

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/369943552>

A Probabilistic Framework for Pruning Transformers Via a Finite Admixture of Keys

Conference Paper · April 2023

DOI: 10.1109/ICASSP49357.2023.10096107

CITATION

1

READS

234

10 authors, including:



Dung D. Le

VinUniversity

34 PUBLICATIONS 132 CITATIONS

SEE PROFILE



Nhat Ho

University of Texas at Austin

133 PUBLICATIONS 681 CITATIONS

SEE PROFILE



Hung Tran

Applied Artificial Intelligence Institute (A2I2)

24 PUBLICATIONS 101 CITATIONS

SEE PROFILE



Richard Baraniuk

Rice University

760 PUBLICATIONS 57,760 CITATIONS

SEE PROFILE

A PROBABILISTIC FRAMEWORK FOR PRUNING TRANSFORMERS VIA A FINITE ADMIXTURE OF KEYS

Tan M. Nguyen^{†*} Tam Nguyen^{§*} Long Bui^{†*} Hai Do[‡] Duy Khuong Nguyen[‡]
 Dung D. Le^{‡‡} Hung Tran-The^{§§} Nhat Ho^{††} Stan J. Osher[†] Richard G. Baraniuk[§]

[†] Department of Mathematics, University of California, Los Angeles, USA

[§] Department of Electrical and Computer Engineering, Rice University, Houston, USA

[‡] FPT Software AI Center, Ha Noi, Vietnam

^{‡‡} College of Engineering and Computer Science, VinUniversity, Vietnam

^{§§} Applied Artificial Intelligence Institute, Deakin University, Geelong, Australia

^{††} Department of Statistics and Data Sciences, The University of Texas, Austin, USA

ABSTRACT

Pairwise dot product-based self-attention is key to the success of transformers which achieve state-of-the-art performance across a variety of applications in language and vision, but are costly to compute. It has been shown that most attention scores and keys in transformers are redundant and can be removed without loss of accuracy. In this paper, we develop a novel probabilistic framework for pruning attention scores and keys in transformers. We first formulate an admixture model of attention keys whose input data to be clustered are attention queries. We show that attention scores in self-attention correspond to the posterior distribution of this model when attention keys admit a uniform prior distribution. We then relax this uniform prior constraint and let the model learn these priors from data, resulting in a new Finite Admixture of Keys (FiAK). The learned priors are used for pruning away redundant attention scores and keys in the baseline transformers, improving the diversity of attention patterns that the models capture. We corroborate the efficiency of transformers pruned with FiAK on the ImageNet object classification and WikiText-103 language modeling tasks. Our experiments demonstrate that transformers pruned with FiAK yield similar or better accuracy than the baseline dense transformers while being much more efficient in terms of memory and computational cost.

Index Terms— Transformers, admixture models, pruning

1 Introduction

Transformers [1] have been becoming the method of choice in computer vision and machine learning [2, 3, 4, 5]. Thanks to their ability to learn from unlabeled data and from different data modalities, transformers have achieved state-of-the-art performance on a wide range of tasks and applications, including image recognition, object detection, and language modeling [6, 7, 8]. At the core of transformers is the self-attention

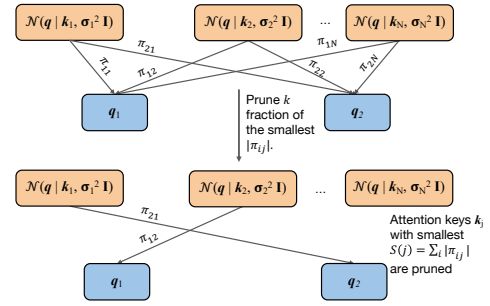


Fig. 1. Our Finite Admixture of Keys (FiAK) models the distribution of the queries q_i in self-attention by an admixture model whose cluster components center around the attention keys k_j , i.e. $p(q_i) = \sum_{j=1}^N \pi_{ij} \mathcal{N}(q_i | k_j, \sigma_j^2 \mathbf{I})$, $i, j = 1, \dots, N$. The prior distributions π_{ij} in the admixture are used to prune redundant attention scores $a_{ij} = \text{softmax}\left(\frac{q_i^\top k_j}{\sqrt{D}}\right)$. The scores $S(j) = \sum_i |\pi_{ij}|$ are used to prune redundant keys k_j . A fraction of attention scores a_{ij} and keys k_j with the smallest $|\pi_{ij}|$ and $S(j)$, respectively, will be pruned away to save memory and computation.

mechanism, which captures the contextual representation of the input sequence by allowing each token in the input sequence to pay attention to other tokens [1, 9]. The capability of self-attention to attain diverse syntactic and semantic representations accounts for the success of transformers [10, 11].

Self-Attention. Given an input $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D_x}$ of N feature vectors, the self-attention transforms it into sequence $\hat{\mathbf{V}} = [\hat{v}_1, \dots, \hat{v}_N]^\top \in \mathbb{R}^{N \times D_v}$ as follows

$$\hat{v}_i = \sum_{j=1}^N \text{softmax}\left(\frac{q_i^\top k_j}{\sqrt{D}}\right) v_j, \text{ for } i = 1, \dots, N, \quad (1)$$

where the scalar $\text{softmax}((q_i^\top k_j)/\sqrt{D})$ can be understood as the attention \hat{v}_i pays to the input feature x_j . The vectors q_i, k_j , and v_j are called the query, key, and value vectors,

*Co-first author.

respectively; these vectors are computed as follows

$$\begin{aligned} [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N]^\top &:= \mathbf{Q} = \mathbf{X}\mathbf{W}_Q^\top \in \mathbb{R}^{N \times D}, \\ [\mathbf{k}_1, \mathbf{k}_2, \dots, \mathbf{k}_N]^\top &:= \mathbf{K} = \mathbf{X}\mathbf{W}_K^\top \in \mathbb{R}^{N \times D}, \\ [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N]^\top &:= \mathbf{V} = \mathbf{X}\mathbf{W}_V^\top \in \mathbb{R}^{N \times D_v}, \end{aligned} \quad (2)$$

where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{D \times D_x}$, and $\mathbf{W}_V \in \mathbb{R}^{D_v \times D_x}$ are the weight matrices. We can write Eqn. 1 into the following form

$$\hat{\mathbf{V}} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}}\right)\mathbf{V} = \mathbf{A}\mathbf{V}, \quad (3)$$

where the softmax function is applied row-wise.

For each query vector \mathbf{q}_i for $i = 1, \dots, N$, an equivalent form of Eqn. 3 to compute the output vector $\hat{\mathbf{v}}_i$ is given by

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j := \sum_{j=1}^N a_{ij} \mathbf{v}_j. \quad (4)$$

The matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and its component a_{ij} for $i, j = 1, \dots, N$ are the attention matrix and attention scores, respectively. Eqn. 3 is also called the ‘‘scaled dot-product attention’’ or ‘‘softmax attention’’. The attention matrix \mathbf{A} after training captures the contextual representation of each token.

Despite the success of transformers in capturing the contextual representation of tokens in the input sequence, it has been shown that the contextual representation learned by the self-attention are redundant and many attention scores and keys explain the same patterns and are not needed [12, 13, 14, 15]. Such redundancy wastes memory and computation during both training and inference while limiting the model’s capacity, posing a challenge to scale up transformers to large-scale tasks.

Contribution. We propose a novel probabilistic model for self-attention, namely the Finite Admixture of Keys (FiAK), that allows pruning attention scores and keys using the prior distributions of attention keys. FiAK models the query distribution $p(\mathbf{q}_i)$ as an admixture of Gaussian distributions $\mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I})$ centering around the attention keys \mathbf{k}_j , $i, j = 1, \dots, N$. Our admixture approach uses different mixture models to represent the queries \mathbf{q}_i and thus helps increase the diversity of attention patterns. Since these mixture models share the same set of component distributions $\mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I})$, FiAK is efficient. The prior distributions of attention keys in FiAK are then used to prune redundant attention scores and keys to improve the memory and computational cost of the model. An illustration of FiAK and our pruning scheme is given in Fig. 1. Our contribution is three-fold:

1. We develop FiAK, a new finite admixture of keys for self-attention that allows key sharing to diversify attention patterns while guaranteeing the model’s efficiency.
2. We design a probabilistic framework for pruning transformers that employs the prior distributions of keys in FiAK to remove redundant attention scores and keys.
3. We demonstrate the advantages of our FiAK-based pruning on ImageNet object classification, and WikiText-103 language modeling tasks.

2 A Finite Admixture of Keys

In this section, we first review the connection between attention scores in self-attention with the posterior distributions from a Gaussian mixture model (GMM) in [16]. We then extend this GMM into a finite admixture of keys (FiAK).

2.1 Background: Attention Scores are Posterior Distributions from a GMM

Given a query $\mathbf{q}_i \in \mathbf{Q}$ and a key $\mathbf{k}_j \in \mathbf{K}$, let t be a K -dimensional binary random variable having a 1-of- K representation in which a particular element t_j is equal to 1 and all other elements are equal to 0. The distribution $p(\mathbf{q}_i | t_j = 1)$ is the likelihood of the query \mathbf{q}_i belongs to the j -th cluster centering around the key \mathbf{k}_j . In particular, let \mathbf{I} be an identity matrix and π_j be the prior distribution $p(t_j = 1)$, the distribution $p(\mathbf{q}_i)$ is given by the following GMM:

$$p(\mathbf{q}_i) = \sum_{j=1}^N \pi_j p(\mathbf{q}_i | t_j = 1) = \sum_{j=1}^N \pi_j \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I}). \quad (5)$$

Following Eqn. 5, the posterior $p(t_j = 1 | \mathbf{q}_i)$ captures how much the query \mathbf{q}_i matches the key \mathbf{k}_j and is computed by

$$\begin{aligned} p(t_j = 1 | \mathbf{q}_i) &= \frac{\pi_j \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2)}{\sum_{j'} \pi_{j'} \mathcal{N}(\mathbf{q}_i | \mathbf{k}_{j'}, \sigma_{j'}^2)} \\ &= \frac{\pi_j \exp[-(\|\mathbf{q}_i\|^2 + \|\mathbf{k}_j\|^2) / 2\sigma_j^2] \exp(\mathbf{q}_i^\top \mathbf{k}_j / \sigma_j^2)}{\sum_{j'} \pi_{j'} \exp[-(\|\mathbf{q}_i\|^2 + \|\mathbf{k}_{j'}\|^2) / 2\sigma_{j'}^2] \exp(\mathbf{q}_i^\top \mathbf{k}_{j'} / \sigma_{j'}^2)}. \end{aligned}$$

Assuming that the query \mathbf{q}_i and the key \mathbf{k}_j are normalized, the prior π_j is uniform, and let $\sigma_j^2 = \sigma^2$, $j = 1, 2, \dots, K$, the posterior $p(t_j = 1 | \mathbf{q}_i)$ can be written in the following form

$$p(t_j = 1 | \mathbf{q}_i) = \frac{\exp(\mathbf{q}_i^\top \mathbf{k}_j / \sigma^2)}{\sum_{j'} \exp(\mathbf{q}_i^\top \mathbf{k}_{j'} / \sigma^2)} = \text{softmax}\left(\mathbf{q}_i^\top \mathbf{k}_j / \sigma^2\right).$$

The equation above becomes Eqn. (4) of the attention score a_{ij} when $\sigma^2 = \sqrt{D}$. Thus, under right assumptions, the attention score a_{ij} between the query \mathbf{q}_i and the key \mathbf{k}_j in a self-attention layer of a transformer plays the role of the posterior distribution $p(t_j = 1 | \mathbf{q}_i)$.

2.2 FiAK: A Finite Admixture of Keys

We extend the GMM of keys for self-attention in Eqn. 5 into a finite admixture of keys so that the attention score a_{ij} can capture more diverse attention patterns and provide a probabilistic framework for pruning transformers.

2.2.1 Finite Admixture Models

A finite mixture distribution of N components for a random array $\mathbf{X} \in \mathbb{R}^{M \times D}$ is given by

$$\mathbf{x}_i \sim \sum_{j=1}^N p_j f(\mathbf{x}; \theta_j), \quad \sum_{j=1}^N p_j = 1, \quad p_j \geq 0, \quad (6)$$

where $\mathbf{x}_i \in \mathbb{R}^D$ is the i -th row of \mathbf{X} randomly sampled from the mixture distribution. f is a chosen probability measure, such as a Gaussian distribution as in Eqn. 5, $p = \{p_1, \dots, p_N\}$ are mixture weights that correspond to the prior π_j , and θ_j denotes the parameter values for the k -th component.

A finite admixture models (FAM) is a generalization of a FMM, in which rows \mathbf{x}_i , $i = 1, \dots, M$, are drawn from different mixture distributions that share N components $f(\mathbf{x}; \theta_j)$, $j = 1, \dots, N$ with different mixture weights

$$\mathbf{x}_i \sim \sum_{j=1}^N p_{ij} f(\mathbf{x}; \theta_j), \quad \sum_{j=1}^N p_{ij} = 1, \quad p_{ij} \geq 0. \quad (7)$$

Comparing to FMM, FAM has better representation capacity thanks to its flexibility in choosing the mixture components. Since all components are shared between mixtures in FAM, FAM is efficient in term of the model size and computational cost for sampling samples from the model.

2.2.2 Finite Admixture of Keys

We propose the finite admixture of keys (FiAK) for the queries in self-attention. In Eqn. 7, let the function $f(\mathbf{x}; \theta_j) = p(\mathbf{q}_i | \mathbf{t}_j = 1) = \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I})$ and $p_{ij} = \pi_{ij} = p_i(\mathbf{t}_j = 1)$ where $\pi_{ij} = p_i(\mathbf{t}_j = 1)$ is the prior distribution $p(\mathbf{t}_j = 1)$ of the mixture corresponding to the query \mathbf{q}_i . FiAK is defined as: **Definition 1** (Finite Admixture of Keys). *Given a set of queries \mathbf{q}_i and keys \mathbf{k}_j in self-attention, $i, j = 1, \dots, N$, the queries \mathbf{q}_i admit a finite admixture of keys if \mathbf{q}_i are sampled from the following finite admixture model:*

$$\mathbf{q}_i \sim \sum_{j=1}^N \pi_{ij} \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I}), \quad \sum_{j=1}^N \pi_{ij} = 1, \quad \pi_{ij} \geq 0. \quad (8)$$

3 Prior-based Pruning via FiAK

Using the prior π_{ij} in FiAK, we propose two novel pruning methods: 1) attention score pruning via FiAK and 2) mixed pruning via FiAK. For comparison with the GMM of keys in Section 2.1, we also derive 3) key pruning via GMM. In all of our proposed methods, attention scores and keys with the smallest importance weights, i.e. $|\hat{\pi}_{ij}|$, $\hat{S}(j)$, and $|\hat{\pi}_j|$ in Algorithm 1, 2, and 3 are pruned away.

Attention Score Pruning. The magnitude of the prior, $|\pi_{ij}|$, in FiAK implies how much the key \mathbf{k}_j is needed to explain the query \mathbf{q}_i . These priors act as importance weights of the keys \mathbf{k}_j given the query \mathbf{q}_i and can be used to prune away the attention score a_{ij} , thus saving memory and computation when computing the self-attention (see Algorithm 1).

Mixed Pruning. To further reduce the computational complexity of the model, we introduce mixed pruning via FiAK in Algorithm 2. In addition to pruning the attention score a_{ij} , we derive the importance weights of the keys \mathbf{k}_j and remove the pairs $(\mathbf{k}_j, \mathbf{v}_j)$ whose importance weights are the smallest. This strategy enables the pruned model to save computation not only at the attention calculation step, but also removes the key vector \mathbf{k}_j and the value vector \mathbf{v}_j , as well as other computations related to these vectors in Eqn. 4.

Key Pruning. We introduce key pruning via GMM (Algorithm 3), which uses the learned prior $|\pi_j|$ in the GMM defined by Eqn. 5 as importance weights to prune the pairs $(\mathbf{k}_j, \mathbf{v}_j)$.

Finetuning the Pruned Network. FiAK introduces additional priors π_{ij} to capture the importance of the attention score

Algorithm 1 Attention Score Pruning via FiAK

Hyperparameter $0 < k < 1$: k fraction of the attention scores a_{ij} to be pruned.

Step 1 Incorporate parameters π_{ij} into the self-attentions.

Step 2 Train the transformer with the additional parameters π_{ij} until convergence.

Step 3 Prune k fraction of the attention scores a_{ij} whose learned coefficients $|\hat{\pi}_{ij}|$ are the smallest.

Step 4 Set the remaining $\hat{\pi}_{ij} = 1$, which corresponds to uniform prior, and finetune the pruned network.

Algorithm 2 Mixed Pruning via FiAK

Hyperparameters $0 < k_1, k_2 < 1$: k_1 fraction of the total attention scores a_{ij} to be pruned; k_2 fraction of pairs (key, value) to be pruned.

Step 1 and Step 2 Same as **Step 1** and **Step 2** of Algorithm 1.

Step 3 Calculate the importance score $\hat{S}(j)$ of each pair $(\mathbf{k}_j, \mathbf{v}_j)$:

$$\hat{S}(j) = \sum_i |\hat{\pi}_{ij}|, \text{ or } \frac{1}{N-j+1} \sum_i |\hat{\pi}_{ij}| \text{ for autoregressive tasks.}$$

Then prune k_2 fraction of the pairs $(\mathbf{k}_j, \mathbf{v}_j)$ with the smallest scores $\hat{S}(j)$.

Step 4 Prune k_1 fraction of the remain unpruned a_{ij} whose corresponding $|\hat{\pi}_{ij}|$ are the smallest $\hat{k}_1 = 1 - \frac{1-k_1}{1-k_2}$.

Step 5 Follow **Step 4** of Algorithm 1.

Algorithm 3 Key Pruning via GMM

Hyperparameter $0 < k < 1$: k fraction of the keys to be pruned.

Step 1 Incorporate parameters π_j into the self-attentions.

Step 2 Train the transformer with the additional parameters π_j until convergence.

Step 3 Prune k fraction of the key-value pairs $(\mathbf{k}_j, \mathbf{v}_j)$, whose corresponding learned mixing-coefficients $|\hat{\pi}_j|$ are the smallest.

Step 4 Set the remaining $\hat{\pi}_j = 1$, i.e. uniform prior, and finetune the pruned network.

a_{ij} . After pruning, those extra parameters can be removed by setting them to 1, which corresponds to using uniform priors. The network is then finetuned for more epochs to obtain competitive accuracy compared to the dense baseline network.

4 Experimental Results

We empirically corroborate the advantages of the models pruned via our proposed FiAK-based pruning methods over the dense baseline model on the ImageNet object classification task. We refer to transformers that use FiAK-based attention defined by Eqn. 8 as FiAKformer and transformers that use GMM-based attention defined by Eqn. 5 as GMMformer.

Model and setting. We use the DeiT-tiny model [5] with 12 layers and 4 attention heads per layer. The model dimension is 192. To train the models, we follow the same setting and configuration as for the baseline [5], with the initialization of the learnable priors π_{ij} and π_j set to be $\frac{1}{\sqrt{N}}$ and $\frac{1}{N}$, respectively, where N is the input sequence’s length.

Results. *Pruned models from attention score and mixed pruning via FiAK attain much better accuracy than the DeiT-tiny baseline while being significantly more efficient (See Ta-*

Table 1. Top-1 and top-5 accuracy (%) of the pruned models from the attention score and mixed pruning via FiAK on the Imagenet dataset compared to the dense baseline DeiT-tiny [5].

Method	Top-1 Acc	Top-5 Acc
<i>Baseline DeiT-tiny</i>	72.23	91.13
GMMformer	72.96	91.64
Key pruning $k = 30\%$	71.57	90.80
FiAKformer	73.50	91.90
Attention-score pruning $k = 50\%$	73.56	91.95
Attention-score pruning $k = 60\%$	73.67	91.91
Attention-score pruning $k = 70\%$	73.09	91.57
Mixed pruning $k_1 = 70\%, k_2 = 15\%$	72.78	91.38
Mixed pruning $k_1 = 70\%, k_2 = 20\%$	72.25	91.14

Table 2. Comparison to other pruning methods on Imagenet task.

Method	FLOPS reduced (%)	Acc-1 (%)
<i>DeiT-tiny</i>	0.00	72.23
Head pruning [12]	23.69	68.59
S^2ViTE [17]	23.69	70.12
Attention-score pruning $k = 70\%$	8.50	73.09
Mixed pruning $k_1 = 70\%, k_2 = 20\%$	13.00	72.25
Mixed pruning $k_1 = 70\%, k_2 = 20\%$ + S^2ViTE [17]	22.76	72.24

ble 1). Attention score pruning via FiAK at different pruning fractions $k = 50\%$, 60% and 70% result in the highest accuracies. In particular, at the pruning fractions $k = 50\%$ and 60% , we observe substantial accuracy improvement over the dense baseline (1.33% and 1.44% in top-1 accuracy, respectively). These two pruned models also outperform the dense FiAKformer. On the other hand, mixed pruning with the same attention score pruning fraction, $k_1 = 70\%$ and different key pruning fractions, $k_2 = 15\%$ and 20% , gain better accuracy compared to the baseline while obtaining the most computation and memory reduction (See Fig. 2). Table 1 also shows the advantage of the FiAK-based pruning over the GMM-based pruning and validate the need of using admixture to model the self-attention and design its effective pruning schemes.

Comparison to Other Pruning Methods. We compare our FiAK-based pruning schemes with other pruning methods for transformers on the ImageNet task (see Table 2 below). Compared to the head pruning [12] and S^2ViTE [17], our schemes prune the model less but increase its accuracy. Combining with the S^2ViTE [17], mixed FiAK pruning can increase the FLOPs reduction up to 22.76% while maintaining similar advantage in accuracy on the ImageNet task.

Other Tasks: Language Modeling on WikiText-103.

To examine the effectiveness of our pruning methods across different data modalities, we experiment with the word-level language modeling task on WikiText-103 [18]. We summarize our results in Table 3.

Efficiency Analysis. We analyze the computation and memory complexity of the pruned models trained for the ImageNet object classification task (see Fig. 2). We observe that the *efficiency advantage of models pruned via FiAK over the*

Table 3. Test perplexity of pruned FiAKformer for the language modeling task on Wikitext-103 dataset.

Method	Perplexity (PPL)
<i>Baseline softmax transformer</i>	34.29
FiAKformer	33.69
Attention score pruning 40%	33.88
Attention score pruning 50%	34.28
Mixed pruning $k_1 = 40\%, k_2 = 10\%$	34.21




Fig. 2. FLOPS and memory ratios at inference the models pruned with FiAK/GMM-based schemes and the DeiT-tiny baseline.

baseline model grows with the sequence length. FiAK-based pruning also wins in real time. On the ImageNet task, the latency for the dense baseline and our attention-score pruned FiAKformer, $k = 70\%$, are 508 and 649 images/second (on GPU) and 76 and 95 images/second (on CPU), respectively.

5 Related Work

Works in pruning transformers can be categorized into two groups: 1) head pruning and 2) token pruning. An early work in head pruning calculates the head sensitivity to decide to prune a head or not [12]. [19] employs layerwise relevance propagation to decide the head importance. The head importance can also be learned in a data-driven manner as in [20]. For token pruning, [21] computes a token’s importance score as average attention score of other tokens to that token. A dropout-based approach that stochastically determines a sequence length at each layer has also been used to prune redundant tokens [22]. [23] learns an attention mask for token pruning adaptively. Our FiAK-based approach is complementary to these methods.

6 Concluding Remarks

In this paper, we propose FiAK, a novel finite admixture of keys for self-attention, that model the distribution of queries q_i in self-attention as an admixture of Gaussian distributions $\mathcal{N}(q_i | k_j, \sigma_j^2 \mathbf{I})$ whose centers are the attention keys $k_j, i, j = 1, \dots, N$. Using the prior distributions of the attention keys in FiAK, we propose a probabilistic pruning framework to remove redundant attention scores and keys in transformers. We verify that models pruned by our FiAK-based pruning methods improve the memory and computational cost over the baseline dense transformers while achieving comparable or better accuracy. Admixture models are equivalent to Latent Dirichlet Allocation (LDA) models under a uniform Dirichlet prior. Extending FiAK into an LDA-based framework for pruning transformers is an interesting research direction.

7 References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [2] Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones, “Character-level language modeling with deeper self-attention,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 3159–3166.
- [3] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov, “Transformer-XL: Attentive language models beyond a fixed-length context,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 2978–2988, Association for Computational Linguistics.
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [5] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou, “Training data-efficient image transformers and distillation through attention,” in *International Conference on Machine Learning*, 2021.
- [6] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever, “Improving language understanding by generative pre-training,” *OpenAI report*, 2018.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Minneapolis, Minnesota, June 2019, pp. 4171–4186, Association for Computational Linguistics.
- [8] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [9] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio, “Learning phrase representations using RNN encoder–decoder for statistical machine translation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, Oct. 2014, pp. 1724–1734, Association for Computational Linguistics.
- [10] Ian Tenney, Dipanjan Das, and Ellie Pavlick, “BERT rediscovered the classical NLP pipeline,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 4593–4601, Association for Computational Linguistics.
- [11] John Hewitt and Percy Liang, “Designing and interpreting probes with control tasks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China, Nov. 2019, pp. 2733–2743, Association for Computational Linguistics.
- [12] Paul Michel, Omer Levy, and Graham Neubig, “Are sixteen heads really better than one?,” in *Advances in Neural Information Processing Systems*, 2019, vol. 32.
- [13] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov, “Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned,” *arXiv preprint arXiv:1905.09418*, 2019.
- [14] Srinadh Bhojanapalli, Ayan Chakrabarti, Himanshu Jain, Sanjiv Kumar, Michal Lukasik, and Andreas Veit, “Eigen analysis of self-attention and its reconstruction from partial computation,” *arXiv preprint arXiv:2106.08823*, 2021.
- [15] Tan Minh Nguyen, Tam Minh Nguyen, Hai Ngoc Do, Khai Nguyen, Vishwanath Saragadam, Minh Pham, Nguyen Duy Khuong, Nhat Ho, and Stanley Osher, “Improving transformer with an admixture of attention heads,” in *Advances in Neural Information Processing Systems*, 2022.
- [16] Tam Minh Nguyen, Tan Minh Nguyen, Dung Duy Le, Duy Khuong Nguyen, Viet-Anh Tran, Richard Baraniuk, Nhat Ho, and Stanley Osher, “Improving transformers with probabilistic attention keys,” in *International Conference on Machine Learning*, 2022.
- [17] Tianlong Chen, Yu Cheng, Zhe Gan, Lu Yuan, Lei Zhang, and Zhangyang Wang, “Chasing sparsity in vision transformers: An end-to-end exploration,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 19974–19988, 2021.
- [18] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher, “Pointer sentinel mixture models,” in *International Conference on Learning Representations*, 2017.
- [19] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov, “Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, July 2019, pp. 5797–5808, Association for Computational Linguistics.
- [20] Jiaoda Li, Ryan Cotterell, and Mrinmaya Sachan, “Differentiable subset pruning of transformer heads,” *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 1442–1459, 2021.
- [21] Saurabh Goyal, Anamitra Roy Choudhury, Saurabh Raje, Venkatesan Chakaravarthy, Yogish Sabharwal, and Ashish Verma, “Power-bert: Accelerating bert inference via progressive word-vector elimination,” in *International Conference on Machine Learning*, 2020.
- [22] Gyuwan Kim and Kyunghyun Cho, “Length-adaptive transformer: Train once with length drop, use anytime with search,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 2021, pp. 6501–6511, Association for Computational Linguistics.
- [23] Sehoon Kim, Sheng Shen, David Thorsley, Amir Gholami, Woosuk Kwon, Joseph Hassoun, and Kurt Keutzer, “Learned token pruning for transformers,” *arXiv preprint arXiv:2107.00910*, 2021.