
IceFormer: Accelerated Inference with Long-Sequence Transformers on CPUs

Yuzhen Mao Martin Ester Ke Li
School of Computing Science
Simon Fraser University
{yuzhenm, ester, keli}@sfu.ca

Abstract

One limitation of existing transformer-based models is that they cannot handle very long sequences as input since their self-attention operations exhibit quadratic time and space complexity. This problem becomes especially acute when transformers are deployed on hardware platforms equipped only with CPUs. To address this issue, we propose a novel method for accelerating self-attention at inference time that works with pretrained transformer models out-of-the-box without requiring retraining. We experiment using our method to accelerate various long-sequence transformers on various benchmarks and demonstrate a greater speedup compared to the baselines.

1 Introduction

Transformers [25] have powered incredible advances in NLP, as exemplified by large language models (LLMs) such as GPT-4 and LLaMa 2. Increasingly LLMs are applied to exceptionally long input sequences, which enables many exciting applications such as long-form content creation, extended conversations, and large document search and analysis [19, 3]. While LLMs can be feasibly trained with expensive hardware accelerators (e.g. GPUs), they need to be deployed on a variety of devices, some of which are only equipped with CPUs.

However, it is currently challenging to deploy LLMs on CPUs due to their high computation cost [9]. A significant computational bottleneck arises from the self-attention mechanism that is integral to Transformers – both time and space complexity are quadratic in the sequence length. This problem is exacerbated in the context of LLMs, which are often used on very long sequences.

To handle long input sequences, there has been substantial research into reducing the quadratic time complexity of self-attention – these methods are collectively known as *efficient Transformers*. However, many do not meet the needs of LLMs and are therefore difficult to apply to LLMs.

An ideal acceleration method for LLMs should meet four criteria: (1) **No retraining** – the method should not require the model retraining since LLMs are very expensive and time-consuming to train; (2) **Generality** – the method can be applied to a broad range of LLMs rather than those trained with particular constraints built-in; (3) **High accuracy** – the method should not introduce large approximation errors, since LLMs have many attention layers and so errors from earlier layers can compound; (4) **Fast inference** – the method should enable fast inference with LLMs.

Meeting all these criteria simultaneously is difficult, and to our knowledge no existing methods can do so. For example, Transformers with fixed attention patterns (e.g., Longformer [4]) also require retraining the model before they can be used. Reformer [18] requires keys to be normalized, which is not the case in most pretrained models. Nyströmformer [27] and LARA [30] do not support causal masks, which are common in LLMs. Low-rank methods such as Performer [6] have poor approximation accuracy without finetuning.

In this paper, we propose an acceleration method named **IceFormer** that simultaneously meets the above four criteria. In particular, (1) IceFormer does not require retraining; (2) it can be applied to most LLMs; (3) it can approximate vanilla attention accurately; (4) it achieves faster inference compared with other methods. We illustrate our method in Figure 1. We devise a method to approximate the attention mechanism with unnormalized keys with k -nearest neighbor search (k -NNS), and leverage a ranking-based method, Prioritized DCI, to identify important keys accurately and quickly.

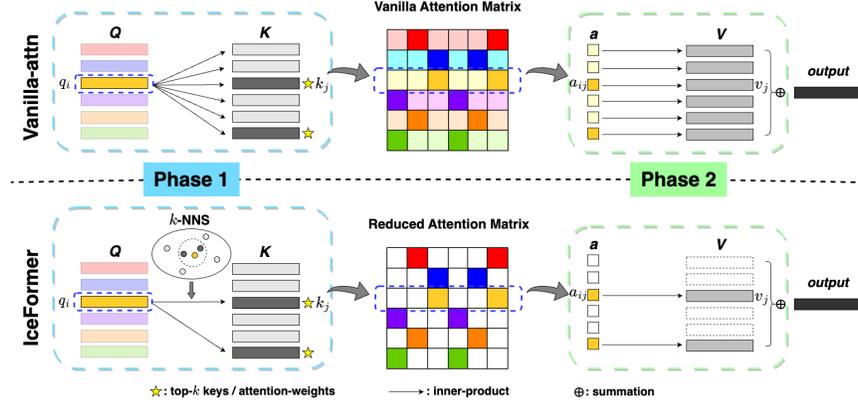


Figure 1: Overview of the vanilla attention [25] (top row) and the proposed method, IceFormer (bottom row). We illustrate with one query and $k = 2$ in k -NNS. In the two attention matrices presented, the top-2 largest attention weights in each row are represented by a dark color. The remaining attention weights are shown in a pale color in the vanilla attention matrix, and are set to zero (depicted in white) in the reduced attention matrix.

We conduct experiments on CPUs on the ZeroSCROLLS benchmark [22] as well as the LongEval benchmark [7]. In both of these benchmarks, our method demonstrates compelling performance during inference when provided with well-trained weights compared with other baseline models. Notably, our approach not only achieves accuracy close to that obtained using the vanilla Transformer, but also surpasses the efficiency of all baseline methods.

2 Notation and Preliminaries

Mathematically, the attention operation takes three matrices as input, $\mathbf{K} \in \mathbb{R}^{m \times d}$, $\mathbf{Q} \in \mathbb{R}^{n \times d}$, $\mathbf{V} \in \mathbb{R}^{m \times d'}$, which denote keys, queries and values respectively, and outputs a matrix $\mathbf{O} \in \mathbb{R}^{n \times d'}$. The i th rows of \mathbf{K} , \mathbf{Q} , \mathbf{V} and \mathbf{O} , denoted as \mathbf{k}_i , \mathbf{q}_i , \mathbf{v}_i and \mathbf{o}_i , represent the i th key, query, value and output respectively.

First the attention operation computes the attention weight matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$. Its entry in the i th row and j th column, denoted as $a_{i,j}$, is computed with the following formula: $a_{i,j} = \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{d}}\right)$. Then the attention operation combines the values with the attention weights in the following way: $\mathbf{o}_i = \sum_{j=1}^m a_{i,j} \mathbf{v}_j$.

Notably, from previous literature [18, 10], it is known that the attention matrix \mathbf{A} is sparse, i.e., in each row of \mathbf{A} , only a few attention weights have significant (large) values, while the majority of the remaining values are close to zero.

3 IceFormer: Accelerated Self-Attention for General Keys without Retraining

In order to build a general-purpose acceleration method that does not require retraining, our approach must not require modifications to the attention mechanism to change attention patterns or the introduction of new model parameters to capture regularities in the attention patterns. This precludes

popular strategies such as attention mechanisms with predefined sparse attention patterns (e.g., Child et al. [5], Beltagy et al. [4], Ainslie et al. [1]) and learned dimensionality reduction of keys and queries (e.g., Wang et al. [26], Choromanski et al. [6]). Consequently, it is difficult to design an acceleration method that exploits known regularities in the attention patterns without imposing the retraining requirement. We therefore aim to design an acceleration method that does not make assumptions on the existence of regularity in the attention patterns. In order to improve on the $O(mn)$ complexity of vanilla attention, we need to adaptively identify the most important keys (i.e., those that receive the highest attention weights) without computing all attention weights. This seems like a chicken-and-egg problem: how can we know which attention weights are highest without comparing them to all the other attention weights?

Remarkably, in the special case of normalized keys, as proposed in Nikita et al. [18], this can be done by leveraging k -nearest neighbour search (k -NNS) to identify the k most important keys for each query. This relies on the following mathematical fact, whose derivation is included in Sect. B.1 of the appendix: if $\|\mathbf{k}_j\|_2 = 1$ for all j , $\arg \max_j a_{i,j} = \arg \min_j \|\mathbf{q}_i - \mathbf{k}_j\|_2^2$. However, this fact only holds when all the keys have the same norm – it is not true when different keys differ in their norms. Intuitively, this is because the norms of keys can modulate the attention weights they receive, all else being equal. So if key A has a larger norm than key B, key A can receive a higher attention weight than key B even if key A is farther from the query than key B. As a result, naïvely applying k -NNS in the general case would fail to identify the most important keys.

In this paper, we develop an acceleration method that does not require retraining or impose any constraints on keys. It is both accurate and computationally efficient, and can also work with attention masks that are common in Transformers, such as causal masks. We will describe the details in the following subsections.

3.1 General Retraining-Free Accelerated Attention

Instead of applying k -NNS to the original keys directly, we will first embed the keys and queries into a higher dimensional space. Inspired by Neyshabur and Srebro [17], we choose the following key and query embedding functions, which we denote as $T_K : \mathbb{R}^d \rightarrow \mathbb{R}^{d+1}$ and $T_Q : \mathbb{R}^d \rightarrow \mathbb{R}^{d+1}$:

$$T_K(\mathbf{k}_j) = [\mathbf{k}_j/c \quad \sqrt{1 - \|\mathbf{k}_j\|_2^2/c^2}]^\top \tag{1}$$

$$T_Q(\mathbf{q}_i) = [\mathbf{q}_i/\|\mathbf{q}_i\|_2 \quad 0]^\top \tag{2}$$

where $c \geq \max_{j'} \|\mathbf{k}_{j'}\|_2$ is at least the maximum norm across all keys.

It turns out that the k most important keys can be identified by performing k -NNS on the key embeddings using the query embedding. We show the derivation in the appendix Sect. B.2.

3.2 Accurate k -NNS for Accelerated Attention

The problem of k -NNS is one of the most well studied problems in theoretical computer science. Many algorithms have been developed, and often significant speedups can be obtained by allowing for mistakes with some probability. Such algorithms are known as randomized algorithms. In the context of LLMs, the number of attention layers is typically high and so errors from earlier layers can compound. Therefore, it is essential for the k -NNS algorithm to achieve high accuracy. Choosing an appropriate k -NNS algorithm is therefore crucial.

Most k -NNS algorithms are bucketing-based, which places keys into discrete buckets and searches over buckets that contain the query. On the other hand, ranking-based algorithms compares the rankings of different keys relative to the query and searches over highly ranked keys. A bucketing-based algorithm effectively uses a fixed threshold on similarity, and so a variable number (including zero) of keys can meet the threshold; on the other hand, a ranking-based algorithm returns a fixed number of keys, which effectively amounts to choosing a variable threshold on similarity based on the distribution of keys, as shown in Figure 3 in the appendix Sect. C. An example of a bucketing-based algorithm is locality-sensitive hashing (LSH) [11], and an example of a ranking-based algorithm is Prioritized DCI [13]. As shown in Figure 3, LSH hashes each key into a bucket associated with the hash value, whereas Prioritized DCI ranks keys along random directions.

For accelerating attention, we posit that ranking-based algorithms are better suited than bucketing-based algorithms, because attention weights depend on how different keys compare to one another,

rather than an absolute evaluation of each key against a fixed threshold. Therefore, ranking-based algorithms is better aligned with how attention weights are calculated and so can yield better recall of truly important keys.

In a Transformer, the keys in an attention layer depend on the output from the preceding attention layer, which means a database needs to be constructed for each attention layer. Therefore, it is important to choose a k -NN algorithm that attains both fast construction and querying. Moreover, in the context of LLMs, many popular models use decoder-only architectures. The attention layers in such architectures use causal masks to prevent the currently generated token to depend on future yet-to-be-generated tokens. Such masked attention is equivalent to excluding the masked out keys from the set of keys the k -NNS algorithm operates over. So each time a token is generated, one key becomes unmasked. Instead of constructing a new database each time a token is generated, it is more efficient to add keys incrementally to the database for k -NNS.

Fortunately, Prioritized DCI is efficient at both the construction and querying stages. If the number of random projection directions p is nearly as large as the intrinsic dimensionality of the data $d' \geq 1$ and the number of nearest neighbours k to look for is small, Prioritized DCI can return the exact k -nearest neighbours for a query with high probability within approximately $\tilde{O}(dk^{p/\bar{d}}m^{1-p/\bar{d}})$ time, where $\tilde{O}(\cdot)$ suppresses log factors. Its preprocessing is lightweight, and so only needs $O(dpm)$ time. If we compare this to the computational complexity of vanilla attention of $O(dmn)$, observe that there is no longer a term that depends on mn , and so there is no longer the quadratic dependence on sequence length. Later in the appendix section C, we also empirically validate the efficiency of Prioritized DCI and found it to be faster than nine other leading k -NNS algorithms.

To support causal masking, we extended the implementation of Prioritized DCI to support incremental updates to the database. This can be done efficiently, since the data structure consists of sorted lists, so insertions and deletions can be performed in $O(\log m)$ time if they are implemented as binary search trees.

4 Experiments

In this section, we analyze the performance of IceFormer in the LLM setting with long prompts as input using the ZeroSCROLLS benchmark [22] and LongEval benchmark [7]. Due to the space limit, we only show LLM experiments, but we also conduct experiments on LRA benchmark [23], and put the results in the appendix Sect. D.

We utilize IceFormer to accelerate the prompt processing process in LLMs. We pick Vicuna-7b-v1.5-16k [29], which is fine-tuned from LLaMa 2 [24] and is one of the top-performing open-source LLMs with a context length up to 16K tokens, for the following experiment. For details on the configuration, please check Sect. E.1 in the appendix.

ZeroSCROLLS Results. We compare IceFormer with vanilla Transformer and H-Transformer-1D on the ZeroSCROLLS benchmark [22] which is specifically designed for LLMs and contains ten diverse natural language tasks that require understanding long input contexts. Each task has a different sequence length varying between 3k and 10k. We measure ZeroSCROLLS scores and latency of the attention module. Table 1 shows that IceFormer achieves up to $2.7\times$ speed-up compared to standard self-attention while attaining at least 99.0% of the vanilla unaccelerated model performance.

LongEval Results & Scalability Analysis. To provide a more comprehensive analysis of IceFormer’s scalability in the LLM setting, we conducted additional experiments on the LongEval benchmark [7], which is designed to measure long-context performance and consists of two tasks: topic retrieval task with prompt length varying from 3k to 16k, and line retrieval task with prompt length varying from 5k to 16k. In Figure 2, we present the averaged latency of the attention module corresponding to different input prompt length as well as the inference accuracy using the vanilla Transformer and IceFormer.

From Figure 2, IceFormer can achieve nearly identical inference accuracy compared with the vanilla Vicuna-7b-v1.5-16k model. As the prompt length increases, there is a corresponding increase in the inference latency for both methods and for both tasks. However, even with very long prompt lengths, IceFormer maintains its scalability and consistently outperforms the vanilla Transformer. Furthermore, as the length of the prompt increases, the difference in the latency (wall clock time) between IceFormer and the vanilla Transformer becomes larger, demonstrating the superior scalability and efficiency of IceFormer in the context of LLMs.

Table 1: The performance of vanilla Transformer, H-Transformer-1D, and IceFormer on the ZeroSCROLLS benchmarks. Numbers in parentheses indicate the relative comparison to the vanilla unaccelerated model, denoted as Transformer. We employ the same abbreviations for metric and task names as specified in the original paper [22]. We refer interested readers to the original paper for the details.

Task (#tokens)	Metric	Transformer	H-Transformer-1D	IceFormer
GvRp (8k)	$R_{geo} \uparrow$	11.0 (100%)	6.8 (61.8%)	11.0 (100%)
	Time (s)	5.39 (1.0 \times)	4.52 (1.2 \times)	2.24 (2.4\times)
SSFD (8k)	$R_{geo} \uparrow$	13.5 (100%)	6.3 (46.7%)	13.5 (100%)
	Time (s)	5.75 (1.0 \times)	4.98 (1.2 \times)	2.14 (2.7\times)
QMsm (9k)	$R_{geo} \uparrow$	16.9 (100%)	10.7 (63.3%)	16.8 (99.4%)
	Time (s)	7.11 (1.0 \times)	5.17 (1.4 \times)	2.67 (2.7\times)
SQAL (8k)	$R_{geo} \uparrow$	18.9 (100%)	7.3 (38.6%)	18.9 (100%)
	Time (s)	5.12 (1.0 \times)	2.33 (2.2 \times)	2.15 (2.4\times)
Qspr (5k)	F1 \uparrow	34.2 (100%)	6.2 (18.1%)	34.0 (99.4%)
	Time (s)	2.49 (1.0 \times)	2.13 (1.2 \times)	1.06 (2.3\times)
Nrtv (10k)	F1 \uparrow	14.7 (100%)	2.0 (13.6%)	14.7 (100%)
	Time (s)	7.64 (1.0 \times)	5.02 (1.5 \times)	3.39 (2.3\times)
QALT (7k)	AC \uparrow	48.8 (100%)	6.8 (13.9%)	48.6 (99.6%)
	Time (s)	4.17 (1.0 \times)	2.31 (1.8 \times)	1.85 (2.3\times)
MuSQ (3k)	F1 \uparrow	18.6 (100%)	16.9 (90.9%)	18.5 (99.5%)
	Time (s)	0.70 (1.0 \times)	0.66 (1.1 \times)	0.49 (1.4\times)
SpDg (7.5k)	ES \uparrow	42.5 (100%)	2.9 (6.8%)	42.3 (99.5%)
	Time (s)	4.72 (1.0 \times)	2.36 (2.0 \times)	2.09 (2.3\times)
BkSS (7.5k)	$C_{idx} \uparrow$	19.5 (100%)	11.7 (60.0%)	19.3 (99.0%)
	Time (s)	4.77 (1.0 \times)	2.41 (2.0 \times)	1.96 (2.4\times)
Avg. (7.5k)	$f \uparrow$	23.9 (100%)	7.8 (32.5%)	23.8 (99.6%)
	Time (s)	4.79 (1.0 \times)	3.19 (1.5 \times)	2.00 (2.4\times)

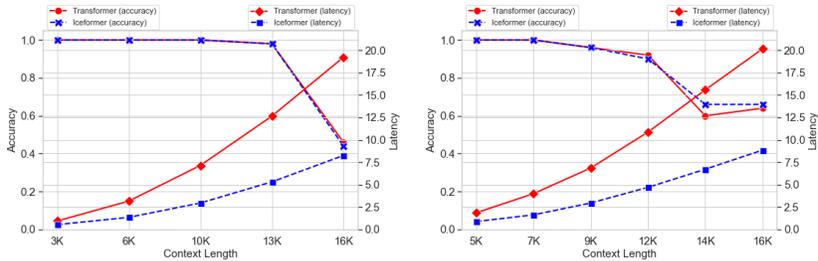


Figure 2: Scalability analysis for IceFormer on the LongEval benchmark. The left figure shows the results of the topic retrieval task; the right figure shows the results of the line retrieval task. X-axis: length of the input prompt; Y-axis (Left): retrieval accuracy; Y-axis (Right): averaged process wall clock time (second) of the attention module.

LongEval Results & Scalability Analysis. To provide a more comprehensive analysis of IceFormer’s scalability in the LLM setting, we conducted additional experiments on the LongEval benchmark [7], which is designed to measure long-context performance and consists of two tasks: topic retrieval task with prompt length varying from 3k to 16k, and line retrieval task with prompt length varying from 5k to 16k. In Figure 2, we present the averaged latency of the attention module corresponding to different input prompt length as well as the inference accuracy using the vanilla Transformer and IceFormer. From the figure, IceFormer can achieve nearly identical inference accuracy compared with the vanilla Vicuna-7b-v1.5-16k model. As the prompt length increases, there is a corresponding increase in the inference latency for both methods and for both tasks. However, even with very long prompt lengths, IceFormer maintains its scalability and consistently outperforms the vanilla Transformer. Furthermore, as the length of the prompt increases, the difference in the latency between IceFormer and the vanilla Transformer becomes larger, demonstrating the superior scalability and efficiency of IceFormer in the context of LLMs.

5 Conclusion

In this paper, we present IceFormer, a new method for improving the inference time efficiency of pre-trained Transformers on the CPU. The experimental findings compellingly illustrate the effectiveness of IceFormer in reducing the quadratic time and space complexity of vanilla Transformers.

References

- [1] Joshua Ainslie, Santiago Ontanon, Chris Alberti, Vaclav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. Etc: Encoding long and structured inputs in transformers. *arXiv preprint arXiv:2004.08483*, 2020.
- [2] Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt. Practical and optimal lsh for angular distance. *Advances in neural information processing systems*, 28, 2015.
- [3] Anthropic. 100k context windows, 2023. URL <https://www.anthropic.com/index/100k-context-windows>.
- [4] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [5] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [6] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020.
- [7] Li Dacheng, Shao Rulin, Xie Anze, Sheng Ying, Zheng Lianmin, Gonzalez Joseph E., Stoica Ion, Ma Xuezhe, and Zhang Hao. How long can open-source llms truly promise on context length?, June 2023. URL <https://lmsys.org/blog/2023-06-29-longchat>.
- [8] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022.
- [9] Dave Dice and Alex Kogan. Optimizing inference performance of transformers on cpus. *arXiv preprint arXiv:2102.06621*, 2021.
- [10] Ankit Gupta, Guy Dar, Shaya Goodman, David Ciprut, and Jonathan Berant. Memory-efficient transformers via top- k attention. *arXiv preprint arXiv:2106.06899*, 2021.
- [11] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, 1998.
- [12] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [13] Ke Li and Jitendra Malik. Fast k-nearest neighbour search via prioritized dci. In *International conference on machine learning*, pages 2081–2090. PMLR, 2017.
- [14] Drew Linsley, Junkyung Kim, Vijay Veerabadrán, Charles Windolf, and Thomas Serre. Learning long-range spatial dependencies with horizontal gated recurrent units. *Advances in neural information processing systems*, 31, 2018.
- [15] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P11-1015>.
- [16] Nikita Nangia and Samuel R Bowman. Listops: A diagnostic dataset for latent tree learning. *arXiv preprint arXiv:1804.06028*, 2018.
- [17] Behnam Neyshabur and Nathan Srebro. On symmetric and asymmetric lshs for inner product search. In *International Conference on Machine Learning*, pages 1926–1934. PMLR, 2015.
- [18] Kitaev Nikita, Kaiser Lukasz, Levskaya Anselm, et al. Reformer: The efficient transformer. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.
- [19] OpenAI. Openai gpt-4, 2023. URL <https://openai.com/gpt-4>.
- [20] Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A Smith, and Lingpeng Kong. Random feature attention. *arXiv preprint arXiv:2103.02143*, 2021.
- [21] Dragomir R Radev, Pradeep Muthukrishnan, Vahed Qazvinian, and Amjad Abu-Jbara. The acl anthology network corpus. *Language Resources and Evaluation*, 47:919–944, 2013.

- [22] Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Berant, and Omer Levy. Zeroscrolls: A zero-shot benchmark for long text understanding, 2023.
- [23] Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. *arXiv preprint arXiv:2011.04006*, 2020.
- [24] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [26] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [27] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14138–14148, 2021.
- [28] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, et al. H₂o: Heavy-hitter oracle for efficient generative inference of large language models. *arXiv preprint arXiv:2306.14048*, 2023.
- [29] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023.
- [30] Lin Zheng, Chong Wang, and Lingpeng Kong. Linear complexity randomized self-attention mechanism. In *International Conference on Machine Learning*, pages 27011–27041. PMLR, 2022.
- [31] Zhenhai Zhu and Radu Soricut. H-transformer-1D: Fast one-dimensional hierarchical attention for sequences. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3801–3815, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.294. URL <https://aclanthology.org/2021.acl-long.294>.

A Related Work

Efficient transformers can be categorized along two axes: method type and compatibility with pretrained models. Along the first axis are sparsity-based methods and low-rank methods. Along the second axis are methods that can and cannot be applied to pretrained unmodified Transformer models without retraining.

Sparsity-based methods employ a sparsified attention mechanism to capture global information and integrate it with local attention results. Some approaches aim to improve the space complexity compared to the vanilla attention mechanism without improving the time complexity, e.g., top- k Attention [10]. Other approaches aim to improve both, e.g., Sparse Transformer [5], Longformer [4], and ETC [1]. A substantial limitation of these models is that the tokens that are attended to are predefined and remain static, which do not adapt to varying input sequences. Because the original attention operation is permitted to attend to any token, these models must be trained with their respective predefined constraints on tokens to be attended to. Reformer [18] can attend to different sets of tokens for different input sequences by using Locality Sensitive Hashing (LSH) [2] to group tokens into chunks and subsequently attending only to tokens within the same chunk as each query and adjacent chunks. However, Reformer imposes two constraints that are not in the original attention operation: keys must be normalized and queries and keys must be the same. Therefore, Reformer must be trained with these constraints built-in. As a result, these methods cannot be applied to pretrained, non-modified, models directly; instead, the models must be retrained with the required constraints before these methods can be used.

Low-rank methods approximate the attention weight matrix with a low-rank matrix to reduce the quadratic time and space complexity. Examples of low-rank transformers include Linformer [26] and Performer [6], which decompose the attention weight matrix into a product of tall and wide matrices consisting of learned linear features or random features of the keys and queries, respectively. However, these transformers typically introduce significant approximation errors because attention weight matrices produced by the original attention operation, especially in the case of long input sequences, typically have high rank. Consequently, models that use these approaches must be trained with low-rank approximations built-in, in order to learn to be robust to the associated approximation errors. As a result, these approaches cannot be applied to pretrained, non-modified, models directly; instead, the models must be retrained with the required approximations before these methods can be used. Other approaches provide more general methodologies that can leverage weights pretrained with standard transformers without retraining. These transformers accelerate the execution of the standard attention operation without altering the underlying architecture. Two examples are Nyströmformer [27] and LARA [30], which replace the softmax structure in the self-attention mechanism with the product of separately activated query and key matrices. Nyströmformer utilizes the Nyström method, while LARA combines randomized attention (RA) and random feature attentions (RFA) [20] to reconstruct the attention weight matrix. In another example, H-Transformer-1D [31] recursively divides the attention weight matrix into blocks and truncates the small singular values of each off-diagonal blocks. All these approaches leverage low-rank approximations, as opposed to sparsity.

Other works propose hardware-specific optimizations without aiming to improve the computational complexity. Examples include Flash-attention [8], which optimizes reads and writes between levels of GPU memory, and H2O [28], which dynamically retains a balance of recent and heavy hitters tokens by a KV cache eviction policy. These strategies are dependent on implementation and are specific to particular hardware platforms (e.g. GPU).

B Proofs

B.1 Proof 1

Here, we provide the full step-by-step derivation of the mathematical equivalence between conducting k -nearest neighbor search on normalized keys and identifying the keys that obtain the highest attention

weight.

$$\arg \max_j a_{i,j} = \arg \max_j \text{softmax}_j \left(\left\{ \frac{\mathbf{q}_i^\top \mathbf{k}_{j'}}{\sqrt{d}} \right\}_{j'=1}^m \right) \quad (3)$$

$$= \arg \max_j \frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{d}} \quad (4)$$

$$= \arg \min_j \|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j + 1 \quad (5)$$

Since $\|\mathbf{k}_{j'}\|_2 = 1$ for all j' , $\|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j + 1 = \|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j + \|\mathbf{k}_j\|_2^2 = \|\mathbf{q}_i - \mathbf{k}_j\|_2^2$,

$$\arg \max_j a_{i,j} = \arg \min_j \|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j + 1 \quad (6)$$

$$= \arg \min_j \|\mathbf{q}_i - \mathbf{k}_j\|_2^2 \quad (7)$$

B.2 Proof 2

Here, we provide the full step-by-step derivation of the result in 3.1 establishing the mathematical equivalence between conducting k -nearest neighbor search on transformed keys and identifying the keys that obtain the highest attention weight.

$$\arg \max_j a_{i,j} = \arg \max_j \text{softmax}_j \left(\left\{ \frac{\mathbf{q}_i^\top \mathbf{k}_{j'}}{\sqrt{d}} \right\}_{j'=1}^m \right) \quad (8)$$

$$= \arg \max_j \frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{d}} \quad (9)$$

$$= \arg \max_j \mathbf{q}_i^\top \mathbf{k}_j \quad (10)$$

$$= \arg \min_j -2\mathbf{q}_i^\top \mathbf{k}_j \quad (11)$$

$$= \arg \min_j 2 - 2\mathbf{q}_i^\top \mathbf{k}_j / c \|\mathbf{q}_i\|_2 \quad (12)$$

$$= \arg \min_j 1 - 2\mathbf{q}_i^\top \mathbf{k}_j / c \|\mathbf{q}_i\|_2 + 1 \quad (13)$$

$$= \arg \min_j \|\mathbf{q}_i\|_2^2 / \|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j / c \|\mathbf{q}_i\|_2 + \|\mathbf{k}_j\|_2^2 / c^2 + 1 - \|\mathbf{k}_j\|_2^2 / c^2 \quad (14)$$

$$= \arg \min_j \mathbf{q}_i^\top \mathbf{q}_i / \|\mathbf{q}_i\|_2^2 - 2\mathbf{q}_i^\top \mathbf{k}_j / c \|\mathbf{q}_i\|_2 + \mathbf{k}_j^\top \mathbf{k}_j / c^2 + 1 - \|\mathbf{k}_j\|_2^2 / c^2 \quad (15)$$

$$= \arg \min_j (\mathbf{q}_i / \|\mathbf{q}_i\|_2 - \mathbf{k}_j / c)^\top (\mathbf{q}_i / \|\mathbf{q}_i\|_2 - \mathbf{k}_j / c) + 1 - \|\mathbf{k}_j\|_2^2 / c^2 \quad (16)$$

$$= \arg \min_j \|\mathbf{q}_i / \|\mathbf{q}_i\|_2 - \mathbf{k}_j / c\|_2^2 + 1 - \|\mathbf{k}_j\|_2^2 / c^2 \quad (17)$$

$$= \arg \min_j \|\mathbf{q}_i / \|\mathbf{q}_i\|_2 - \mathbf{k}_j / c\|_2^2 + \left(0 - \sqrt{1 - \|\mathbf{k}_j\|_2^2 / c^2} \right)^2 \quad (18)$$

$$= \arg \min_j \|T_Q(\mathbf{q}_i) - T_K(\mathbf{k}_j)\|_2^2 \quad (19)$$

$$= \arg \min_j \|T_Q(\mathbf{q}_i) - T_K(\mathbf{k}_j)\|_2 \quad (20)$$

C Different k -NNS Algorithms Comparison

C.1 Fast k -NNS for Accelerated Attention

We compare the recall of true nearest neighbors and total construction and querying time of ten k -NNS algorithms including Prioritized-DCI on fashion-mnist-784 dataset in Figure 4. As shown, Prioritized-DCI can achieve the best recall-latency trade-off across all the tested algorithms, which illustrates its suitability in the setting when the construction and querying time are both important.

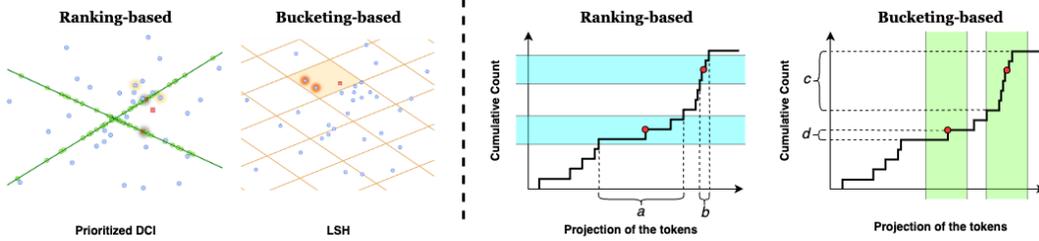


Figure 3: Difference between ranking-based and bucketing-based k -NNS.

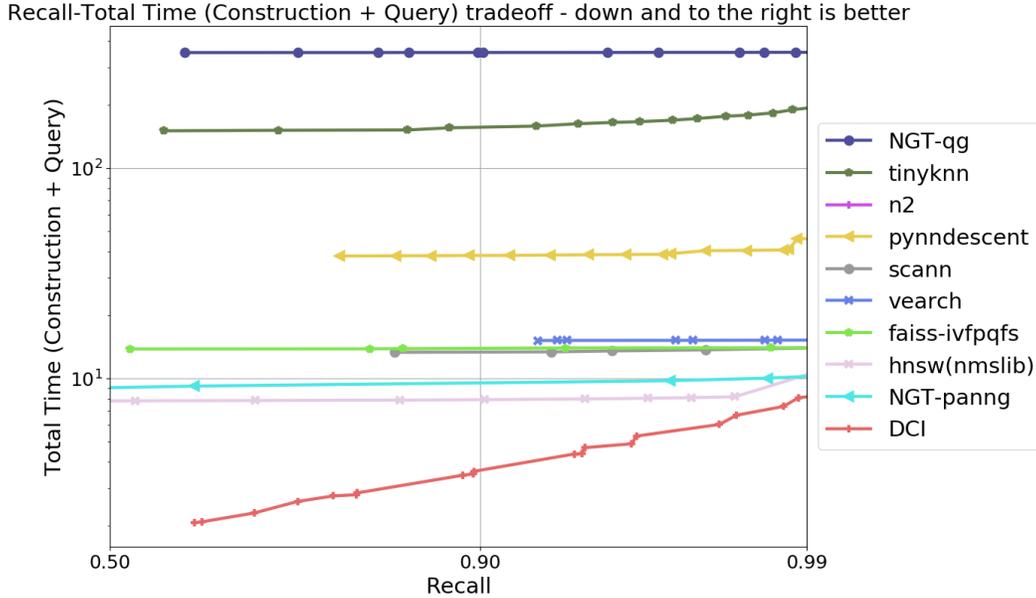


Figure 4: Comparison between ten k -NNS algorithms on fashion-mnist-784 dataset. There are in total 60,000 keys and 10,000 queries with 784 dimensions. The task is to find top-10 closest neighbors from the entire set of keys for every query. X-axis: Average recall across all the queries; Y-axis: Total latency (seconds) including database construction and querying. DCI denotes the Prioritized-DCI.

D Long Range Arena (LRA) Benchmark Experiments

In this section, we empirically analyze the performance of IceFormer on the LRA benchmark [23] which is under the long-context scenario and is evaluated in lots of previous literature [31, 27, 30].

D.1 Datasets and Metrics

LRA consists of five different tasks: ListOps [16], document retrieval (Retrieval) [21], text classification (Text) [15], CIFAR-10 image classification (Image) [12] and Pathfinder [14]. Specifically, all the five tasks consist of sequences with at most 4k tokens. We summarize the dataset information in the appendix for more details. In this experiment, we follow the train/test splits from Tay et al. [23] and report the test dataset classification accuracy, average running time of the attention module, and CPU memory usage during inference for each task.

In our LRA experiments, for Retrieval, Text and Pathfinder (64×64 version), we directly use the dataset from LRA codebase¹. Because the original datasets for ListOps and Image only contain short sequences, we generate longer samples for ListOps using the same code from the LRA codebase with 4000 as the maximum length; for Image task, we use a version of the CIFAR-10 dataset super-resolved

¹<https://github.com/google-research/long-range-arena/tree/main>

to 64×64^2 instead of the original low-resolution 32×32 CIFAR-10 dataset. We follow the exact same train/test split as the original LRA paper [23]. The details of the LRA dataset is listed in Table 2.

Table 2: LRA Dataset Details.

Task	ListOps	Text	Retrieval	Image	Pathfinder
Max length	3,991	4,000	4,000	4,096	4,096
Avg. length	2,232	1,267	3,917	4,096	4,096
number of classes	10	2	2	10	2
Accuracy by chance	0.100	0.500	0.500	0.100	0.500

D.2 Base Model Configuration

We follow the experimental setup of prior work [31] for training the base model. However, since we were not able to successfully train base Transformer models to satisfactory accuracy on Image and Pathfinder datasets using the original setting, we decreased the number of heads and layers for these two tasks. The details of the base model for each task are outlined in Table 3.

Table 3: Configurations of the base models for different tasks.

Task	ListOps	Text	Retrieval	Image	Pathfinder
head embedding size	512	512	512	512	512
feed-forward size	2048	2048	2048	1024	1024
number of heads	8	8	8	4	4
number of layers	6	6	6	4	4

D.3 Baselines.

In addition to the vanilla Transformer, we compare with Nyströmformer [27], H-Transformer-1D [31], LARA [30] and Reformer [18]. In order to compare with Reformer, we train a Transformer model with shared \mathbf{Q} and \mathbf{K} according to Nikita et al. [18]. For fair comparisons, we use the LRA evaluation benchmark implemented in PyTorch by [27], and only replace the self-attention module while making other parts of each model exactly the same as the vanilla Transformer.

D.4 Implementation Details.

For each task, we begin by training a base model using GPU with a vanilla Transformer architecture. Then we replace the vanilla attention module with one of the five efficient attention modules mentioned earlier and directly apply the pre-trained weights for inference. To ensure fair comparison, we adjust the batch size to 1, eliminating the need for a padding mask since our proposed IceFormer automatically ignores padding masks during inference. Note that because of the additional shared-KQ constraint, for the Pathfinder task, our attempts to train a shared-KQ Transformer were unsuccessful. As a result, we have excluded the corresponding results from the subsequent analysis. Additionally, during the inference, we utilize a total of 4 CPU threads. The machine used for LRA benchmark is Intel(R) Core(TM) i7-6850K.

D.5 Inference Results.

We conducted a comparison of IceFormer with five baseline methods on five LRA tasks. Ideally, the accuracy of the vanilla Transformer (non-shared-KQ) serves as an upper bound for the approximated accuracy of the other four models (IceFormer (non-shared-KQ), Nyströmformer, H-Transformer-1D, and LARA). Similar for the shared-KQ Transformer. Also, the attention module inference time of the vanilla Transformer would be the longest, with other efficient Transformers achieving shorter inference times at the cost of sacrificing prediction accuracy. Table 4 presents the prediction accuracy and inference time of the attention module for each method. The hyper-parameter settings are listed in the appendix. In general, our proposed IceFormer consistently outperforms all efficient Transformers, offering the best accuracy approximation while requiring the least inference time across all five tasks. This demonstrates the generalizability and effectiveness of our model.

²<https://www.kaggle.com/datasets/joaopauloschuler/cifar10-64x64-resized-via-cai-super-resolution>

Table 4: The performance of vanilla Transformer, and approximate attention baselines on the LRA benchmarks.

Method	shared-KQ	ListOps		Text		Retrieval		Image		Pathfinder	
		Acc	Time (s)								
Transformer [25]	✗	0.4255	2.9208	0.6019	0.6933	0.6586	8.3588	0.4132	4.9303	0.7514	0.9620
	✓	0.4145	1.8608	0.5986	0.6603	0.6681	6.7946	0.3844	5.9804	/	/
Reformer [18]	✓	0.4121	1.4281	0.5941	0.2288	0.6467	1.4751	0.3726	3.6927	/	/
LARA [30]	✗	0.4125	0.6146	0.5831	0.2348	0.6401	1.8605	0.3094	2.6720	0.7380	0.5961
Nyströmformer [27]	✗	0.4128	0.7994	0.5838	0.3542	0.6540	2.4179	0.3754	1.7644	0.6758	0.3211
H-Transformer-1D [31]	✗	0.3265	1.9301	0.5944	0.4811	0.5808	3.5605	0.2286	1.2586	0.5286	0.5708
IceFormer (ours)	✗	0.4153	0.3766	0.5978	0.0921	0.6541	0.8337	0.4046	0.5076	0.7442	0.3058
	✓	0.4124	0.4678	0.6001	0.0903	0.6602	0.8480	0.3752	0.9581	/	/

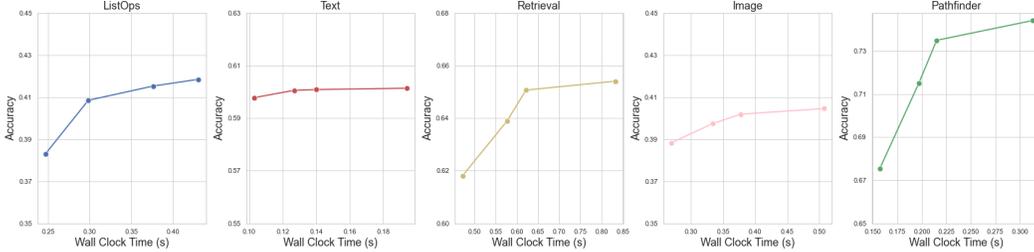


Figure 5: Changes of speed and accuracy on five LRA tasks. X-axis: the averaged wall clock time of attention module. Y-axis: the model prediction accuracy. We test different values of k : {3, 5, 8, 10}.

D.6 Speed & Accuracy Trade-off.

For IceFormer, increasing the extent of approximation generally improves model efficiency but can lead to a decrease in prediction performance. Here, we study how the extent of approximation affects inference speed and accuracy by varying the number of returned candidates of IceFormer, k , from 3 to 10 for each task and present the results in Figure 5. From the figure, we observe that across all tasks, when k becomes larger, IceFormer achieves improved prediction accuracy but becomes less efficient.

D.7 Memory Complexity Analysis.

Table 5 summarizes the maximum memory usage for each method during inference. We employ the same hyper-parameters as in Table 4 and maintain a batch size of 1 to eliminate the need for padding masks. The table reveals that IceFormer consistently exhibits the lowest peak memory usage across all tasks. In comparison to the vanilla Transformer, IceFormer achieves memory savings of up to 0.862 GB.

Table 5: Peak memory usage (GB) on LRA benchmark. The peak memory usage is the total memory usage of the whole program, which includes the memory for the Prioritized DCI database/index.

Method	shared-KQ	ListOps	Text	Retrieval	Image	Pathfinder
Transformer [25]	✗	3.729	4.327	5.031	3.778	3.926
	✓	3.631	4.265	4.877	3.740	/
Reformer [18]	✓	3.623	3.983	4.250	3.687	/
LARA [30]	✗	3.584	4.129	4.566	3.772	3.943
Nyströmformer [27]	✗	3.478	3.982	4.375	3.463	3.845
H-Transformer-1D [31]	✗	3.883	4.328	4.543	3.553	3.603
IceFormer (ours)	✗	3.374	3.834	4.169	3.304	3.465
	✓	3.306	3.756	4.053	3.286	/

D.8 Approximation Quality.

In order to assess how well various efficient Transformers approximate the outputs of the vanilla modified attention module, we measure the approximation error by computing the L2-norm of the difference between their attention module outputs and those of the standard vanilla attention module. The averaged approximation errors for different efficient Transformers, utilizing the same

hyper-parameter settings of Table 7, are summarized in Table 6. As indicated in the table, IceFormer consistently achieves the lowest approximation errors across all LRA tasks, providing further evidence of its approximation efficacy.

Table 6: Quality of the approximation on LRA benchmark. The approximation error of the attention module output is reported for each method across all the tasks.

Method	shared-KQ	ListOps	Text	Retrieval	Image	Pathfinder
Reformer [18]	✓	3.823	3.926	5.452	2.130	/
LARA [30]	✗	2.395	9.456	10.025	22.066	9.261
Nyströmformer [27]	✗	5.758	10.269	6.523	18.789	10.442
H-Transformer-1D [31]	✗	6.110	10.605	5.676	53.926	12.228
IceFormer (ours)	✗	2.140	3.891	1.825	6.873	8.749
	✓	1.562	1.686	2.499	2.127	/

D.9 Hyperparameters for the Baselines and the Proposed Method

For LARA and Nyströmformer, we tuned the parameter *num_landmarks* by optimizing over the range {64, 128, 256, 512, 1024}. For H-Transformer-1D, we tuned the parameter *block_size* by optimizing over the range {64, 128, 256, 512, 1024}. For Reformer, we tuned the parameters *num_hash* and *bucket_size*: we considered the values of *num_hash* in range {1, 2, 4} and the values of *bucket_size* in range {64, 128, 256, 512, 1024}. For IceFormer, we tuned the parameter *top_k* over the range {3, 5, 8, 10, 15, 20}. In general, a larger value for *bucket_size*, *num_landmarks*, *block_size*, or *top_k* indicates less aggressive approximation, meaning that the model performance is closer to that of the vanilla Transformer. We select the values of the hyperparameters that lead to the best accuracy-time trade-off for each model, and list them in Table 7.

Table 7: Hyperparameter settings for different methods.

Method	hyper-parameter	ListOps	Text	Retrieval	Image	Pathfinder
Reformer [18]	num_hash	1	1	1	1	/
	bucket_size	512	128	256	1024	/
LARA [30]	num_landmarks	256	256	512	1024	1024
Nyströmformer [27]	num_landmarks	256	256	512	512	1024
H-Transformer-1D [31]	block_size	1024	512	1024	256	1024
IceFormer	top_k	8	3	10	10	10
	top_k	10	3	10	20	/

E More Details on the LLM Experiment

E.1 LLM Experiments Setting

The LLM that we use in the main paper, vicuna-7b-v1.5-16k, contains 32 attention layers, each of which contains 32 attention heads with dimensionality equals to 128. Its maximum input sequence length is 16,384. We observe varying levels of sparsity across different layers of LLMs, and this sparsity remains consistent across different prompts. Therefore, in all the LLMs experiments in the main paper, we apply IceFormer to approximate relatively sparse layers ranging from the sixteenth to the thirty-first in vicuna-7b-v1.5-16k. This selection encompassed a total of 16 layers, equivalent to half of the total number of layers in the model. We run all the experiment under CPU setting, and the machine used for the LLM experiments is AMD Ryzen 9 5950X. During the inference, we utilize a total of 24 CPU threads.

The *k* in the *k*-NNS of IceFormer for each task of the ZeroSCROLLS benchmark and the LongEval benchmark is defined as:

$$k = \max(\min(\lfloor n * \alpha \rfloor, 50), 30)$$

where *n* is the number of input tokens, $\lfloor x \rfloor$ is the floor function, and α is a hyper-parameter set by the users. In the ZeroSCROLLS benchmark, we set α equals to 4e-3 for tasks SSFD and QMsm; 5e-3 for tasks GvRp, SQAL, Qspr, Nrtv, MuSQ and BkSS; 6e-3 for tasks QALT and SpDg. In the LongEval benchmark, we set α equals to 5e-3 for all the settings of both two tasks.

E.2 Causal Masks and Other Inference-Time Efficient Transformers

In the main paper, we did not compare IceFormer with LARA and Nyströmformer on LLMs. In this section, we elaborate on the problems of causal masks for these two methods.

Most random-feature-based models such as LARA and Nyströmformer group different tokens into different clusters, known as *landmarks*. In order to enable causal masking in these models, not only does the masking need to be applied at the landmark level to prevent the leakage of information from future tokens, an additional set of masks is also required to mask out different numbers of tokens within the same landmark for different queries. The latter is not supported natively and is especially difficult to implement. As a result, it is difficult to apply LARA and Nyströmformer to models that have causal masks.

F Text outputs of IceFormer + LLM

In this section, we provide the text outputs of IceFormer when applied to the LLM (vicuna-7b-v1.5-16k) in Figure 6&7.

Prompt (4k tokens)
<p>A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Please help me to summarize the following:</p> <p>"There is nothing better than sitting on the couch after a long day. Before that, though, one of the first thoughts running through a person's head probably is, "I wonder what I should eat for supper?" After working for roughly eight hours or longer, no one wants to do is stand over a hot stove for an hour when he or she could be spending that time getting caught up on their favorite Netflix show. Instead of firing up the grill, she pushes a few buttons on a microwave, and voilà! A delicious plate of frozen food, be it pizza bites or a corn dog, and cooked to perfection in seconds. Although this is what happens on a regular night in most households, it is a tradition that needs to be put to a stop before it is too late. It's time for people to face the hard and unwelcome facts. Frozen food causes obesity, not just in American, but around the globe.</p> <p>[...]</p> <p>Diabetes is caused by the lack of a hormone produced by the pancreas needed to regulate the amount of glucose in the blood. The reason for why Type 1 diabetes occurs in humans is unknown. Type 2 diabetes is frequent in children that are overweight, less active, and often, these children have a parent with diabetes (Hubbard, 2017). "In my career, Type 1 diabetes was a rare disease in children and Type 2 disease didn't exist. And I'm not that old," said Dr. Robin S. Goland said, who has been in practice for almost 25 years and is the co-director of the Naomi Berrie Diabetes Center at Columbia University Medical Center in New York (Saint Louis, 2014). Diabetes could be prevented in children for the most part but this is challenging to do when almost every frozen and pre-packaged food item is seriously unhealthy and goes against nutritional needs.</p> <p>ASSISTANT:</p>
IceFormer Output Example (k=30)
<p>The article discusses the history and popularity of frozen foods, including the development of quick-freeze methods by Clarence Birdseye and the marketing strategies used by companies like Swanson to promote frozen meals. Frozen foods have become popular due to their convenience, affordability, and availability in grocery stores worldwide. However, the article also notes that frozen meals are often ultra-processed and high in calories, leading to concerns about their impact on health. Additionally, the rise of fast food restaurants has contributed to the increasing consumption of frozen and processed foods in the US.</p>

Figure 6: Output of IceFormer (Top-30) + LLM with 4k input tokens. We ask the LLM to summarize an article titled "Frozen Food: The World's Favorite Killer".

Prompt (8k tokens)
<p>A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user's questions. USER: Please help me to summarize the following:</p> <p>Rapidly growing population along with digitalization and globalization result in constantly increasing number of internet users. (ITU 2019) Closely related to this development is the expansion of ecommerce business. Especially in the times of isolation, caused by the Corona crisis, which we currently live in, we spend much more time on the internet than we used to, and this seems to be great opportunity for the growth of online shopping industry. Lockdowns in the countries across the whole world have had considerable influence on the numbers of online purchases. The statistics compiled by the Organisation for Economic Co-operation and Development confirm this assumption showing that the e-commerce retail in the United States and in the United Kingdom rose between the first quarter and the second quarter of 2020 (first lockdown in the UK on the 26th March 2020, in the U.S. on the 19th March 2020) (from 11.8% to 16.1% respectively from 20.3% to 31.3%). The increasing numbers of online purchases signify higher revenues and at the same time add the importance to this sector.</p> <p>[...]</p> <p>If advertisers wish to be successful in their OBA practices, they should also attempt to find out more about users' reasons for being online. According to the interactive advertising model (Rogers and Thorson 2000), advertisements, which address these motives can achieve better results, because consumers tend to pay more attention towards them, and they are also considered to be easier to remember. (Boerman et al. 2017) Consumers are more willing to accept OBA practices, if they feel that they might support them in the pursuit of their objectives. (Kirmani and Campbell 2004)</p> <p>ASSISTANT:</p>
IceFormer Output Example (k=40)
<p>The rapidly growing population and digitalization have led to an increasing number of internet users, with the ecommerce business expanding in particular during the Corona crisis. Online purchases have risen, and retail volumes have decreased due to lockdowns. Companies are investing more capital in online marketing and advertising, but users often feel irritated by online advertisements, especially when they are not relevant. Online behavioral advertising (OBA) targets consumers based on collected personal information, but it can be effective only if it is highly personalized. However, this can also cause feelings of anxiety and fear, and there have been several scandals related to the misuse of personal data collected online. The importance of the ecommerce sector and the scandals related to online marketing have led to a need to discuss and investigate OBA practices and their influence on consumers.</p>

Figure 7: Output of IceFormer (Top-40) + LLM with 8k input tokens. We ask the LLM to summarize an article titled “Research of How Online Behavioral Advertising Influences Consumers”.