Full-Caching Inference with Five-Shot Demonstrations. Anchor-caching refers to saving only the keys/values caches of anchor tokens with the AnSAN technique, while full-caching denotes saving caches for all prefix tokens. The tasks are arranged according to the demonstration lengths. The experiments are the same as those of Table 1. These results suggest that inference speed differences for short texts are minimal but become more pronounced for longer texts. However, full-caching inference demands more GPU memory to store the complete keys/values caches, which is not ideal for environments with limited computational resources.

Model	OBQA	WG	ARC-e	ARC-c	PIQA	HS	SCIQ	BoolQ	AVG.
<i>Zero-Shot Performance</i>									
Llama2-7B	31.4	69.1	76.3	43.4	78.1	57.1	93.7	77.7	65.8
Llama2-13b	35.2	72.1	79.4	48.5	79.1	60.0	94.5	80.6	68.7
AnLLM-AC-7B	31.6	68.5	74.4	42.5	78.3	54.7	93.8	77.0	65.1
AnLLM-AC-13B	35.2	70.7	77.9	46.9	78.6	58.1	94.7	78.1	67.5
<i>Five-Shot Performance</i>									
Llama2-7B	37.2	73.7	79.8	50.0	78.7	58.3	96.8	78.4	69.1
Llama2-13b	38.2	76.3	82.2	52.6	80.0	61.4	97.5	83.5	71.5
AnLLM-AC-7B	37.2	72.3	79.8	49.0	78.6	56.9	96.8	77.5	68.5
+AnSAN	35.6	70.6	79.2	47.9	78.7	55.6	95.7	76.6	67.5
AnLLM-AC-13B	36.6	72.5	81.6	53.7	79.2	59.6	97.5	79.6	70.0
+AnSAN	36.0	74.0	81.6	52.0	79.1	58.4	96.3	78.8	69.5

Table 4: Accuracy of 13B LLMs on Question Answering Benchmarks. Compared to 7B AnLLMs, the 13B AnLLMs exhibit superior performance, with up to 2.0 accuracy enhancements, suggesting that AnLLMs possess excellent scalability to larger model architectures.

Input	Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request. ---- ### Instruction: Translate the following sentences from German to English. ---- ### Input: Nachdem Werte in einen anderen Teil des Speichers eingeschrieben wurden, wird das CMOS RAM in der gleichen Weise wie das Communications RAM geprüft. Wurde der Test bestanden, werden alle Speicherstellen auf ihren früheren Wert eingestellt. LED-Wert: 00 0011 Wenn das Gerät mit dem o.a. Display hält, liegt ein Fehler vor. Prüfen Sie in diesem Falle U 85 und U 86 und die damit verbundenen Stromkreise bzw. die Dekodierung. ---- ### Response:
Output	After values have been written to another part of the CMOS RAM, the CMOS RAM is tested in the same way as the communications RAM. If the test is successful, all storage locations will be reset to their former value. LED value: 00 0011 If the device is displaying this value, there is a fault to be found. In this case, check U85 and U86 and the associated power circuits, as well as decoding.
COMET-DA	82.2
Length	293
MaxKV	170
C_{\downarrow}	42%

(a) An Example of the AnLLM-EP Model in De2En Translation Task.

Input	Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request. <AC> ---- ### Instruction: Translate the following sentences from German to English. <AC> ---- ### Input: Der Spielablauf wird ausgelöst, wenn der Druckschalter START betätigt wird. <AC> Der LED-Schalter bleibt erleuchtet und die Spielgeräusche werden hörbar. <AC> Wird der letzte Spieler-Panzer vom feindlichen Angreifer eliminiert, ist ein Spiel beendet. <AC> Das Spiel beginnt mit langsam fahrenden Panzern, die auf dem Bildschirm sichtbar werden. <AC> Die Position der jeweiligen feindlichen Panzer oder Raketen kann der Spieler dadurch bestimmen, dass er den eingblendeten Text im oberen linken Hinweisdisplay-Winkel seines Cockpits beachtet. <AC> Ein Hinweis kann z.B. <AC> lauten: GEGNER VON RECHTS. <AC> Zusätzlich wird auf dem weitreichenden Radarschirm, begleitet von einem Sonar-BLIP-Ton, ein Punkt sichtbar, der die feindliche Stellung darstellt. <AC> ---- ### Response:
Output	The game is triggered when the start switch is pressed. <AC> The LED switch remains illuminated and the game sounds can be heard. <AC> When the last enemy player's tank is destroyed, the game is over. <AC> The game begins with tanks moving slowly move across the screen. <AC> The player's tank's position can be determined by the player by paying attention to the respective enemy tank or missile icon in the upper left corner of his cockpit's display screen. <AC> An example of an instruction can be: ENEMY ON THE RIGHT. <AC> In addition, on the wide-range radar screen, accompanied by a sonar BLIP sound, a point is visible, which indicates the enemy position.
COMET-DA	81.8
Length	431
MaxKV	272
C_{\downarrow}	37%

(b) An Example of the AnLLM-AC Model in De2En Translatio Task.

Table 5: Case Study of Real-time Inference. During the inference process, AnLLM-EP generates "endpoint" as the anchor token, whereas AnLLM-AC produces "<AC>" as the anchor token. Once upon an anchor token, we execute the REDUCTION as shown in Line 16 to reduce the keys/values caches.

Gender diversity, or more often its lack thereof, among participants to software development activities has been thoroughly studied in recent years. In particular, the presence of, effects of, and countermeasures for gender bias in Free/Open Source Software (FOSS) have received a lot of attention over the past decade. Geographic diversity is on the other hand the kind of diversity that stems from participants in some global activity coming from different world regions and cultures. Geographic diversity in FOSS has received relatively little attention in scholarly works. In particular, while seminal survey-based and point-in-time medium-scale studies of the geographic origins of FOSS contributors exist, large-scale longitudinal studies of the geographic origin of FOSS contributors are still lacking. Such a quantitative characterization would be useful to inform decisions related to global development teams and hiring strategies in the information technology (IT) market, as well as contribute factual information to the debates on the economic impact and sociology of FOSS around the world. ...

(a) A Training Data Example for the AnLLM-EP Model. The endpoints in the text serve as the anchor tokens.

Gender diversity, or more often its lack thereof, among participants to software development activities has been thoroughly studied in recent years. <AC> In particular, the presence of, effects of, and countermeasures for gender bias in Free/Open Source Software (FOSS) have received a lot of attention over the past decade. <AC> Geographic diversity is on the other hand the kind of diversity that stems from participants in some global activity coming from different world regions and cultures.<AC> Geographic diversity in FOSS has received relatively little attention in scholarly works. <AC> In particular, while seminal survey-based and point-in-time medium-scale studies of the geographic origins of FOSS contributors exist, large-scale longitudinal studies of the geographic origin of FOSS contributors are still lacking. <AC> Such a quantitative characterization would be useful to inform decisions related to global development teams and hiring strategies in the information technology (IT) market, as well as contribute factual information to the debates on the economic impact and sociology of FOSS around the world. <AC> ...

(b) A Training Data Example for the AnLLM-AC Model. The newly added tokens <AC> in the text serve as the anchor tokens.

Table 6: Training Data Examples for the AnLLM-EP and AnLLM-AC models. For the AnLLM-EP model, the endpoints are the natural anchor tokens. For the AnLLM-AC model, we manually append <AC> tokens to sequences as the anchor tokens.

Choice 1: Slacklining: A group of people have stretched a tightrope across a gym. They *take turns trying to balance and walk on the rope.*

Choice 2: Slacklining: A group of people have stretched a tightrope across a gym. They *slide down with it, jumping and spinning in the air.*

Choice 3: Slacklining: A group of people have stretched a tightrope across a gym. They *cross it together, swinging back and fourth in anticipation.*

Choice 4: Slacklining: A group of people have stretched a tightrope across a gym. They *drop an orange rope at the end.*

(a) A Zero-Shot Testing Data Example of the HellaSwag Task. The log-likelihood of the red texts is computed as the choice probabilities.

Choice 1: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. Demonstration 2 Demonstration 3 Demonstration 4 Demonstration 5 Slacklining: A group of people have stretched a tightrope across a gym. They *take turns trying to balance and walk on the rope.*

Choice 2: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. Demonstration 2 Demonstration 3 Demonstration 4 Demonstration 5 Slacklining: A group of people have stretched a tightrope across a gym. They *slide down with it, jumping and spinning in the air.*

Choice 3: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. Demonstration 2 Demonstration 3 Demonstration 4 Demonstration 5 Slacklining: A group of people have stretched a tightrope across a gym. They *cross it together, swinging back and fourth in anticipation.*

Choice 4: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. Demonstration 2 Demonstration 3 Demonstration 4 Demonstration 5 Slacklining: A group of people have stretched a tightrope across a gym. They *drop an orange rope at the end.*

(b) A Five-Shot Testing Data Example of the HellaSwag Task for the ALLM-EP Inference.

Choice 1: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. <AC> Demonstration 2 <AC> Demonstration 3 <AC> Demonstration 4 <AC> Demonstration 5 <AC> Slacklining: A group of people have stretched a tightrope across a gym. They *take turns trying to balance and walk on the rope.*

Choice 2: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. <AC> Demonstration 2 <AC> Demonstration 3 <AC> Demonstration 4 <AC> Demonstration 5 <AC> Slacklining: A group of people have stretched a tightrope across a gym. They *slide down with it, jumping and spinning in the air.*

Choice 3: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. <AC> Demonstration 2 <AC> Demonstration 3 <AC> Demonstration 4 <AC> Demonstration 5 <AC> Slacklining: A group of people have stretched a tightrope across a gym. They *cross it together, swinging back and fourth in anticipation.*

Choice 4: Ballet: We see a pregnant lady doing ballet in a studio. The lady spins and does a pliea. <AC> Demonstration 2 <AC> Demonstration 3 <AC> Demonstration 4 <AC> Demonstration 5 <AC> Slacklining: A group of people have stretched a tightrope across a gym. They *drop an orange rope at the end.*

(c) A Five-Shot Testing Data Example of the HellaSwag Task for the ALLM-AC Inference.

Table 7: Testing Data Examples for the AnLLM-EP and AnLLM-AC models. The log-likelihood of the red italicized texts is calculated as the choice probabilities.