
Improving Transformers with Probabilistic Attention Keys

Tam Nguyen^{*1} Tan M. Nguyen^{*2} Dung D. Le³ Duy Khuong Nguyen¹ Viet-Anh Tran⁴ Richard G. Baraniuk⁵
Nhat Ho^{**6} Stanley J. Osher^{**2}

Abstract

Multi-head attention is a driving force behind state-of-the-art transformers, which achieve remarkable performance across a variety of natural language processing (NLP) and computer vision tasks. It has been observed that for many applications, those attention heads learn redundant embedding, and most of them can be removed without degrading the performance of the model. Inspired by this observation, we propose Transformer with a Mixture of Gaussian Keys (Transformer-MGK), a novel transformer architecture that replaces redundant heads in transformers with a mixture of keys at each head. These mixtures of keys follow a Gaussian mixture model and allow each attention head to focus on different parts of the input sequence efficiently. Compared to its conventional transformer counterpart, Transformer-MGK accelerates training and inference, has fewer parameters, and requires fewer FLOPs to compute while achieving comparable or better accuracy across tasks. Transformer-MGK can also be easily extended to use with linear attention. We empirically demonstrate the advantage of Transformer-MGK in a range of practical applications, including language modeling and tasks that involve very long sequences. On the Wikitext-103 and Long Range Arena benchmark, Transformer-MGKs with 4 heads attain comparable or better performance to the baseline transformers with 8 heads.

^{*}Equal contribution. Co-first authors are listed in alphabetical order ^{**}Co-last authorship ¹FPT Software AI Center, Ha Noi, Vietnam ²Department of Mathematics, University of California, Los Angeles, USA ³College of Engineering and Computer Science, VinUniversity, Ha Noi, Vietnam ⁴Deezer Research, France ⁵Department of Electrical and Computer Engineering, Rice University, Houston, USA ⁶Department of Statistics and Data Sciences, The University of Texas at Austin, USA. Correspondence to: Tan Nguyen <tanmnguyen89@ucla.edu>.

1. Introduction

Transformers (Vaswani et al., 2017) have become the state-of-the-art model for sequence processing tasks, solving many challenging problems in natural language processing and computer vision (Al-Rfou et al., 2019; Dai et al., 2019; Williams et al., 2018; Devlin et al., 2018; Brown & et al., 2020; Howard & Ruder, 2018; Rajpurkar et al., 2016; Dehghani et al., 2018; So et al., 2019; Dosovitskiy et al., 2020; Touvron et al., 2020; Nguyen et al., 2021a). These models can also transfer the learned knowledge from a pre-trained model to tasks that involve different data modalities and have limited supervision (Radford et al., 2018; 2019; Devlin et al., 2018; Yang et al., 2019; Liu et al., 2019). The success of transformers is rooted in the *self-attention* mechanism as their fundamental building blocks for modeling (Cho et al., 2014; Parikh et al., 2016; Lin et al., 2017). For each token, self-attention computes a weighted average of the feature representations of other tokens where the weight is proportional to a similarity score between each pair of tokens. This mechanism allows a token to pay attention to other tokens in the sequence and attain a contextual representation (Bahdanau et al., 2014; Vaswani et al., 2017; Kim et al., 2017). It has been shown that the representation capacity of the attention mechanism (Tenney et al., 2019) and its capability of capturing diverse syntactic and semantic relationships (Tenney et al., 2019; Vig & Belinkov, 2019; Clark et al., 2019; Voita et al., 2019a; Hewitt & Liang, 2019) is key to the impressive performance of transformers in practice.

1.1. Self-Attention

For a given input sequence $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D_x}$ of N feature vectors, self-attention transforms \mathbf{X} into the output sequence \mathbf{H} in the following two steps:

Step 1. The input sequence \mathbf{X} is projected into the query matrix \mathbf{Q} , the key matrix \mathbf{K} , and the value matrix \mathbf{V} via three linear transformations

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q^\top; \mathbf{K} = \mathbf{X}\mathbf{W}_K^\top; \mathbf{V} = \mathbf{X}\mathbf{W}_V^\top,$$

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 denote $\mathbf{Q} := [\mathbf{q}_1, \dots, \mathbf{q}_N]^\top$, $\mathbf{K} := [\mathbf{k}_1, \dots, \mathbf{k}_N]^\top$, and $\mathbf{V} := [\mathbf{v}_1, \dots, \mathbf{v}_N]^\top$, where the vectors $\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i$ for $i = 1, \dots, N$ are the query, key, and value

vectors, respectively.

Step 2. The output sequence $\mathbf{H} := [\mathbf{h}_1, \dots, \mathbf{h}_N]^\top$ is then computed as follows

$$\mathbf{H} = \text{softmax}\left(\mathbf{Q}\mathbf{K}^\top/\sqrt{D}\right)\mathbf{V} := \mathbf{A}\mathbf{V}, \quad (1)$$

where the softmax function is applied to each row of the matrix $(\mathbf{Q}\mathbf{K}^\top)/\sqrt{D}$. For each query vector \mathbf{q}_i for $i = 1, \dots, N$, an equivalent form of Eqn. (1) to compute the output vector \mathbf{h}_i is given by

$$\mathbf{h}_i = \sum_{j=1}^N \text{softmax}\left(\mathbf{q}_i^\top \mathbf{k}_j/\sqrt{D}\right) \mathbf{v}_j := \sum_{j=1}^N a_{ij} \mathbf{v}_j. \quad (2)$$

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. The self-attention computed by Eqn. (1) and (2) is called the scaled dot-product attention or softmax attention. In our paper, we call a transformer that uses this attention the softmax transformer. The structure that the attention matrix \mathbf{A} learns from training determines the ability of the self-attention to capture contextual representation for each token.

Multi-head Attention Each output sequence \mathbf{H} forms an attention head. In multi-head attention, multiple heads are concatenated to compute the final output. Let H be the number of heads and $\mathbf{W}^O \in \mathbb{R}^{HD_v \times HD_v}$ be the projection matrix for the output. The multi-head attention is defined as $\text{MultiHead}(\{\mathbf{Q}, \mathbf{K}, \mathbf{V}\}_{i=1}^H) = \text{Concat}(\mathbf{H}_1, \dots, \mathbf{H}_H) \mathbf{W}^O$.

Even though multi-head attention extends single-head attention to capture diverse attention patterns and improve the performance of transformers, it has been shown that transformers for practical tasks including sequence classification and language modeling learn redundant heads (Michel et al., 2019). These redundant heads compute similar attention mappings. Having many of them in the model limits the representation capacity of the transformer while wasting parameters, memory and computation, impeding the application of transformers to many important large-scale tasks.

1.2. Contribution

We establish the correspondence between self-attention in transformer and a Gaussian mixture model (GMM) and propose Transformer with a Mixture of Gaussian Keys (Transformer-MGK), a novel class of transformers that can avoid the head redundancy. At the core of Transformer-MGK is replacing the attention key \mathbf{k}_j in each head by a GMM to allow the query \mathbf{q}_i , as well as its associated token, to attend to more diverse positions in the input sequence, thereby increasing the representation of each attention head and reducing the chance of learning redundant heads. Our contribution is four-fold:

1. We construct a GMM and show that attention scores in self-attention match posterior distribution in our model, providing a probabilistic framework to study self-attention in transformers.
2. Under our probabilistic framework for self-attention, we introduce an additional mixture of Gaussian to model each attention key. We empirically show that this mixture of Gaussian keys (MGK) can capture a diversity of attention patterns, thus alleviating head redundancy.
3. We extend our MGK to use with linear attentions and propose the mixture of linear keys (MLK) for efficient computation and better memory footprint.
4. We empirically show that Transformer-MGK and Transformer-MLK are comparable or better than the corresponding baseline transformers with softmax and linear attentions while only using half the number of attention heads and reducing both model complexity measured by the number of parameters and computational cost in terms of FLOPs.

Organization: We structure this paper as follows: In Section 2, we establish the connection between GMM and self-attention and then present our Transformer-MGK and its extensions including Transformer-MLK. In Section 3, we validate and empirically analyze the efficiency and accuracy of Transformer-MGK/MLK. We discuss related works in Section 4. The paper ends up with concluding remarks. More experimental details are provided in the Appendix.

2. Transformer with a Mixture of Gaussian Keys

2.1. Attention Score as a Posterior Distribution

We first consider a query $\mathbf{q}_i \in \mathbf{Q}$ and a key $\mathbf{k}_j \in \mathbf{K}$. Let \mathbf{t} be a N -dimensional binary random variable having a 1-of- N representation in which a particular element \mathbf{t}_j is equal to 1 and all other elements are equal to 0. We use \mathbf{t}_j to indicate the position j of the key \mathbf{k}_j . Let \mathbf{I} be the identity matrix, we model the distribution $p(\mathbf{q}_i)$ by the following GMM:

$$p(\mathbf{q}_i) = \sum_{j=1}^N \pi_j p(\mathbf{q}_i | \mathbf{t}_j = 1) = \sum_{j=1}^N \pi_j \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2 \mathbf{I}), \quad (3)$$

where π_j is the prior $p(\mathbf{t}_j = 1)$. Given the query \mathbf{q}_i , how likely \mathbf{q}_i matches the key \mathbf{k}_j is given by posterior $p(\mathbf{t}_j = 1 | \mathbf{q}_i)$. This posterior is computed as follows

$$\begin{aligned} p(\mathbf{t}_j = 1 | \mathbf{q}_i) &= \frac{\pi_j \mathcal{N}(\mathbf{q}_i | \mathbf{k}_j, \sigma_j^2)}{\sum_{j'} \pi_{j'} \mathcal{N}(\mathbf{q}_i | \mathbf{k}_{j'}, \sigma_{j'}^2)} \\ &= \frac{\pi_j \exp\left[-(\|\mathbf{q}_i\|^2 + \|\mathbf{k}_j\|^2)/2\sigma_j^2\right] \exp\left(\mathbf{q}_i \mathbf{k}_j^\top / \sigma_j^2\right)}{\sum_{j'} \pi_{j'} \exp\left[-(\|\mathbf{q}_i\|^2 + \|\mathbf{k}_{j'}\|^2)/2\sigma_{j'}^2\right] \exp\left(\mathbf{q}_i \mathbf{k}_{j'}^\top / \sigma_{j'}^2\right)}. \end{aligned}$$

We further assume that the query \mathbf{q}_i and the key \mathbf{k}_j are normalized, and the prior π_j is uniform. We will justify these assumptions in our Remarks at the end of this section. We also let $\sigma_j^2 = \sigma^2$, $j = 1, 2, \dots, K$. Then the posterior $p(\mathbf{t}_j = 1 | \mathbf{q}_i)$ can be written as

$$p(\mathbf{t}_j = 1 | \mathbf{q}_i) = \exp(\mathbf{q}_i \mathbf{k}_j^\top / \sigma^2) / \sum_{j'} \exp(\mathbf{q}_i \mathbf{k}_{j'}^\top / \sigma^2). \quad (4)$$

The right-hand side of Eqn. (4) matches the attention score given in Eqn. (2) when $\sigma^2 = \sqrt{D}$. Thus, we show that under right assumptions, the attention score between the query \mathbf{q}_i and the key \mathbf{k}_j in an attention unit of a transformer is the posterior $p(\mathbf{t}_j = 1 | \mathbf{q}_i)$, which indicates the *responsibility* that the key \mathbf{k}_j takes for ‘explaining’ the query \mathbf{q}_i , which in turn decide, for example, how much a token at position i pays attention to a token at position j in the input sequence.

Remark 1. *The assumption that the query \mathbf{q}_i and the key \mathbf{k}_j are normalized is realistic and not artificial. In many applications, those two vectors are normalized. (Schlag et al., 2021) points out that such normalization is to avoid instability occurring during the training.*

Remark 2. *In practice, the prior is chosen to be uniform when there is no prior knowledge available.*

2.2. Transformer with a Mixture of Gaussian Keys: Each Key is Again a Gaussian Mixture Model

As we have seen from Eqn. (4), the key k_j is used to explain the query q_i via the posterior $\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i)$. Via this simple connection, each query q_i is treated to be as a sample from the mixture of N keys $\sum_{j=1}^N \pi_j \mathcal{N}(q_i | k_j, \sigma_j^2 I_d)$. However, the assumption that the distribution $\mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1)$ at each subpopulation is Gaussian in Eqn. (3) can be quite strong as there is no guarantee that this assumption is valid in practice. In particular, it may happen that the distribution of each subpopulation is asymmetric or skewed or even multimodal. Therefore, using the Gaussian distribution for each subpopulation can potentially limit the explanation power and diversity of each subpopulation/ key. It indicates that we should use more expressive distributions to represent $\mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1)$.

Mixture of Gaussian keys: To improve the explanation power of each key k_j , potentially increase the representation of each attention head, and reduce the chance of learning redundant heads, we would like to model it as mixture of Gaussian distributions. We refer to this model as *Transformer with a Mixture of Gaussian Keys* (Transformer-MGK). In particular, in Transformer-MGK we model each key \mathbf{k}_j at position j as a mixture of M Gaussians $\mathcal{N}(\mathbf{k}_{j,r}, \sigma_{j,r}^2 \mathbf{I})$, $r = 1, 2, \dots, M$. Here we are overloading the notation a little bit and use $\mathbf{k}_{j,r}$ and $\sigma_{j,r}^2 \mathbf{I}$ to denote the mean and covariance matrix of the r^{th} Gaussian at position j . Let \mathbf{z} be a M -dimensional binary random variable having a 1-of- M

representation. We use \mathbf{z}_r to indicate the r^{th} Gaussian in the mixture. Let $\pi_{j,r} \equiv \mathbb{P}(\mathbf{z}_r = 1 | \mathbf{t}_j = 1)$, our MGK can be written as

$$\begin{aligned} \mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1) &= \sum_r \mathbb{P}(\mathbf{z}_r = 1 | \mathbf{t}_j = 1) \mathbb{P}(\mathbf{q}_i | \mathbf{z}_r = 1, \mathbf{t}_j = 1) \\ &= \sum_r \pi_{j,r} \mathcal{N}(\mathbf{q}_i | \mathbf{k}_{j,r}, \sigma_{j,r}^2 \mathbf{I}). \end{aligned} \quad (5)$$

Our motivation of using mixture of Gaussian distributions to represent the distribution of each subpopulation in Eqn. (5) stems from the following important approximation result:

Theorem 1. *Assume that P is probability distribution on $[-a, a]^d$ for some $a > 0$ and admits density function p such that p is differentiable and bounded. Then, for any given variance $\sigma > 0$ and for any $\epsilon > 0$, there exists a mixture of K components $\sum_{i=1}^K \pi_i \mathcal{N}(\theta_i, \sigma^2 \mathbf{I})$ where $K \leq (C \log(1/\epsilon))^d$ for some universal constant C such that*

$$\sup_{x \in \mathbb{R}^d} |p(x) - \sum_{i=1}^K \pi_i \phi(x | \theta_i, \sigma^2 \mathbf{I})| \leq \epsilon,$$

where $\phi(x | \theta, \sigma^2 \mathbf{I})$ is the density function of multivariate Gaussian distribution with mean θ and covariance matrix $\sigma^2 \mathbf{I}$.

The proof of Theorem 1 is in Appendix C. The result of Theorem 1 suggests that regardless of the real form of $\mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1)$, we can use finite mixture of Gaussian distributions to approximate $\mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1)$. It allows us to have richer approximation of $\mathbb{P}(\mathbf{q}_i | \mathbf{t}_j = 1)$ than by using a simple Gaussian distribution in Eqn. (3). Similar to the derivation above, the posterior $p(\mathbf{t}_j = 1 | \mathbf{q}_i)$ in Transformer-MGK can be written as

$$\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i) = \frac{\sum_r \pi_{j,r} \exp(\mathbf{q}_i \mathbf{k}_{j,r}^\top / \sigma_{j,r}^2)}{\sum_{j'} \sum_r \pi_{j',r} \exp(\mathbf{q}_i \mathbf{k}_{j',r}^\top / \sigma_{j',r}^2)}. \quad (6)$$

Furthermore, in Transformer-MGK, we relax the assumption that the queries and keys are normalized. Thus, when computing $\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i)$, we compute the Gaussian kernels between the queries and keys instead of their dot products. The posterior $\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i)$ in Transformer-MGK is then given by

$$\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i) = \frac{\sum_r \pi_{j,r} \exp(-\|\mathbf{q}_i - \mathbf{k}_{j,r}\|^2 / 2\sigma_{j,r}^2)}{\sum_{j'} \sum_r \pi_{j',r} \exp(-\|\mathbf{q}_i - \mathbf{k}_{j',r}\|^2 / 2\sigma_{j',r}^2)}. \quad (7)$$

As proven in Section 2.1, this posterior corresponds to the attention score. Thus, Eqn. (7) is the formula for computing the attention score in Transformer-MGK. We compute the output vector \mathbf{h}_i of the self-attention in Transformer-MGK

as follows

$$\mathbf{h}_i = \sum_j \left(\frac{\sum_r \pi_{jr} \exp(-\|\mathbf{q}_i - \mathbf{k}_{jr}\|^2/2\sigma_{jr}^2)}{\sum_{j'} \sum_r \pi_{j'r} \exp(-\|\mathbf{q}_i - \mathbf{k}_{j'r}\|^2/2\sigma_{j'r}^2)} \right) \mathbf{v}_j.$$

2.3. Inference and Learning via the Expectation Maximization Algorithm

Let $\gamma_{ir} \equiv \mathbb{P}(z_r = 1 | \mathbf{q}_i, \mathbf{t}_j = 1)$, in MGK, we apply the E-step inference in the Expectation-Maximization (EM) algorithm to estimate this posterior given the query \mathbf{q}_i . The posterior γ_{ir} is also known as the *responsibility* that the component $\mathcal{N}(\mathbf{k}_{jr}, \sigma_{jr}^2 \mathbf{I})$ takes to account for the observation, which in MGK is the query \mathbf{q}_i . Below we propose two approaches to estimate this responsibility.

Soft E-step Using soft E-step inference, the EM algorithm makes a soft assignment, in which each query is associated with all clusters. The responsibilities are then given by

$$\gamma_{ir} = \frac{\pi_{jr} \exp(-\|\mathbf{q}_i - \mathbf{k}_{jr}\|^2/2\sigma_{jr}^2)}{\sum_{r'} \pi_{j'r'} \exp(-\|\mathbf{q}_i - \mathbf{k}_{j'r'}\|^2/2\sigma_{j'r'}^2)}. \quad (8)$$

At learning time, the responsibilities estimated by Eqn. (8) are used to update the prior π_{jr} , i.e. $\pi_{jr} = N_{jr}/N$, where N is the number of queries and $N_{jr} = \sum_{i=1}^N \gamma_{ir}$. These updated priors π_{jr} are then used in Eqn. (7) to compute attention scores.

Hard E-step Hard E-step performs a hard assignment of queries to key clusters, in which each query is associated uniquely with one cluster. This is similar to the K -means algorithm (Lloyd, 1982) and corresponds to the MGK at the limit when the variance parameter σ_{jr}^2 goes to 0. Following the derivation of K -means from a GMM in (Bishop, 2006), Eqn. (7) becomes

$$\mathbb{P}(\mathbf{t}_j = 1 | \mathbf{q}_i) = \frac{\max_r \exp(-\|\mathbf{q}_i - \mathbf{k}_{jr}\|^2/2\sigma_{jr}^2)}{\sum_{j'} \max_r \exp(-\|\mathbf{q}_i - \mathbf{k}_{j'r}\|^2/2\sigma_{j'r}^2)}.$$

Remark 3. *The hard E-step inference allows the attention score to be computed more efficiently because the priors π_{jr} no longer play an active role in the algorithm and can be completely ignored.*

Learning via Stochastic Gradient Descent (SGD) In order to increase the efficiency of the model, in MGK, we fix the variance parameter σ_{jr}^2 to be \sqrt{D} as in the standard softmax attention and make the cluster means, i.e. the keys, \mathbf{k}_{jr} learnable parameters. We also make the prior π_{jr} learnable parameters as one of the design options. In that case, both \mathbf{k}_{jr} and π_{jr} are learned via SGD. This update via SGD can be considered as a generalized M-step (Bishop, 2006).

Design Options for Keys (Option A) We follow the standard setting in the softmax transformer and make the keys

\mathbf{k}_{jr} a linear projection of the input \mathbf{x}_j , i.e. $\mathbf{k}_{jr} = \mathbf{x}_j \mathbf{W}_{K_r}^\top$, where $\mathbf{x}_j \in \mathbb{R}^{1 \times D_x}$, $\mathbf{W}_{K_r} \in \mathbb{R}^{D \times D_x}$ and $r = 1, 2, \dots, M$. (Option B) Alternatively, we also make the keys \mathbf{k}_{jr} shifted version of each other to save computation, i.e. $\mathbf{k}_{jr} = \mathbf{x}_j \mathbf{W}_K^\top + \mathbf{b}_r$, where $\mathbf{W}_K \in \mathbb{R}^{D \times D_x}$.

2.4. Transformer with a Mixture of Linear Keys

The MGK can be easily extended to use with linear attentions. We call that model Transformer with a Mixture of Linear Keys (Transformer-MLK). In this section, we adopt the formulation of linear attentions from (Katharopoulos et al., 2020) to derive Transformer-MLK. Similar approach can be taken to derive Transformer-MLK when using with other linear attentions such as those in performers (Choromanski et al., 2021) and fast-weight transformers (Schlag et al., 2021). In Transformer-MLK, the Gaussian kernel in softmax attention is linearized as the product of feature maps $\phi(\cdot)$ on the vectors \mathbf{q}_i and \mathbf{k}_j . The associative property of matrix multiplication is then utilized to derive the following efficient computation of the attention map

$$\mathbf{h}_i = \frac{\phi(\mathbf{q}_i)^\top \sum_j \sum_r \pi_{jr} \phi(\mathbf{k}_{jr}) \mathbf{v}_j^\top}{\phi(\mathbf{q}_i)^\top \sum_j \sum_r \pi_{jr} \phi(\mathbf{k}_{jr})}.$$

Replacing $\sum_j \sum_r \pi_{jr} \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_{jr}) \mathbf{v}_j^\top$ with $\phi(\mathbf{q}_i)^\top \sum_j \sum_r \pi_{jr} \phi(\mathbf{k}_{jr}) \mathbf{v}_j^\top$, as in linear transformers, reduces the memory and computational cost of computing the attention map in Transformer-MLK from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$, making Transformer-MLK scalable to very long sequences.

2.5. Reduction in Model Complexity and Computational Cost from Transformer-MGK

We provide a detailed analysis of reduction in model complexity and computational cost of Transformer-MGK over the baseline softmax attention in Appendix B.

3. Experimental Results

In this section, we numerically justify the efficiency of Transformer-MGK/MLK and empirically study the advantage of using mixture of keys on various benchmarks, including different tasks in the Long Range Arena (LRA) (Tay et al., 2021) (Section 3.1) and language modeling on Wikitext-103 (Merity et al., 2017) (Section 3.2). We aim to show that: (i) Transformer-MGK/MLK with half the number of heads is comparable or better than the baseline softmax and linear transformers with the full number of heads while being more efficient in both computational cost and memory footprints; (ii) Mixture of keys helps reduce the redundancy in multi-head transformers and benefits learning of the long-term dependency in long input sequences; (iii) Using the same number of heads, Transformer-MGK/MLK

Table 1. Test Accuracy (%) of Transformer-MGK compared with the baseline softmax transformer on the LRA benchmark. Our Transformer-MGKs outperform softmax transformers while using half the number of heads, having less parameters, and requiring less FLOPs (see Figure 4). Results are averaged over 5 runs.

Model	ListOps	Text	Retrieval	Average
<i>Softmax 12 heads</i>	36.64	65.62	82.18	61.48
<i>Softmax 8 heads</i>	37.03	65.71	81.74	61.49
sMGK 4 heads	37.25	65.51	82.79	61.85
MGK 4 heads	36.98	65.69	82.23	61.63
<i>Softmax 4 heads</i>	36.89	65.26	81.54	61.23
sMGK 2 heads	37.35	65.17	82.20	61.57
MGK 2 heads	36.88	65.37	81.83	61.36
<i>Softmax 2 heads</i>	36.76	64.90	79.1	60.25
sMGK 1 head	37.31	65.04	81.23	61.19
MGK 1 head	37.13	65.40	80.63	61.05
Softmax 1 head	36.81	64.48	77.9	59.73

significantly outperforms the baseline softmax and linear transformers. Especially in the case of Transformer-MLK, it helps reduce the performance gap between softmax and linear transformers.

Throughout this section, we compare Transformer-MGK/MLK with the softmax and linear transformers that have the same or double the number of attention heads. In all experiments, for our Transformer-MGK/MLK models, we set $M=2$ where M is the number of Gaussians, i.e. keys, at each timestep. Among the design options for Transformer-MGK mentioned in Section 2.3, we use the one with Soft-E step but make the parameter π_{jr} and k_{jr} learnable and fix the variance σ_{jr}^2 to be constants. We study both implementations for keys: (A) k_{jr} is a linear projection of the input x_j , i.e., $k_{jr} = x_j \mathbf{W}_{K_r}^\top$ and (B) k_{jr} are shifted version of each other, i.e., $k_{jr} = x_j \mathbf{W}_K^\top + b_r$.

In this section, we refer to the Transformer-MGK/MLK whose keys are implemented by (A) as Transformer-MGK/MLK, and whose keys are implemented by (B) as Transformer-sMGK/sMLK. We empirically compare these models with other design options for Transformer-MGK in Section 3.4. Details on datasets, models, and training are provided in Appendix A.1. Our PyTorch (Paszke et al., 2019) code are available at <https://github.com/minhtannnguyen/transformer-mgk>.

3.1. Long Range Arena (LRA) Benchmark

Models and baselines We compare our 1-head, 2-head, 4-head Transformer-MGK and MLK with the baseline softmax (Vaswani et al., 2017) and linear transformers (Katharopoulos et al., 2020) that have 1 head, 2 heads, 4 heads, and 8 heads. Each model consists of two layers, and we adopt the model and training setting from (Xiong et al., 2021) in our experiments.

Results We summarize our results in Table 1. Transformer-

Table 2. Test Accuracy (%) of Transformer-MLK compared with the linear transformer on the LRA. Our Transformer-MLKs achieve comparable/better accuracy than the baselines while using half the number of heads, having less parameters, and requiring less FLOPs (see Figure 4 for details). Results are averaged over 5 runs.

Model	ListOps	Text	Retrieval	Average
<i>Linear 12 heads</i>	20.26	65.87	81.97	56.03
<i>Linear 8 heads</i>	19.17	65.85	81.18	55.40
sMLK 4 heads	20.11	65.74	81.53	55.79
MLK 4 heads	20.06	65.7	81.34	55.7
<i>Linear 4 heads</i>	19.37	65.81	81.65	55.61
sMLK 2 heads	19.88	65.61	81.66	55.71
MLK 2 heads	20.12	65.72	80.80	55.54
<i>Linear 2 heads</i>	18.35	65.94	80.94	55.07
sMLK 1 head	18.87	65.57	80.37	54.93
MLK 1 head	18.34	65.70	81.09	55.04
<i>Linear 1 head</i>	18.60	65.70	80.6	54.96

MGKs with half the number of heads consistently achieve better test accuracy than the baseline softmax attention across tasks. Since fewer heads are needed, transformer-MGKs use less parameters and need less FLOPs to compute than the baselines. We provide a detailed efficiency analysis for Transformer-MGKs in Figure 4. More interestingly, these efficiency advantages of Transformer-MGK over the baseline become more significant as the number of heads in the baseline model grows. When using the same number of heads as the baseline models, Transformer-MGKs further improve over those baselines. Among the models, Transformer-sMGK performs the best across LRA tasks.

We also compare the performance of Transformer-MLK with the baseline linear transformers in Table 2. Like Transformer-MGK, Transformer-MLK yields comparable or better results than the baseline using only half the number of heads with less parameters and FLOPs. When using the same number of heads, Transformer-MLK helps improve the linear transformer further.

We provide results of the 12-head baselines in Table 1 and 2 for reference. It is interesting to notice from Table 1 and 2 that even our 2-head Transformer-MGK/MLK models achieve better or equivalent results to the 12-head and 8-head baselines. A comparison between the 12-head baselines with our 6-head Transformer-MGK/MLK models on the retrieval task is provided in Table 9 in Appendix A.10.

In Figure 1, we compare the training loss and test accuracy curves of our 1-head and 2-head Transformer-MGK/MLK with the 2-head softmax and 2-head linear transformers on the document retrieval task. This retrieval task has the longest average sequence-length and attention span among the LRA tasks (Tay et al., 2021). On this task, as shown in Figure 1, our Transformer-MGKs/MLKs are always better than the baseline models throughout the training. This observation corroborates our models’s capability of captur-

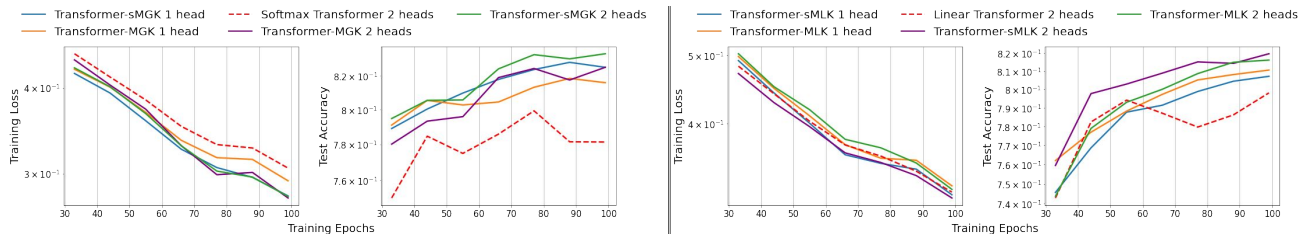


Figure 1. Training loss and test accuracy of Transformer-MGK/MLK vs. softmax/linear transformer on the retrieval task, which has the longest average sequence-length and attention span among the LRA tasks (Tay et al., 2021). The impressive performance of Transformer-MGK/MLK on this challenging task validates the capability of our models to capture long-range dependencies via learning a diversity of attention patterns.

ing long-range dependencies in very long input sequences.

3.2. Language Modeling on WikiText-103

Next we confirm the advantage of our models on a large-scale application. We consider the word-level language modeling task on WikiText-103 (Merity et al., 2017) for our experiments in this section.

Models and baselines We compare 4 and 8-head Transformer-MGKs/MLKs with 8-head softmax (Vaswani et al., 2017) and linear transformers (Katharopoulos et al., 2020). Each model consists of 16 layers. Our experiments follow the setting for small/medium models from (Schlag et al., 2021).

Results As shown in Table 3, our Transformer-MGKs outperform the baseline softmax transformers. Even when using half the number of attention heads (i.e., 4 vs. 8 heads as in the baselines), the Transformer-MGK still achieves better test perplexities (PPL) than the baseline. Adding more heads into Transformer-MGKs improves their performance. Similarly, Transformer-MLKs attain comparable test/validation PPL to the baseline linear transformers when using half the number of attention heads. When using the same number of attention heads as in the baseline, Transformer-MLKs consistently achieve better performance. Note that reducing the number of heads from 8 to 4 in the baseline models significantly decreases their performance with more than 1.5 reduction in test/validation PPL for the softmax transformer and more than 1.0 reduction in test/validation PPL for the linear transformer. Our proposed Transformer-MGK and Transformer-MLK helps close this gap.

To further examine the scalability of our models, we apply the MGK on a stronger baseline, which is the 8-head medium softmax transformer in (Schlag et al., 2021). This model has 90M parameters, 16 layers, 8 attention heads per layer, and hidden size of 256. The size of our baseline model is close to BERT_{Base} (Devlin et al., 2019), which has 110M parameters, 12 layers, 12 attention heads per layer, and hidden size of 768. Applying our MGK on top of this baseline and using only 4 heads instead of 8, we significantly improve the test PPL from 29.60 to 28.86 while reducing

Table 3. Perplexity (PPL) on WikiText-103 of Transformer-MGK and MLK compared to the baselines. Both Transformer-MGK and MLK achieve comparable or better PPL than the baselines while using only half the number of heads. When using the same number of heads, our models significantly improve the baselines.

Method	Valid PPL	Test PPL
<i>Softmax 8 heads (small)</i>	33.15	34.29
MGK 4 heads (small)	33.28	34.21
sMGK 8 heads (small)	32.92	33.99
MGK 8 heads (small)	32.74	33.93
<i>Softmax 4 heads (small)</i>	34.80	35.85
<i>Linear 8 heads (small)</i>	38.07	39.08
MLK 4 heads (small)	38.49	39.46
MLK 8 heads (small)	37.78	38.99
<i>Linear 4 heads (small)</i>	39.32	40.17
<i>Softmax 8 heads (medium)</i>	27.90	29.60
MGK 4 heads (medium)	27.58	28.86

the model size and computational cost, demonstrating the advantages of our scaled models.

3.3. Neural Machine Translation on IWSLT’14 German to English and WMT’14

We further examine the advantages of our methods on the IWSLT’14 German-English machine translation task (Cetolo et al., 2014). Table 4 shows that the 2-head Transformer-MGK and sMGK models achieve the BLEU score of 34.34 and 34.69, respectively, which are comparable to and better than the BLEU score of 34.42 of the 4-head softmax transformer baseline.

We also compare the 8-head Transformer-MGK/sMGK with the 16-head softmax baseline on the WMT’14 task in Table 4. This baseline consists 12 layers with 201M trainable parameters. Transformer-MGK/sMGK obtain the BLEU score of 29.03 and 29.07, respectively, which are comparable with the baseline BLEU score of 29.38.

3.4. Empirical Analysis

We conduct empirical analysis based on the Transformer-MGK trained for the document retrieval task the Language modeling task on WikiText-103, and the IWSLT14 De-En

Table 4. Machine translation BLEU scores of 2-head Transformer-MGKs on the IWSLT14 De-En task are better than or equivalent to that of the 4-head baseline. Similarly, the BLEU scores of 8-head Transformer-MGKs on the WMT’14 task are comparable to that of the 16-head baseline.

Method	Task	BLEU score
<i>Softmax 4 heads</i>	IWSLT’14	34.42
Transformer sMGK 2 head	IWSLT’14	34.69
Transformer MGK 2 head	IWSLT’14	34.34
<i>Softmax 16 heads</i>	WMT’14	29.38
Transformer sMGK 8 head	WMT’14	29.07
Transformer MGK 8 head	WMT’14	29.03

machine translation task. Results for Transformer-MLKs and other other results for WikiText-103 task are provided in the Appendix.

The ability of MGK to approximate the probability distribution of queries We compute the negative log-likelihood (NLL) score of MGK in comparison with the baseline that uses only 1 Gaussian distribution centered around the key k_j , $j = 1, \dots, N$. Both models are trained for the document retrieval task. The NLLs of MGK are 0.180 and 0.182, which are smaller than the baseline’s NLLs of 0.187 and 0.292, in layer 1 and 2, respectively, indicating that MGK fits the distribution of the attention queries better.

Ablation study on the variances of the MGK In our experiments, we set $M=2$ and $\sigma_{j1}^2=\sqrt{D}$ as in the baseline softmax transformer and $\sigma_{j2}^2=3\sqrt{D}$ where D is the feature dimension defined in Section 1.1. We conduct an ablation study on the IWSLT14 De-En machine translation task in which we vary σ_{j2}^2 . We observe that the BLEU score of the trained models do not differ much between different values of σ_{j2}^2 . In particular, BLEU scores are 34.31, 34.36, 34.61, 34.34, and 34.27 for $\sigma_{j2}^2 = 0.5\sqrt{D}, \sqrt{D}, 2\sqrt{D}, 3\sqrt{D}$ and $4\sqrt{D}$, respectively. We also make σ_{j1}^2 and σ_{j2}^2 learnable and observe that the resultant BLEU score is worse (32.86).

Transformer-MGK helps avoid learning redundant heads We visually compare attention matrices learned by Transformer-MGKs and the baseline softmax transformer on the document retrieval task in Figure 2. In particular, we randomly select an attention matrix at each head in each layer and visualize that attention matrix for each model in comparison. Figure 2(Left) shows that the queries in Transformer-MGKs can attend to a variety of keys and equivalently to other tokens at different positions in the input sequence. This diversity in attention pattern helps reduce the chance that the model learns similar and redundant attention matrices at different heads significantly.

Another metric to measure the representation capacity of an attention matrix is its rank. Attention matrices with high ranks can capture more diverse attention patterns compared to those with low ranks (Nguyen et al., 2021b). We study

Table 5. Performance of Transformer-MGK using different inference/learning techniques on LRA benchmark.

Model	ListOps	Text	Retrieval	Average
sMGK + Hard-E 1 head	37.25	64.7	81.29	61.08
sMGK + Soft-E 1 head	37.05	64.68	81.44	61.05
sMGK 1 head	37.31	65.04	81.23	61.19
MGK + Hard-E 1 head	19.40	65.40	80.72	55.17
MGK + Soft-E 1 head	33.85	65.25	80.73	59.94
MGK 1 head	37.13	65.40	80.63	61.05

the rank of the attention matrix from the Transformer-MGK and the softmax transformer trained for the retrieval task. In particular, we randomly select 1000 different attention matrices at each layer from each model. Then, we perform singular value decomposition (SVD) to compute the rank of each matrix and threshold the singular values smaller than 10^{-6} . Figure 2(Right) presents the distribution of the rank of attention matrices at each layer of the Transformer-MGK and the softmax transformer. We observe that attention matrices in Transformer-MGK has higher rank than those in the softmax transformer. Thus, our attention with MGK is capable of capturing more diverse and complex attention patterns than the baseline softmax attention.

Comparing different inference and learning techniques

Table 5 compares the performance of Transformer-MGKs using different design options mentioned in Section 2.3 on the LRA benchmark. In particular, we consider the following three design options: A) Soft-E step, parameters π_{jr} and k_{jr} are learnable via SGD, and variance σ_{jr}^2 are constants, B) Soft-E step, parameter π_{jr} is updated according to the M-step update, k_{jr} are learnable via SGD, and variance σ_{jr}^2 are constants, and C) Hard-E step, π_{jr} and k_{jr} are learnable via SGD, and variance σ_{jr}^2 are constants. Note that Transformer-MGKs with setting A are the default models we use in all experiments above. In Table 5, Transformer-MGK + Hard-E is the Transformer-MGK with setting C, Transformer-MGK + Soft-E is the Transformer-MGK with setting B, and Transformer-MGK only is the Transformer-MGK with setting A. It is worth noting that Transformer-sMGK + Hard-E obtains comparable results to the models with the best performance in each task even though it is the most efficient model in our study.

Transformer-MGK reduces model complexity and computational cost

Figure 3A and 3B present the reduction ratio of train and test FLOPS, respectively, of our 4-head Transformer-MGK vs. the 8-head Softmax baseline as functions of model dimension D and sequence length N . In Fig. 3C, we show the reduction ratio of model size of Transformer-MGK vs. the baseline as D varies. We observe that the efficiency advantage of Transformer-MGK over the baseline grows with D and N , making it more suitable and superior for large-scale applications. Note that the model size in term of the number of parameters does not depend

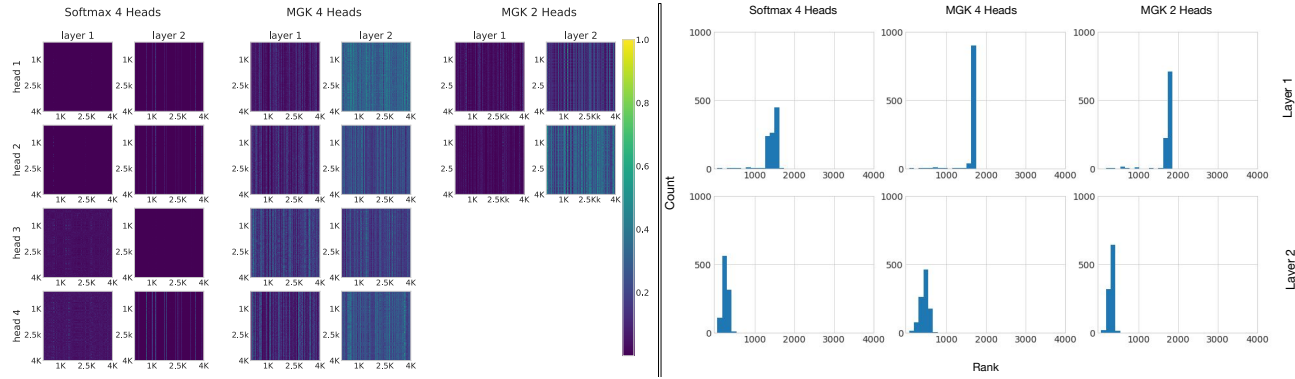


Figure 2. (Left) Visualization of attention matrices in the baseline 4-head softmax transformer (left), 4-head Transformer-MGK (middle), and 2-head Transformer-MGK (right) trained on the document retrieval task. Attention matrices from our Transformer-MGKs have more diverse pattern than those from the baseline softmax transformer, reducing the risk of learning redundant heads. (Right) Rank distribution of attention matrices shows that attention matrices in Transformer-MGK have higher rank than those in the softmax transformer and thus can capture more diverse attention patterns.

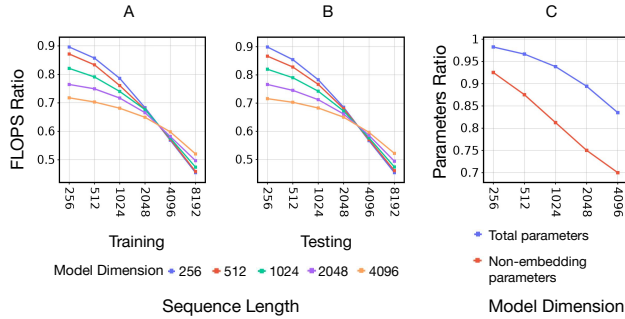


Figure 3. Training (A), Inference (B) FLOP ratios and number-of-parameter ratio (C) between 4-head Transformer-MGK with the 8-head softmax baseline across different D model dimensions and N sequence lengths, for the language modeling task on WikiText-103. 4-head Transformer-MGK is significantly efficient, both in computation and model complexity, than the baseline as D and N increases, indicating the benefits of our method for long-range and large-scale tasks.

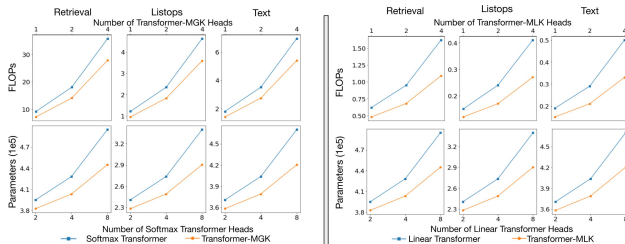


Figure 4. Computational cost (FLOPs) and the number of parameters of Transformer-MGK vs the baseline softmax transformer (Left) and Transformer-MLK vs. the baseline linear transformer (Right) trained on the Retrieval task. The efficiency advantage of Transformer-MGK/MLK over the baselines in both metrics grows with the number of head.

on the sequence length N . The efficiency analysis results on this language modeling task for Transformer-MLK and the analysis metrics are provided in Fig. 10 in Appendix A.5.

Also, Figure 4 compares the computational cost, measured

in FLOPS, and model complexity, measured in the number of parameters, between our Transformer-MGK/MLK that has half the number of heads and the full-head softmax/linear transformer as functions of the number of heads. The models are trained for the Retrieval task in the LRA benchmark. The more heads being used, the more advantage Transformer-MGK/MLK has over the softmax/linear transformer. For much larger transformer models, this saving is significant.

4. Related Work

Efficient Transformers Efficient transformers can be classified into several categories, as summarized in (Roy et al., 2021). Among these categories are models which design the attention matrix to have sparse structure (Parmar et al., 2018; Liu et al., 2018; Qiu et al., 2019; Child et al., 2019; Beltagy et al., 2020; Du et al., 2021). Another category includes models that combine two or more different access patterns to improve the coverage (Child et al., 2019; Ho et al., 2019). Access patterns can also be made learnable in a data-driven fashion (Kitaev et al., 2020; Roy et al., 2021; Tay et al., 2020). Other efficient transformers take advantage of a side memory module to access multiple tokens at once (Lee et al., 2019; Sukhbaatar et al., 2019; Asai & Choi, 2020; Beltagy et al., 2020). Inspired by the use of momentum methods in speeding up neural network training (Sutskever et al., 2013; Bengio et al., 2013; Kingma & Ba, 2015; Wang et al., 2022) and designing neural network architectures (Li et al., 2018; He et al., 2020; Nguyen et al., 2020; Xia et al., 2021; Sander et al., 2021), (Wang et al., 2021a) incorporates adaptive momentum into the linear transformers to improve its accuracy while maintaining the linear computational and memory complexity. Furthermore, dynamic inference previously proposed for deep neural networks (Wang et al., 2018a;b; 2020b; Han et al., 2021) can also be used to accelerate the inference in transformers (Bakhtiarnia et al.,

2021; Wang et al., 2021b). Low-rank and kernelization approximation are also utilized to enhance the memory and computational efficiency of computing self-attention (Tsai et al., 2019; Wang et al., 2020a; Katharopoulos et al., 2020; Choromanski et al., 2021; Shen et al., 2021; Nguyen et al., 2021b; Peng et al., 2021). In addition to the aforementioned efficient transformers, multi-query attention that shares keys and values between different attention heads (Shazeer, 2019) has also been studied to reduce the memory-bandwidth cost and increase the speed for incremental transformer inference (see Appendix A.11). It is worth noting that efficient retrieval methods have been employed to increase the efficiency of machine learning models including transformers (Chen et al., 2020; Le & Lauw, 2020; 2021; Mittal et al., 2022). Last but not least, methods such as using auxiliary losses (Al-Rfou et al., 2019) and adaptive input embedding (Baevski & Auli, 2019) have been explored to speed up the convergence of training transformers. Our MGK/MLK can be easily incorporated into these methods above to further improve their accuracy and efficiency.

Redundancy in Transformers Latest works suggest that most of the neurons and heads in the pre-trained transformer are redundant and can be removed when optimizing towards a downstream task (Dalvi et al., 2020; Michel et al., 2019; Durrani et al., 2020). Other works also study the contextualized embeddings in pretrained networks under this redundancy due to overparameterization and show that the representations learned within these models are highly anisotropic (Mu & Viswanath, 2018; Ethayarajh, 2019). An emerging body of work is proposed to distill and prune the model, including (Sanh et al., 2019; Sun et al., 2019; Voita et al., 2019b; Sajjad et al., 2020). Our MGK/MLK approach can be combined with these distilling and pruning methods.

Mixture Models for Transformers Several works have used mixture models to study and enhance transformers. Switch transformers (Fedus et al., 2021) employ the routing algorithm in Mixture of Experts (MoE) to reduce the communication and computational costs in transformers. (Nguyen et al., 2018; Patel et al., 2016; 2015) derive a probabilistic framework based on GMMs for deep neural networks that can be extended to study transformers and attention-based architectures. (Zhang & Feng, 2021) develops a Gaussian mixture attention that models each attention score as a GMM. In contrast, the MGK models each key k_j at timestep j as a GMM. In other words, our MGK is a generative model approach in which we build a generative model for the key k_j and derive the attention score from the posterior distribution while the Gaussian mixture attention in (Zhang & Feng, 2021) is a discriminative approach (Bishop, 2006). Other works that use mixture models with deep neural networks and transformers include (Cho et al., 2020; Guo et al., 2019; Huang et al., 2020).

5. Concluding Remarks

In this paper, we proposed Transformer-MGK, a class of transformers that use Gaussian mixture model to represent the key vectors in self-attention. Transformer-MGK reduces the redundancy among heads in transformer. Furthermore, attention heads in the Transformer-MGK have better representation capability than those in the baseline, allowing the Transformer-MGK to achieve comparable or better performance than the baseline softmax transformer while using only half of the number of heads. Comparing to the baseline, the Transformer-MGK uses fewer parameters and requires less FLOPs to compute. We extend the Transformer-MGK into the Transformer-MLK to use linear attentions for better efficiency. We empirically validate the advantage of the Transformer-MGK/MLK over the baseline softmax and linear transformers on various benchmarks including tasks in the LRA benchmark, WikiText-103 language modeling, and IWSLT’14/WMT’14 machine translation. Both Transformer-MGK/MLK and the softmax/linear transformers assume that features in the attention queries and keys are independent. (Nguyen et al., 2022) suggests that the Fourier integral estimators can be used to efficiently capture the dependence structure between those embedding features. The Fourier integral estimators can be incorporated into Transformer-MGK/MLK to increase the representation power of the models. In our work, we make the means and the variance of the cluster learnable variables and constants, respectively. It is interesting to explore how to leverage the M-step update in the EM algorithm to update those parameters. Since the EM algorithm can have sub-linear convergence rate due to the weak separation of the means and variance (Dwivedi et al., 2020b;a), we may need to develop more efficient optimization algorithms than the EM for estimating the means and variance. The current works of (Ho et al., 2022; Ren et al., 2022) demonstrate that an exponential learning rate schedule for the gradient descent can be utilized for the means and variance to obtain the linear convergence rate of the parameters. It is of practical importance to investigate the performance of that algorithm in the Transformer-MGK/MLK. Finally, we leave the application of MGK/MLK for improving the vision transformer (Dosovitskiy et al., 2020; Touvron et al., 2020) as future work.

Acknowledgements

This material is based on research sponsored by the AFOSR MURI FA9550-18-1-0502, the ONR grant N00014-20-1-2093, the MURI N00014-20-1-2787, and the NSF under Grant# 2030859 to the Computing Research Association for the CIFellows Project (CIF2020-UCLA-38). NH acknowledges support from the NSF IFML 2019844 and the NSF AI Institute for Foundations of Machine Learning. RGB was supported by NSF grants CCF-1911094, IIS-1838177, IIS-1730574, and a Vannevar Bush Faculty Fellowship.

References

- Al-Rfou, R., Choe, D., Constant, N., Guo, M., and Jones, L. Character-level language modeling with deeper self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3159–3166, 2019.
- Asai, A. and Choi, E. Challenges in information seeking qa: Unanswerable questions and paragraph retrieval. *arXiv preprint arXiv:2010.11915*, 2020.
- Bacharoglou, A. G. Approximation of probability distributions by convex mixtures of Gaussian measures. *Proceedings of the American Mathematical Society*, 138: 2619–2628, 2010.
- Baevski, A. and Auli, M. Adaptive input representations for neural language modeling. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ByxZX20qFQ>.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bakhtiarnia, A., Zhang, Q., and Iosifidis, A. Multi-exit vision transformer for dynamic inference. *arXiv preprint arXiv:2106.15183*, 2021.
- Beltagy, I., Peters, M. E., and Cohan, A. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Bengio, Y., Boulanger-Lewandowski, N., and Pascanu, R. Advances in optimizing recurrent networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pp. 8624–8628. IEEE, 2013.
- Bishop, C. M. Pattern recognition. *Machine learning*, 128 (9), 2006.
- Brown, T. and et al. Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901, 2020.
- Cettolo, M., Niehues, J., Stüker, S., Bentivogli, L., and Federico, M. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of the International Workshop on Spoken Language Translation, Hanoi, Vietnam*, volume 57, 2014.
- Chen, B., Liu, Z., Peng, B., Xu, Z., Li, J. L., Dao, T., Song, Z., Shrivastava, A., and Re, C. Mongoose: A learnable lsh framework for efficient neural network training. In *International Conference on Learning Representations*, 2020.
- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. 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)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL <https://www.aclweb.org/anthology/D14-1179>.
- Cho, S. M., Park, E., and Yoo, S. Meantime: Mixture of attention mechanisms with multi-temporal embeddings for sequential recommendation. In *Fourteenth ACM Conference on Recommender Systems*, pp. 515–520, 2020.
- Choromanski, K. M., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J. Q., Mohiuddin, A., Kaiser, L., Belanger, D. B., Colwell, L. J., and Weller, A. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ua6zuk0WRH>.
- Clark, K., Khandelwal, U., Levy, O., and Manning, C. D. What does BERT look at? an analysis of BERT’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 276–286, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4828. URL <https://www.aclweb.org/anthology/W19-4828>.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q. V., and Salakhutdinov, R. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- Dalvi, F., Sajjad, H., Durrani, N., and Belinkov, Y. Analyzing redundancy in pretrained transformer models. *arXiv preprint arXiv:2004.04010*, 2020.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, L. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Du, N., Huang, Y., Dai, A. M., Tong, S., Lepikhin, D., Xu, Y., Krikun, M., Zhou, Y., Yu, A. W., Firat, O., et al. Glam: Efficient scaling of language models with mixture-of-experts. *arXiv preprint arXiv:2112.06905*, 2021.
- Durrani, N., Sajjad, H., Dalvi, F., and Belinkov, Y. Analyzing individual neurons in pre-trained language models. *arXiv preprint arXiv:2010.02695*, 2020.
- Dwivedi, R., Ho, N., Khamaru, K., Wainwright, M. J., Jordan, M. I., and Yu, B. Sharp analysis of expectation-maximization for weakly identifiable models. *AISTATS*, 2020a.
- Dwivedi, R., Ho, N., Khamaru, K., Wainwright, M. J., Jordan, M. I., and Yu, B. Singularity, misspecification, and the convergence rate of EM. *Annals of Statistics*, 44: 2726–2755, 2020b.
- Ethayarajh, K. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*, 2019.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- Ghosal, S. and van der Vaart, A. Entropies and rates of convergence for maximum likelihood and bayes estimation for mixtures of normal densities. *Ann. Statist.*, 29: 1233–1263, 2001.
- Guo, M., Zhang, Y., and Liu, T. Gaussian transformer: a lightweight approach for natural language inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 6489–6496, 2019.
- Han, Y., Huang, G., Song, S., Yang, L., Wang, H., and Wang, Y. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Hanson, S. J. A stochastic version of the delta rule. *Physica D: Nonlinear Phenomena*, 42(1-3):265–272, 1990.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- Hewitt, J. and Liang, P. 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)*, pp. 2733–2743, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1275. URL <https://www.aclweb.org/anthology/D19-1275>.
- Ho, J., Kalchbrenner, N., Weissenborn, D., and Salimans, T. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
- Ho, N., Ren, T., Sanghavi, S., Sarkar, P., and Ward, R. An exponentially increasing step-size for parameter estimation in statistical models. *arXiv preprint arXiv: 2205.07999*, 2022.
- Howard, J. and Ruder, S. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 328–339, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1031. URL <https://www.aclweb.org/anthology/P18-1031>.
- Huang, Y., Gornet, J., Dai, S., Yu, Z., Nguyen, T., Tsao, D., and Anandkumar, A. Neural networks with recurrent generative feedback. *Advances in Neural Information Processing Systems*, 33:535–545, 2020.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pp. 5156–5165. PMLR, 2020.
- Kim, Y., Denton, C., Hoang, L., and Rush, A. M. Structured attention networks. *arXiv preprint arXiv:1702.00887*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Kitaev, N., Kaiser, L., and Levskaya, A. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.
- Le, D. D. and Lauw, H. Efficient retrieval of matrix factorization-based top-k recommendations: A survey of recent approaches. *Journal of Artificial Intelligence Research*, 70:1441–1479, 2021.

- Le, D. D. and Lauw, H. W. Stochastically robust personalized ranking for lsh recommendation retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 4594–4601, 2020.
- Lee, J., Lee, Y., Kim, J., Kosiorek, A., Choi, S., and Teh, Y. W. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, pp. 3744–3753. PMLR, 2019.
- Li, H., Yang, Y., Chen, D., and Lin, Z. Optimization algorithm inspired deep neural network structure design. In *Asian Conference on Machine Learning*, pp. 614–629. PMLR, 2018.
- Lin, Z., Feng, M., dos Santos, C. N., Yu, M., Xiang, B., Zhou, B., and Bengio, Y. A structured self-attentive sentence embedding. *CoRR*, abs/1703.03130, 2017.
- Liu, P. J., Saleh, M., Pot, E., Goodrich, B., Sepassi, R., Kaiser, L., and Shazeer, N. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*, 2018.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Lloyd, S. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/P11-1015>.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Byj72udxe>.
- Michel, P., Levy, O., and Neubig, G. Are sixteen heads really better than one? In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Mittal, S., Rappaport, S. C., Rish, I., Bengio, Y., and Lajoie, G. Compositional attention: Disentangling search and retrieval. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=IwJPj2MBcIa>.
- Mu, J. and Viswanath, P. All-but-the-top: Simple and effective postprocessing for word representations. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HkuGJ3kCb>.
- Nangia, N. and Bowman, S. ListOps: A diagnostic dataset for latent tree learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pp. 92–99, New Orleans, Louisiana, USA, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-4013. URL <https://www.aclweb.org/anthology/N18-4013>.
- Nguyen, T., Ho, N., Patel, A., Anandkumar, A., Jordan, M. I., and Baraniuk, R. G. A Bayesian perspective of convolutional neural networks through a deconvolutional generative model. *arXiv preprint arXiv:1811.02657*, 2018.
- Nguyen, T., Baraniuk, R., Bertozzi, A., Osher, S., and Wang, B. Momentumrnn: Integrating momentum into recurrent neural networks. *Advances in Neural Information Processing Systems*, 33:1924–1936, 2020.
- Nguyen, T., Nguyen, P., Pham, H., Bui, T., Nguyen, T., and Luong, D. Sp-gpt2: Semantics improvement in vietnamese poetry generation. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1576–1581. IEEE, 2021a.
- Nguyen, T., Pham, M., Nguyen, T., Nguyen, K., Osher, S. J., and Ho, N. Transformer with Fourier integral attentions. *arXiv preprint arXiv:2206.00206*, 2022.
- Nguyen, T. M., Suliafu, V., Osher, S. J., Chen, L., and Wang, B. Fmmformer: Efficient and flexible transformer via decomposed near-field and far-field attention. *arXiv preprint arXiv:2108.02347*, 2021b.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- Parikh, A., Täckström, O., Das, D., and Uszkoreit, J. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2249–2255, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1244. URL <https://www.aclweb.org/anthology/D16-1244>.
- Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., and Tran, D. Image transformer. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of

- Proceedings of Machine Learning Research*, pp. 4055–4064. PMLR, 10–15 Jul 2018. URL <http://proceedings.mlr.press/v80/parmar18a.html>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Patel, A. B., Nguyen, T., and Baraniuk, R. G. A probabilistic theory of deep learning. *arXiv preprint arXiv:1504.00641*, 2015.
- Patel, A. B., Nguyen, M. T., and Baraniuk, R. A probabilistic framework for deep learning. *Advances in neural information processing systems*, 29:2558–2566, 2016.
- Peng, H., Pappas, N., Yogatama, D., Schwartz, R., Smith, N., and Kong, L. Random feature attention. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=QtTKTdVrFBB>.
- Qiu, J., Ma, H., Levy, O., Yih, S. W.-t., Wang, S., and Tang, J. Blockwise self-attention for long document understanding. *arXiv preprint arXiv:1911.02972*, 2019.
- Radev, D. R., Muthukrishnan, P., Qazvinian, V., and Abu-Jbara, A. The acl anthology network corpus. *Language Resources and Evaluation*, 47(4):919–944, 2013.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. *OpenAI report*, 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264. URL <https://www.aclweb.org/anthology/D16-1264>.
- Ren, T., Cui, F., Sanghavi, S., and Ho, N. Beyond EM algorithm on over-specified two-component location-scale Gaussian mixtures. *arXiv preprint arXiv: 2205.11078*, 2022.
- Roy, A., Saffar, M., Vaswani, A., and Grangier, D. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9:53–68, 2021. doi: 10.1162/tacl.a.00353. URL <https://www.aclweb.org/anthology/2021.tacl-1.4>.
- Sajjad, H., Dalvi, F., Durrani, N., and Nakov, P. Poor man’s bert: Smaller and faster transformer models. *arXiv e-prints*, pp. arXiv–2004, 2020.
- Sander, M. E., Ablin, P., Blondel, M., and Peyré, G. Momentum residual neural networks. In *International Conference on Machine Learning*, pp. 9276–9287. PMLR, 2021.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- Schlag, I., Irie, K., and Schmidhuber, J. Linear transformers are secretly fast weight programmers. In *International Conference on Machine Learning*, pp. 9355–9366. PMLR, 2021.
- Shazeer, N. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*, 2019.
- Shen, Z., Zhang, M., Zhao, H., Yi, S., and Li, H. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 3531–3539, 2021.
- So, D. R., Liang, C., and Le, Q. V. The evolved transformer. *arXiv preprint arXiv:1901.11117*, 2019.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Sukhbaatar, S., Grave, E., Lample, G., Jegou, H., and Joulin, A. Augmenting self-attention with persistent memory. *arXiv preprint arXiv:1907.01470*, 2019.
- Sun, S., Cheng, Y., Gan, Z., and Liu, J. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pp. 1139–1147. PMLR, 2013.
- Tay, Y., Bahri, D., Yang, L., Metzler, D., and Juan, D.-C. Sparse Sinkhorn attention. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning Research*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9438–9447. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/tay20a.html>.
- Tay, Y., Dehghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D.

- Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=qVyeW-grC2k>.
- Tenney, I., Das, D., and Pavlick, E. BERT rediscovers the classical NLP pipeline. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4593–4601, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1452. URL <https://www.aclweb.org/anthology/P19-1452>.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., and Jégou, H. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020.
- Tsai, Y.-H. H., Bai, S., Yamada, M., Morency, L.-P., and Salakhutdinov, R. Transformer dissection: An unified understanding for transformer’s attention via the lens of kernel. *arXiv preprint arXiv:1908.11775*, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Vig, J. and Belinkov, Y. Analyzing the structure of attention in a transformer language model. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 63–76, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4808. URL <https://www.aclweb.org/anthology/W19-4808>.
- Voita, E., Talbot, D., Moiseev, F., Sennrich, R., and Titov, I. 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*, pp. 5797–5808, Florence, Italy, July 2019a. Association for Computational Linguistics. doi: 10.18653/v1/P19-1580. URL <https://www.aclweb.org/anthology/P19-1580>.
- Voita, E., Talbot, D., Moiseev, F., Sennrich, R., and Titov, I. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. *arXiv preprint arXiv:1905.09418*, 2019b.
- Wang, B., Xia, H., Nguyen, T., and Osher, S. How does momentum benefit deep neural networks architecture design? a few case studies. *arXiv preprint arXiv:2110.07034*, 2021a.
- Wang, B., Nguyen, T., Sun, T., Bertozzi, A. L., Baraniuk, R. G., and Osher, S. J. Scheduled restart momentum for accelerated stochastic gradient descent. *SIAM Journal on Imaging Sciences*, 15(2):738–761, 2022.
- Wang, S., Li, B., Khabsa, M., Fang, H., and Ma, H. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020a.
- Wang, X., Yu, F., Dou, Z.-Y., Darrell, T., and Gonzalez, J. E. Skipnet: Learning dynamic routing in convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 409–424, 2018a.
- Wang, Y., Nguyen, T., Zhao, Y., Wang, Z., Lin, Y., and Baraniuk, R. Energynet: Energy-efficient dynamic inference. 2018b.
- Wang, Y., Shen, J., Hu, T.-K., Xu, P., Nguyen, T., Baraniuk, R., Wang, Z., and Lin, Y. Dual dynamic inference: Enabling more efficient, adaptive, and controllable deep inference. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):623–633, 2020b.
- Wang, Y., Huang, R., Song, S., Huang, Z., and Huang, G. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Williams, A., Nangia, N., and Bowman, S. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122, June 2018. doi: 10.18653/v1/N18-1101. URL <https://www.aclweb.org/anthology/N18-1101>.
- Xia, H., Suliafu, V., Ji, H., Nguyen, T., Bertozzi, A., Osher, S., and Wang, B. Heavy ball neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 34, 2021.
- Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y., and Singh, V. Nyströmformer: A Nyström-based Algorithm for Approximating Self-Attention. 2021.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., and Le, Q. V. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
- Zhang, S. and Feng, Y. Modeling concentrated cross-attention for neural machine translation with Gaussian mixture model. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 1401–1411, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.121. URL <https://aclanthology.org/2021.findings-emnlp.121>.

Supplement to “Transformer with a Mixture of Gaussian Keys”

In this supplementary material, we provide experimental details and additional experiments of Transformer-MGK and Transformer-MLK.

A. Additional Experiments

A.1. Experiment details

In this section, we provide model and training details for experiments in Section 3. All our experiments are conducted on a server with 4 NVIDIA A100 GPUs.

A.1.1. LONG RANGE ARENA BENCHMARK

Datasets and metrics We consider the following tasks in the LRA benchmark: Listops (Nangia & Bowman, 2018), byte-level IMDb reviews text classification (Maas et al., 2011), and byte-level document retrieval (Radev et al., 2013). These tasks involve long sequences of length $2K$, $4K$, and $4K$, respectively. We follow the setup/evaluation protocol in (Tay et al., 2021) and report the test accuracy for individual task and the average result across all tasks.

Models and baselines We use the softmax transformer (Vaswani et al., 2017) and linear transformer (Katharopoulos et al., 2020) as our baselines. All models have 2 layers, 64 embedding dimension, and 128 hidden dimension. The number of heads in each layer are set to 1, 2, 4, and 8. For Transformer-MGK/MLKs and their shifted versions, we share π_{j_r} for all position j and learn it for each head. The initial value for each π_{j_r} is set to 0.5. For Transformer-sMGK, we learn \mathbf{b}_r and initialize its elements from a standard normal distribution. Each σ_{j_r} is a constant with value \sqrt{D} where D is the dimension of each head, which is the same as in the baselines models

Details about the Long Range Arena (LRA) benchmarks can be found in the original paper (Tay et al., 2021). Our implementation is based on the public code by (Xiong et al., 2021), and we follow their training procedures. The training setting and additional baseline model details are provided in the configuration file used in (Xiong et al., 2021) and available at <https://github.com/mlpen/Nystromformer/blob/main/LRA/code>.

A.1.2. LANGUAGE MODELING ON WIKITEXT-103

Datasets and metrics WikiText-103 consists of articles from Wikipedia and is a dataset with long contextual dependencies. The training set is made up of about $28K$ articles containing $103M$ running words; this corresponds to text blocks of about 3600 words. The validation and test sets are composed of $218K$ and $246K$ running words, respectively. Each of them contains 60 articles and about $268K$ words. Our experiment follows the standard setting (Merity et al., 2017; Schlag et al., 2021) and splits the training data into L -word independent long segments. For evaluation, we use a batch size of 1, and go through the text sequence with a sliding window of size L . We consider only the last position for computing perplexity (PPL) except in the first segment, where all positions are evaluated as in (Al-Rfou et al., 2019; Schlag et al., 2021).

Models and baselines Our language modeling implementation is based on the public code <https://github.com/IDSIA/lmtool-fwp> by (Schlag et al., 2021). We use their small and medium model configurations for models in our experiments. In particular, for small models, we set the key, value, and query dimension to 128, and the training and evaluation context length to 256. For medium models, we set the key, value, and query dimension to 256, and the training and evaluation context length to 384. In both configurations, we set the number of heads to 8, the feed-forward layer dimension to 2048, and the number of layers to 16. For Transformer-MGK/MLK, we use our 4 and 8-head versions to compare with the 8-head baselines. π_{i_r} , \mathbf{b}_r and σ_{j_r} for this task follow our settings for LRA experiments. Other than those, our language modeling share the same configurations as the baselines.

We train our models for language modeling on 2 A100, 40GB each with a batch size of 96, and each model is trained for 120 epochs. We apply 10% dropout (Hanson, 1990; Srivastava et al., 2014) and use the Adam optimizer (Kingma & Ba, 2015) with an initial learning rate of 0.00025 and 2000 steps for learning rate warm-up.

A.1.3. NEURAL MACHINE TRANSLATION ON IWSLT’14 GERMAN TO ENGLISH

Datasets and metrics The IWSLT14 German-English dataset (Cettolo et al., 2014) contains $153K$ training, $7K$ validation, and $7K$ test TED-talks scripts German-English translated sentences. We follow the same preprocessing steps as in (Ott

et al., 2019). We use the BiLingual Evaluation Understudy (BLEU) score as our evaluation metric for this task. All trained models are evaluated using the evaluation protocol in (Ott et al., 2019).

Models and baselines The baseline model we use in this machine translation task is an encoder-decoder transformer with six encoder/decoder layers, four attention heads per layer. The embedding dimension and the hidden size are 512 and 1024, respectively, for both encoder and decoder. These architecture configurations are the same for Transformer-MGK/sMGK except for the number of attention heads per layer, which is reduced by half. We share π_{j_r} across all position j and learn it for each head. The initial value for each π_{j_r} is set to 0.5. For Transformer-sMGK, we learn \mathbf{b}_r and initialize its elements from a standard normal distribution. Each σ_{j_r} is a constant with the value \sqrt{D} where D is the dimension of each head, which is the same as in the baselines models. Our experiments follow the settings from (Ott et al., 2019), and our implementation is based on the public code <https://github.com/pytorch/fairseq/tree/main/examples/translation>.

A.2. More Empirical Analysis of Transformer-MGKs/MLKs trained for Language Modeling

In Table 3 in the main text, we show the improvements in valid and test perplexity of our Transformer-MGKs/MLKs compared with the baseline softmax and linear transformers. In particular, Transformer-MGK/MLKs with the same number of heads as the baselines, e.g. 8 heads, significantly improve the baselines during training while Transformer-MGK/MLKs with 4 head achieve comparable or better performance than the 8-head baselines. In this section, we provide more empirical analysis to shed light on those results. Figure 5 shows the validation perplexity curve of our models versus the softmax and linear transformers.

Figure 6 visualizes the attention matrices from a randomly selected sample for the trained softmax transformer with 8 heads and Transformer-MGKs with 8 heads and 4 heads. These visualizations show that Transformer-MGKs attend to more diverse positions in all heads and layers than the softmax attention baseline. We also compare the rank distributions of these attention matrices computed from 1000 samples at each layer in the model. Figure 7 presents the rank histograms of the 8-head softmax attention and 4-head and 8-head MGK attentions for the 1st and 5th layer. It is clear that attention matrices from the Transformer-MGKs have higher ranks than those in the softmax transformer, which implies that Transformer-MGK can attend to more diverse regions than the softmax transformer without the need of using more attention heads.

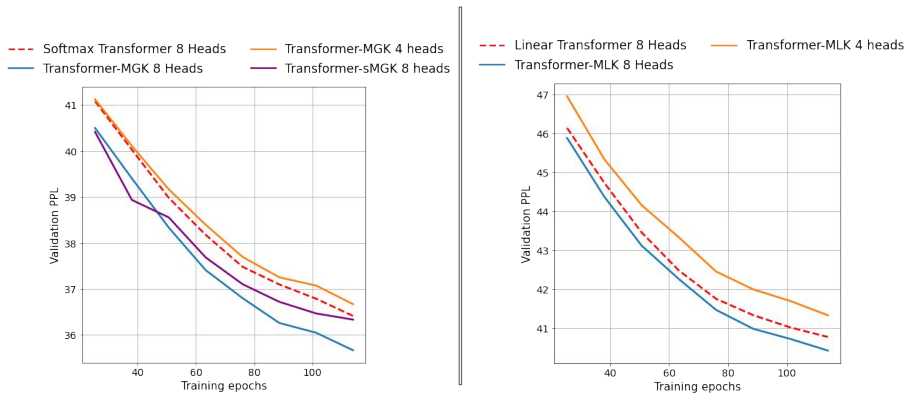


Figure 5. Validation perplexity of Transformer-MGK vs. the softmax transformer (Left) and Transformer-MLK vs. the linear transformer (Right) for language modeling on WikiText-103.

A.3. Additional training results for LRA

In this section, we provide additional experimental results on the LRA benchmark. Figure 8 compares the computational cost measured by FLOPs and model complexity in term of the number of parameters of different inference and learning methods for Transformer-MGK. The computational costs of Transformer-MGK/sMGK and Transformer-MGK Hard-E/Soft-E are on par with each other, while Transformer-sMGK uses fewer parameters than the other without trade-off in performance 5 for all tasks. The naming is as explained in Section 3.4 in the main text. In addition, Figure 9 visualizes the attention matrices in the 4-head linear transformer baseline, 4-head Transformer-MLK, and 2-head Transformer-MLK trained on the document retrieval task.

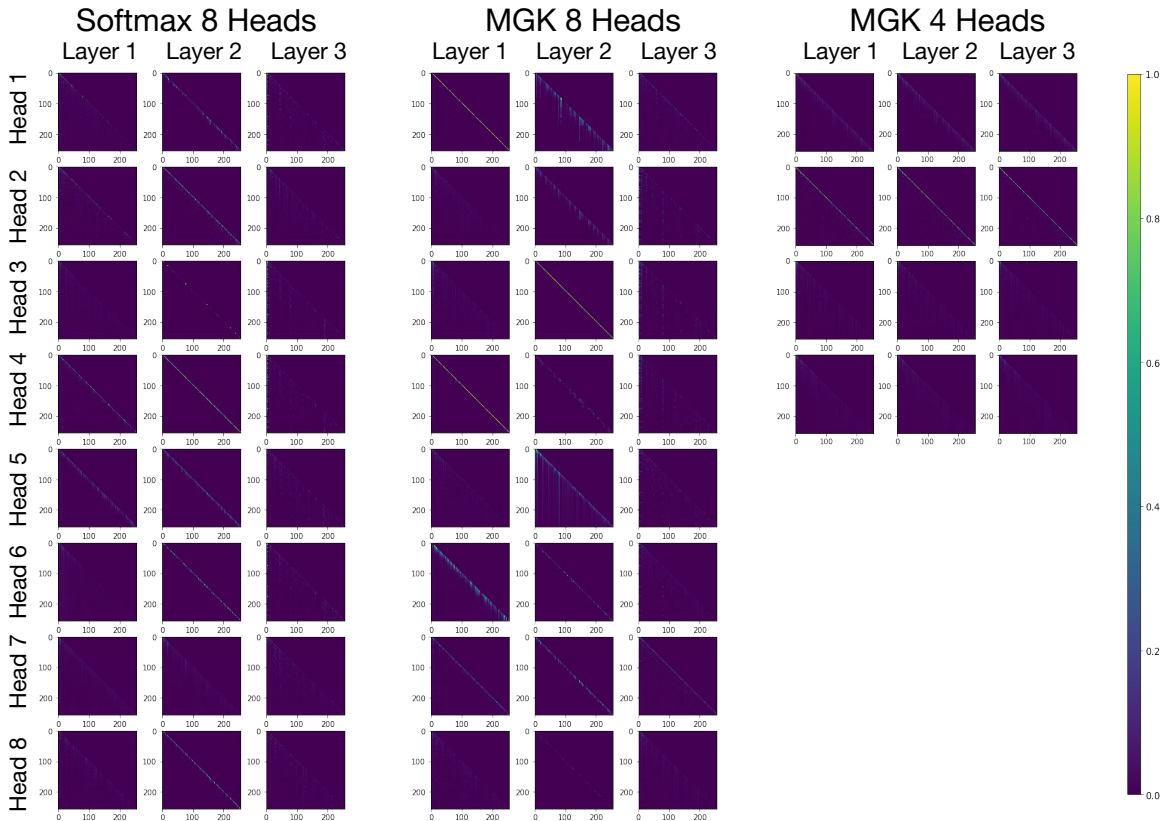


Figure 6. Visualization of attention matrices from 8-head softmax transformer (Left), 8-head Transformer-MGK (Middle), and 4-head Transformer-MGK (Right) trained on WikiText-103 language modeling. Here, we plot the attention matrices for all heads and layers in the models. The sequence length is 256 and the size of each matrix is 256×256 .

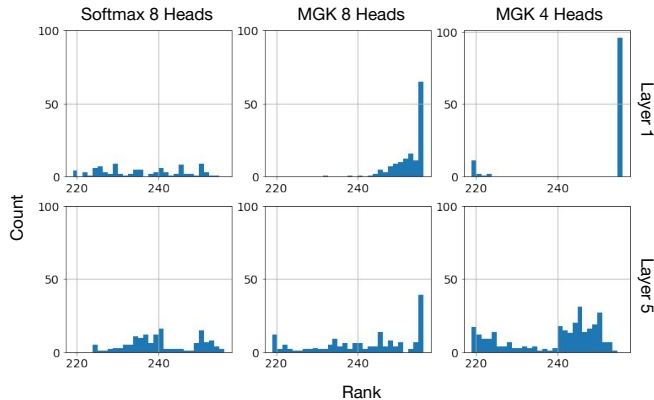


Figure 7. Rank distributions of attention matrices from 8-head softmax transformer (Left), 8-head Transformer-MGK (Middle), and 4-head Transformer-MGK (Right) trained on WikiText-103 language modeling. The rank histograms are computed from 1000 attention matrices at each layer. The attention matrices from Transformer-MGKs have higher ranks than those from softmax transformers. This implies that Transformer-MGKs have more diverse attention patterns, which allows us to reduce the number of heads in Transformer-MGKs.

A.4. Ablation Study on the Impact of the Mixture of Keys, the Gaussian Distance, and the Key Shifting

In this section, we conduct an ablation study of the Transformer-MGK on the LRA retrieval task to investigate where the performance improvement is from. In particular, we would like to understand the impact of the following factors on the performance of Transformer-MGK: 1) the mixture of keys, 2) the Gaussian distance, and 3) the key shifting. We summarize our empirical results in Table 6 and discuss the impact of 1, 2 and 3 below.

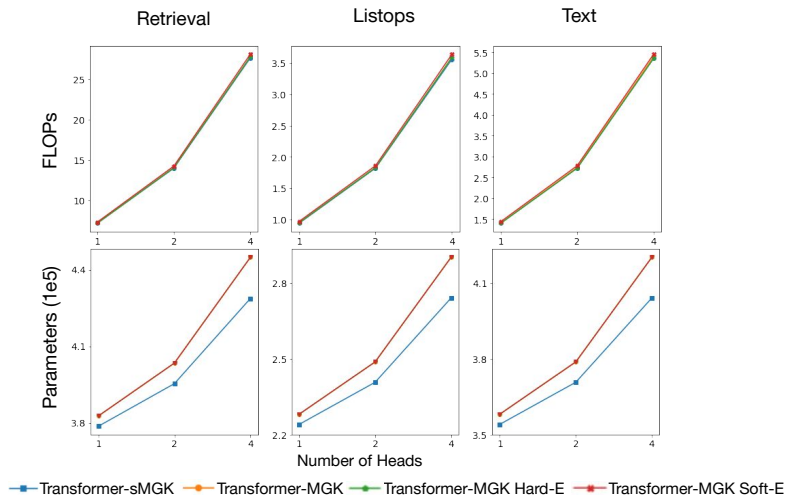


Figure 8. Model complexity (Top) and computational cost (Bottom) of different inference and learning methods for Transformer-MGK trained on the document retrieval task. While computational costs are almost the same, Transformer-sMGK has more advantage in model size, comparing to Transformer-MGK, Transformer-MGK Hard-E, and Soft-E. The naming is as explained in Section 3.4 in the main text.

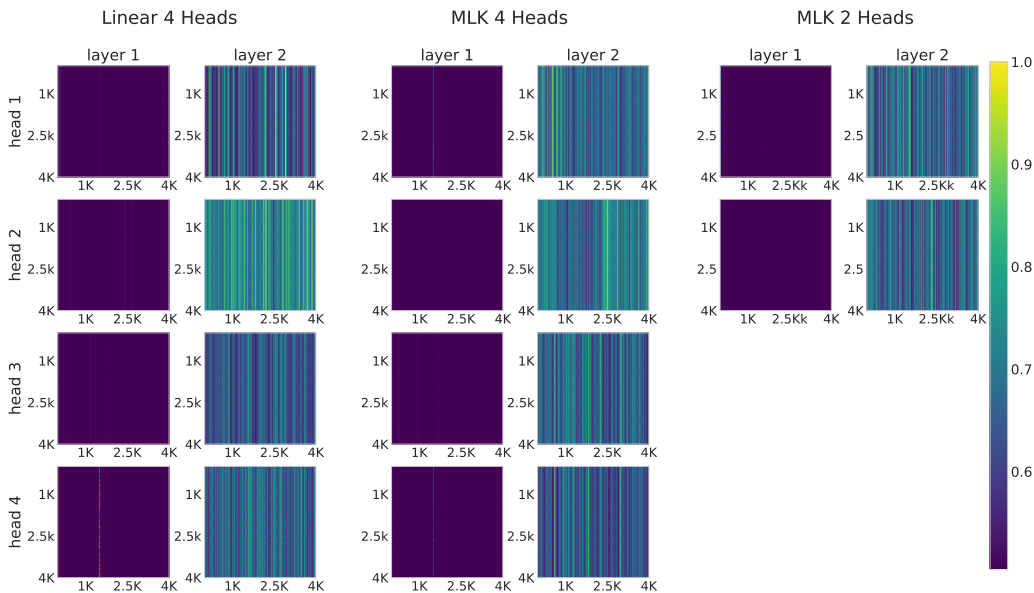


Figure 9. Visualization of attention matrices in the 4-head linear transformer baseline (Left), 4-head Transformer-MLK (Middle), and 2-head Transformer-MLK (Right) trained on the document retrieval task. Here, the sequence length is 4000, and the size of each matrix is 4000×4000 .

Impact of the Mixture of Keys We apply our mixture of keys (MGK) approach to the softmax transformer using the dot product between queries and keys instead of the Gaussian distance as in our paper. We name this model Softmax MGK. We compare the Softmax MGK that has 1 head (Softmax MGK 1 head in Table 6) with the baseline softmax transformers that use 1 and 2 heads (Softmax 1 head and Softmax 2 heads in Table 6). Results in Table 6 show that the Softmax MGK 1 head outperforms both the baseline softmax transformers of 1 and 2 heads. Note that our Softmax MGK 1 head is more efficient than the baseline of 2 heads in terms of the number of parameters and the number of FLOPs. These results confirm the benefit of using MGK.

Impact of using the Gaussian distance Next, we compare the softmax MGK with the Gaussian MGK. Here the Gaussian MGK is the Transformer-MGK proposed and discussed in our paper, which computes the attention scores using the MGK approach and the Gaussian distance between the queries and keys. Results in Table 6 suggest that the Gaussian MGK 1

Table 6. Ablation study on the impact of the mixture of keys, the Gaussian distance, and the key shifting on the LRA retrieval task. We denote the softmax Transformer by Softmax. All Softmax models (i.e., Softmax 2 heads, Softmax 1 head, Softmax MGK 1 head, and Softmax sMGK 1 head) use dot product to compute the attention scores. All Gaussian models (i.e., Gaussian MGK 1 head, Gaussian sMGK 1 head, and Gaussian 1 head) use Gaussian distance to compute the attention scores. We denote MGK with key shifting by sMGK. Here MGK is used to denote our approach of using a mixture of keys at each timestep.

Method	Accuracy (%)
<i>Softmax 2 heads</i>	79.10
Softmax sMGK 1 head	79.81
Softmax MGK 1 head	79.23
<i>Softmax 1 head</i>	77.90
Gaussian sMGK 1 head	81.23
Gaussian MGK 1 head	80.63
<i>Gaussian 1 head</i>	80.38

head improves over the Softmax MGK 1 head (80.63% vs. 79.23%). This result justifies the advantage of using Gaussian distance over dot product to compute the attention scores.

Impact of key shifting Finally, we apply key shifting to both Softmax MGK and Gaussian MGK (Softmax sMGK and Gaussian sMGK in Table 6). From Table 6, we observe that the Softmax sMGK 1 head and Gaussian sMGK 1 head outperform the Softmax MGK 1 head and Gaussian MGK 1 head, respectively. These results, again, corroborate the benefit of using key shifting.

We also include the result for the Gaussian 1 head model in Table 6. This Gaussian 1 head model is similar to the Softmax 1 head model but uses the Gaussian distance to compute the attention scores. Comparing the results of the Gaussian 1 head model, the Softmax 1 head model, and the Gaussian MGK 1 head model reported in Table 6 further confirms the advantage of using MGK and Gaussian distance.

A.5. Efficiency analysis on WikiText-103

Efficiency analysis on WikiText-103 In this section, we provide greater details on the metrics and results of our model efficiency analysis. Extending our method to Linear Transformer (Transformer-MLK) significantly benefits the efficiency of the linear baseline, making it especially applicable for long-range and large-scale tasks.

Analysis metrics Our analysis reports the number of FLOPS as model computational efficiency for both training and inference. The number of total and non-embedding parameters is used to measure model complexity. Since the aim of our method is to reduce the number of parameters in the main model, which does not include the number of input-embedding parameters. Hence, it is also essential to provide non-embedding parameter reduction as a criterion for model size efficiency. Here, we investigate the computation and memory benefits of our model in large-scale tasks through increasing $D \in \{256, 512, 1024, 2048, 4096\}$ and $N \in \{128, 256, 512, 1024, 2048, 4096, 8192\}$. For the FLOP calculation, we use `fvcore`. All measurements are calculated when running the model through data of batch size 1.

Transformer-MLK reduces model complexity and computational cost

Fig. 10A and Fig. 10B shows the reduction in FLOPS ratio between 4-head Transformer-MLK vs. the 8-head Linear baseline as the function of model dimension D and the sequence length N . Fig. 10C presents the model size benefits (number of total/non-embedding parameter ratios) between our model and the baseline as D increases. These figures indicate that as we scale up the tasks and the model, our Transformer-MLK is significantly more advantageous than the Linear baseline.

A.6. Time Overhead and Memory Footprint Analysis

Table 7 compares the GPU memory footprint and computational time overhead (seconds/iteration) of our 4-head Transformer-MGKs/MLKs with those of the 8-head softmax/linear transformer baselines at test time. All models are trained on the LRA retrieval task, and we use a batch size of 32 for each iteration. Our MGK/MLK models save memory and reduce wall-clock time significantly compared to the baselines. Using key shifting in our models, i.e. Transformer-sMGKs/sMLKs, helps improve the efficiency further.

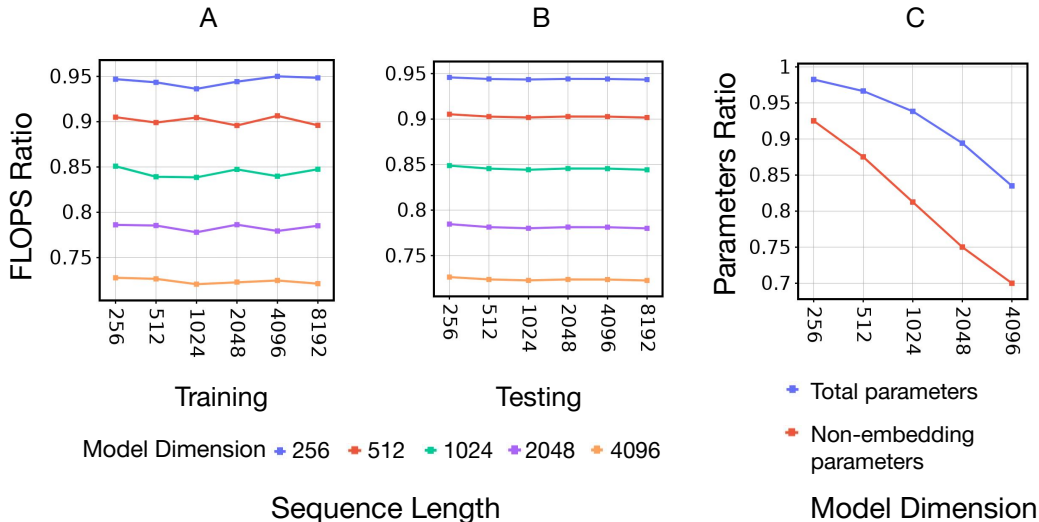


Figure 10. Training (A), Inference (B) FLOP ratios and number-of-parameter ratio (C) between 4-head Transformer-MLK with the 8-head Linear baseline across different D model dimensions and N sequence lengths, for the language modeling task on WikiText-103. 4-head Transformer-MLK significantly reduces computation (both training and inference FLOPS) and model complexity (both the number of total and non-embedding parameters), than the baseline as we scale up model dimension and the sequence length, showing the advantage of Transformer-MLK for long-range and large-scale tasks.

Table 7. Comparing the GPU memory footprint and computational time overhead (seconds/iteration) between our 4-head Transformer-MGKs/MLKs and the 8-head softmax/linear transformer baselines trained on the LRA retrieval task at test time. Our Transformer-MGKs/MLKs save much more memory and have significantly smaller wall-clock time compared to the baselines. Here we use a batch size of 32 for each iteration.

Method	Memory (Gb)	Time overhead (seconds/iteration)
<i>Softmax 8 heads</i>	58.00	0.357
Transformer-MGK 4 heads	53.58	0.278
Transformer-sMGK 4 heads	45.50	0.227
<i>Linear 8 heads</i>	3.38	0.055
Transformer-MLK 4 heads	2.90	0.043
Transformer-sMLK 4 heads	2.83	0.042

Table 8. The learned mixing coefficient π_{jr} of all heads and layers in the 1-head Transformer-MGKs trained on the LRA retrieval task. Here we use the same π_{j1}, π_{j2} for all time step $j = 1, \dots, N$.

Method	Layer 1		Layer 2	
	π_{j1}	π_{j2}	π_{j1}	π_{j2}
Transformer-sMGK 1 heads	0.488	0.512	0.500	0.500
Transformer-MGK 1 heads	0.502	0.498	0.497	0.503

A.7. Learning Curves Showing Convergence

In this section, we replot Figure 1 and 5 to show the convergence of our trainings. In Figure 1, the results do not seem to converge since we plot the loss and the accuracy in log-scale. In Figure 4, the results do not seem to converge since we zoom into the specific range on the y and x-axes. Figure 11 and 12 are the replotted versions of Figure 1 and 5, respectively. In Figure 11, the training loss/accuracy curves stop early because we use early stopping to avoid overfitting. The test loss/accuracy curves in this figure already converge.

A.8. Weight Matrices of the Keys, Keys and Mixing Coefficient

In this section, we analyze the learned π_{jr} , \mathbf{k}_{jr} , and \mathbf{W}_{K_r} , $j = 1, \dots, N$ and $r = 1, \dots, M$ in the Transformer-MGK trained on the LRA retrieval task. In all of our experiments, we set $M = 2$. In Figure 13 and 14, we visualize the weight

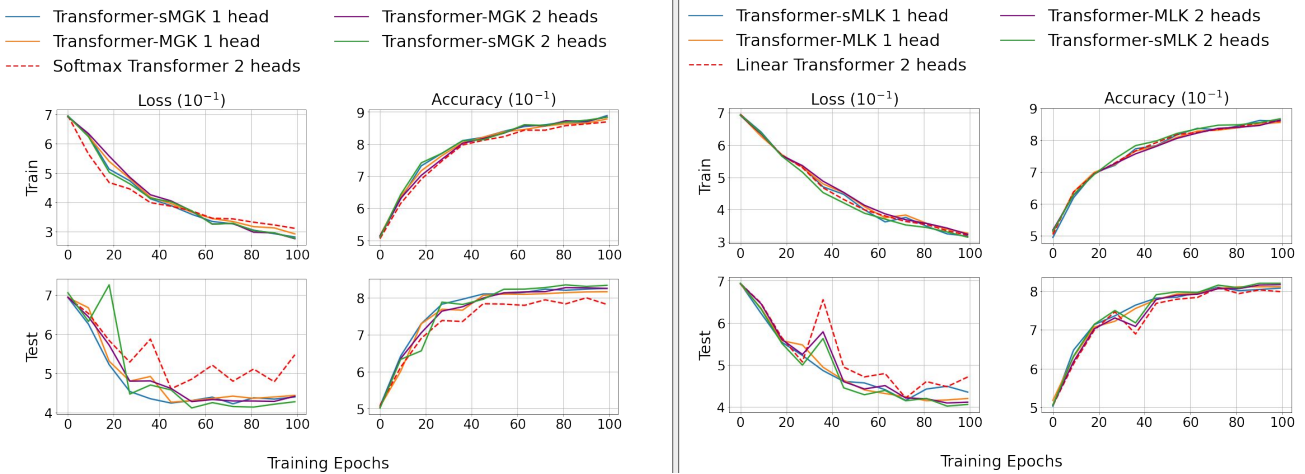


Figure 11. Training and test loss/accuracy of Transformer-MGK vs. softmax transformer (Left) and Transformer-MLK vs. linear Transformer (Right) on the retrieval task, which has the longest average sequence-length and attention span among the LRA tasks (Tay et al., 2021). In training, we apply early stopping to avoid overfitting. That explains why the training loss and accuracy curves stop early. The test loss/accuracy curves already converge. The impressive performance of Transformer-MGK/MLK on this challenging task validates the capability of our models to capture long-range dependencies via learning a diversity of attention patterns.

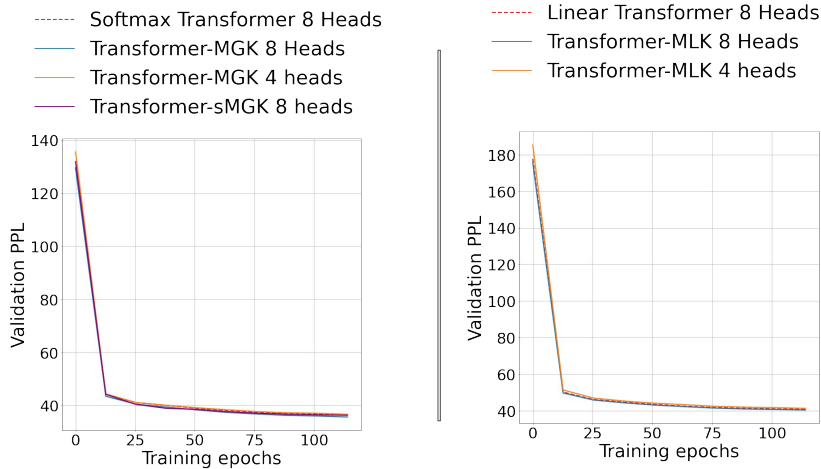


Figure 12. Validation perplexity of the Transformer-MGK vs. the softmax transformer (Left) and the Transformer-MLK vs. the linear transformer (Right) for language modeling on WikiText-103. Training converges on this task after 500000 iterations, equivalent to 115 epochs .

matrices \mathbf{W}_K that computes the keys and the keys \mathbf{K} , respectively, for all heads and layers in the 2-head softmax transformer baseline, the 1-head Transformer-MGK with 2 keys, and the 1-head Transformer-sMGK with 2 keys trained on the LRA retrieval task. Note that for the keys, we only plot the first 100 tokens. Also, Table 8 summarizes the learned mixing coefficient π_{j_r} of all heads and layers in the 1-head Transformer-MGK trained on the same retrieval task. Here we use the same π_{j_1}, π_{j_2} for all time step $j = 1, \dots, N$.

A.9. Additional Computational Complexity (FLOPs) Analysis at Training Time

Figure 15 demonstrates the computational cost (FLOPs) for each training iteration of the Transformer-MGK vs. the baseline softmax transformer (Left) and the Transformer-MLK vs. the baseline linear transformer (Right) on the LRA retrieval task. The efficiency advantage of Transformer-MGK/MLK over the baselines grows with the number of heads.

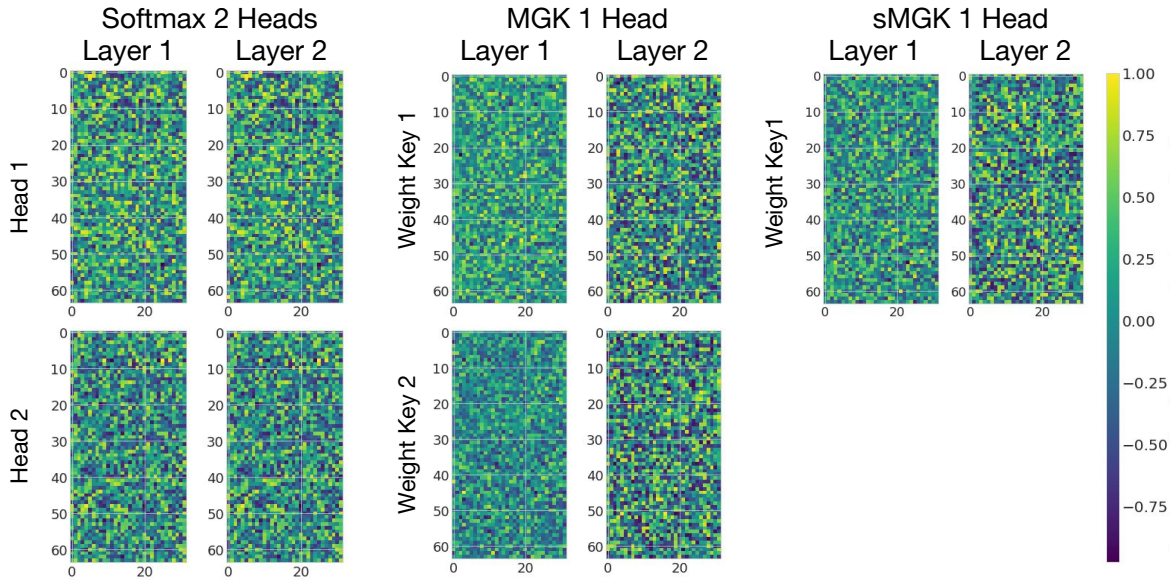


Figure 13. Weight matrices \mathbf{W}_K for computing the keys, for all heads and layers, in the 2-head softmax transformer baseline (Left), the 1-head Transformer-MGK with 2 keys (Middle), and the 1-head Transformer-sMGK with 2 keys (Right) trained on the LRA retrieval task. Here, the dimension of each head $D = 32$ and that of input x_i is $D_x = 64$. Hence, each weight matrix has the shape of $(64, 32)$.

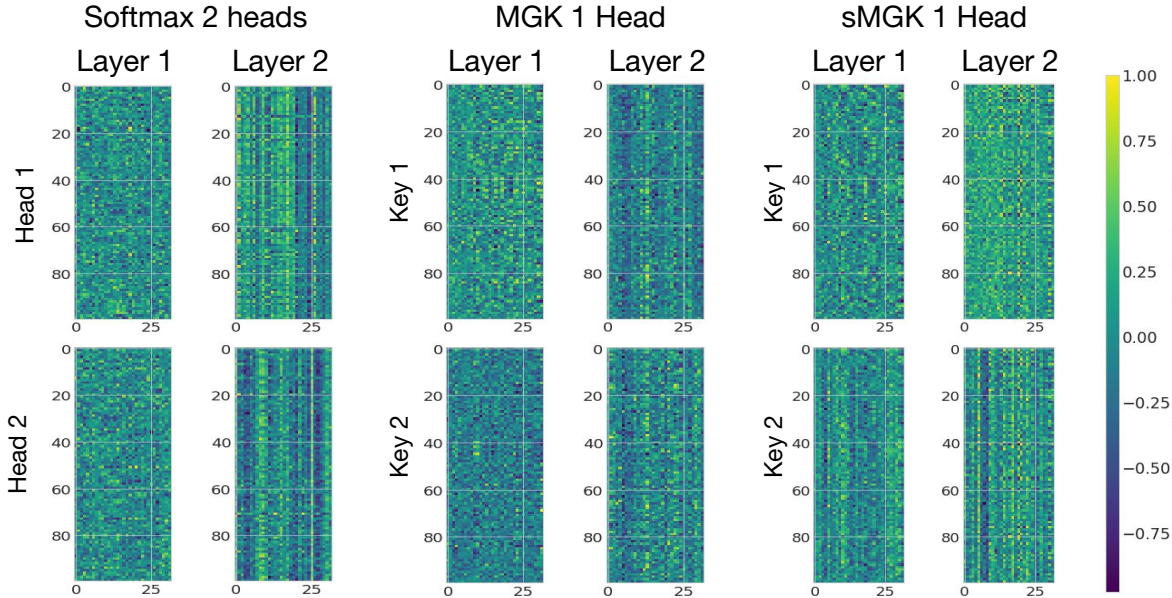


Figure 14. Key embeddings \mathbf{K} for all heads and layers of the 2-head softmax transformer baseline (Left), the 1-head Transformer-MGK with 2 keys (Middle), and the 1-head Transformer-sMGK with 2 keys (Right) trained on the LRA retrieval task. Here the dimension D of each head is 32, and we plot the key embeddings of the first 100 tokens in a randomly chosen sequence. Hence, each key matrix has the shape of $(100, 32)$.

Figure 16 shows the computational cost per training iteration (measured in FLOPs) of different inference and learning methods for Transformer-MGK trained on the document retrieval task. Transformer-sMGK, Transformer-MGK, Transformer-MGK Hard-E, and Soft-E have similar computational costs.

A.10. Scaling to 12-Head Baseline Models for the Retrieval task.

To further study the scalability of our model, in this section, we investigate the performance of our 6-head Transformer-MGKs/MLKs in comparison with the 12-head baseline softmax/linear transformers on the retrieval task. Table 9 indicates

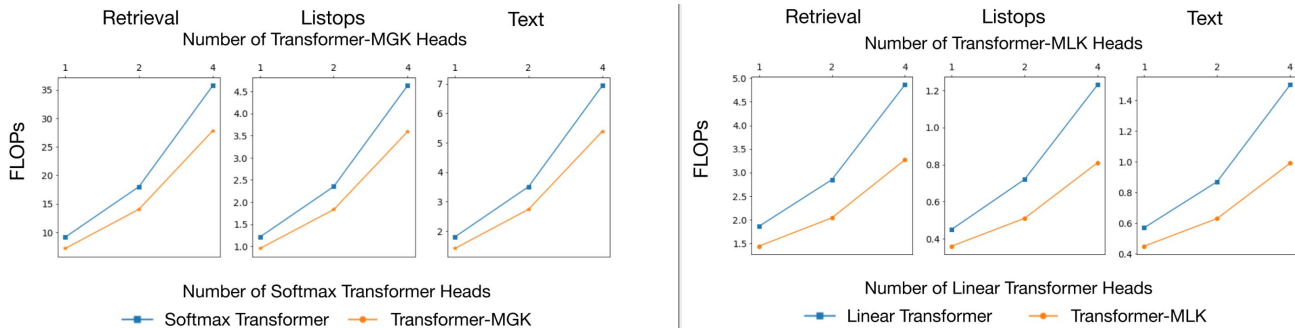


Figure 15. Computational cost (FLOPs) for each training iteration of Transformer-MGK vs. the baseline softmax transformer (Left) and Transformer-MLK vs. the baseline linear transformer (Right) on the LRA retrieval task. The efficiency advantage of Transformer-MGK/MLK over the baselines grows with the number of heads. Here, the batch size of each training iteration is 32.

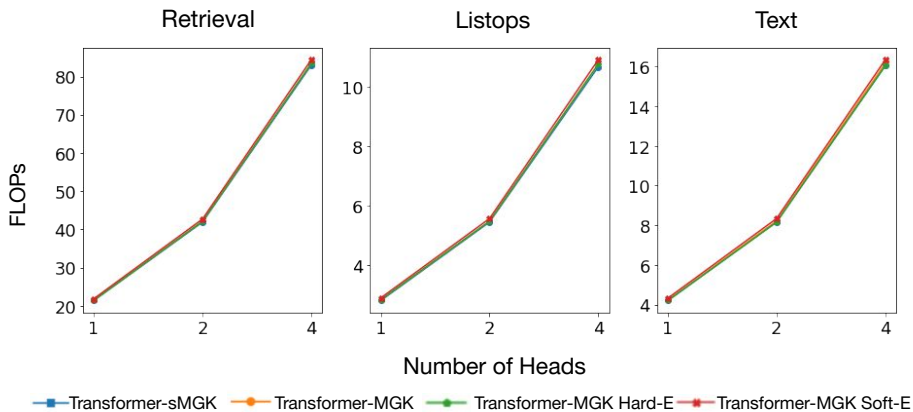


Figure 16. Computational cost (measured in FLOPs) per training iteration of different inference and learning methods for Transformer-MGK trained on the LRA retrieval task. Computational costs are almost the same for Transformer-sMGK, Transformer-MGK, Transformer-MGK Hard-E, and Soft-E. The naming is as explained in Section 3.4 in the main text.

Table 9. Test Accuracy (%) of 6-head Transformer-MGKs/MLKs compared with the baseline 12-head softmax and linear transformers on the retrieval task. Our 6-head Transformer-MGKs/MLKs significantly outperform softmax and linear transformers, respectively, while being more efficient in terms of computational cost, model size, and memory usage.

Method	Accuracy (%)
<i>Softmax 12 heads</i>	82.18
Transformer sMGK 6 head	83.31
Transformer MGK 6 head	83.05
<i>Linear sMGK 12 head</i>	81.97
Transformer sMLK 6 head	82.80
Transformer MLK 6 head	82.11

that our 6-head Transform-MGKs/MLKs significantly outperform 12-head softmax/linear transformers, respectively. Moreover, comparing these results to those in Table 1 and Table 2, although the 12-head softmax/linear transformers improve over the 8-head ones, their accuracies are still worse than or only equivalent to those of our 4-head and even 2-head Transformer-MGK/MLK models.

A.11. Comparison to Multi-query Attention

In this section, we compare our MGK approach to the multi-query attention (Shazeer, 2019). The multi-query attention shares the same set of keys and values at different heads to reduce the memory-bandwidth cost during incremental inference, which allows faster inference since the size of the reloaded "keys" and "values" tensors are significantly reduced. On

Table 10. Perplexity (PPL) on Wikitext-103 of 4-head Multi-query Transformer-MGK and the 8-head Multi-query Transformer baseline. Combining multi-query attention with transformer-MGK results in a better test performance than the baseline while being more efficient since Transformer-MGK enables the model to reduce the number of heads by half.

Method	Valid PPL	Test PPL
<i>Multi-query attention 8 heads (small)</i>	35.08	36.03
Multi-query Transformer MGK 4 head (small)	35.16	35.79
<i>Softmax 8 heads (small)</i>	33.15	34.29
Transformer MGK 4 head (small)	33.28	34.21

another hand, our Transformer-MGK models the key at each head as a Gaussian mixture model, which leads to the use of multiple keys at each head and allows us to decrease the number of attention heads. This helps reduce the computations and parameters needed to calculate additional queries and values. If using key shifting (see option (B) in paragraph Design Options for Keys in Section 2.3), the MGK approach also helps reduce the computations and parameters needed to calculate additional keys. The advantages of Transformer-MGK hold in both training and inference, including incremental inference. We have provided a detailed analysis on the computational complexity and the number of parameters of the Transformer-MGK in comparison with the corresponding softmax transformer in Appendix B. Combining the multi-query attention and our MGK approach is interesting since each method has its own advantage and they are complementary to each other. In particular, we can let the transformer model share the same set of values and mixtures of keys at different heads. This approach can potentially have the advantages of both multi-query attention and MGK. Table 10 shows the Perplexity (PPL) on Wikitext-103 of 4-head Multi-query Transformer-MGK and the 8-head Multi-query Transformer baseline. Combining multi-query attention with transformer-MGK benefits both the model’s performance (Test PPL) compared to the baselines and the its efficiency since Transformer-MGK allows the model to reduce the number of heads by half.

B. An Analysis on the Computational Complexity and the Number of Parameters in Transformer-MGK and the Softmax Transformer

In this section, we compare the computational complexity and the number of parameters in transformer-MGK with M keys at each timestep and H/M heads to the baseline softmax transformer that has 1 key at each timestep and H heads. Here we choose M such that H is a multiple of M and use keys design option (A) that make the key k_{jr} a linear projection of the input x_j , i.e. $k_{jr} = x_j \mathbf{W}_{K_r}^\top$, where $x_j \in \mathbb{R}^{1 \times D_x}$, $\mathbf{W}_{K_r} \in \mathbb{R}^{D \times D_x}$ and $r = 1, 2, \dots, M$ (see Design Options for Keys in Section 2.3 for more details). To simplify notation and without loss of generality, we let $M = 2$ as in our experiments and assume that $D_v = D$, i.e., the values have the same feature dimension as the queries and the keys. To simplify the computation, we also do not take the softmax operator into account since this operator yields similar costs when applied in Transformer-MGK and softmax transformer.

B.1. Computational Complexity

(i) Softmax H -head attention: The number of computations in a softmax H -head attention is $N^2 H(4D - 1) + NHD(6D_x + 2HD - 5)$.

Explanation: To calculate the query matrix \mathbf{Q} , the key matrix \mathbf{K} , and the value matrix \mathbf{V} in Step 1 in Section 1.1 at each head, we need $3NDD_x$ multiplications and $3ND(D_x - 1)$ additions. In total, these need $3ND(2D_x - 1)$ computations. Next, to compute the product \mathbf{QK}^\top in Eqn. (1), we need $N^2 D$ multiplications and $N^2(D - 1)$ additions. Similarly, the product \mathbf{AV} requires $N^2 D$ multiplications and $N(N - 1)D$ additions. In total, computing the output sequence \mathbf{H} in Eqn. (1) at each head requires $3ND(2D_x - 1) + N^2 D + N^2(D - 1) + N^2 D + N(N - 1)D = N^2(4D - 1) + ND(6D_x - 4)$ computations. The total computation for all H heads is then

$$\begin{aligned} & H(N^2(4D - 1) + ND(6D_x - 4)) + NHD(2HD - 1) \\ & = N^2 H(4D - 1) + NHD(6D_x + 2HD - 5), \end{aligned}$$

where the extra $NHD(2HD - 1)$ is from the linear projection by \mathbf{W}^O .

(ii) Mixture of 2 Gaussian keys attention with $H/2$ -head: The number of computations in a mixture of 2 Gaussian keys attention with $H/2$ -head is $N^2 H(3D - 0.5) + NHD(4D_x + HD - 4)$.

Explanation: Similar to the above derivation, in a Mixture of M Gaussian keys attention, to compute the output sequence \mathbf{H} we need $N^2((2M+2)D-1) + ND((M+2)(2D_x-1)-1)$ computations. Note that, to compute the Gaussian distances between the queries q_i and the keys k_j as in the mixture of M Gaussian keys attention and to compute their dot product as in the softmax attention, we need the similar number of computations. Therefore, the total computation for all H/M heads is then

$$\begin{aligned} & (H/M)(N^2((2M+2)D-1) + ND((M+2)(2D_x-1)-1)) + NHD(2(H/M)D-1) \\ &= N^2H \left(\frac{2(M+1)D-1}{M} \right) + NHD \left(\frac{2(M+2)}{M}D_x + \frac{2}{M}HD - \frac{3M+2}{M} \right). \end{aligned}$$

Let $M = 2$, then the total computation of the mixture of 2 Gaussian keys attention is given by $N^2H(3D-0.5) + NHD(4D_x + HD - 4)$.

Soft-max H-head attention versus mixture of 2 Gaussian keys attention with H/2-head: Given the results in (i) and (ii), when compared to the baseline softmax H -head attention, our mixture of 2 Gaussian keys attention with $H/2$ -head saves

$$N^2H(D-0.5) + NHD(2D_x + HD - 1)$$

computations. When N is large, this difference is significant. In conclusion, the mixture of 2 Gaussian keys attention with $H/2$ -heads has cheaper computational complexity than that of soft-max H -head attention.

B.2. The Number of Parameters

(iii) Softmax H -head attention: The number of parameters in a softmax H -head attention is $3HDD_x + (HD)^2$.

Explanation: $3HDD_x$ is from the linear projects to calculate the query matrix \mathbf{Q} , the key matrix \mathbf{K} , and the value matrix \mathbf{V} in Step 1 in Section 1.1. $(HD)^2$ is from the linear project to compute the final output as in Eqn. (??).

(iv) Mixture of 2 Gaussian Keys attention with $H/2$ -head: The number of parameters in a Mixture of 2 Gaussian Keys attention with $H/2$ -head is $2HDD_x + 0.5(HD)^2 + H$.

Explanation: The linear projections to calculate \mathbf{Q} and \mathbf{V} contribute $HDD_x/2$ parameters each. The linear projection to calculate \mathbf{K} contributes HDD_x . The linear project to compute the final output has dimension $(H/2)D \times HD$, so it contributes $0.5(HD)^2$ parameters. H more parameters is from the prior π_{j_r} . These priors contribute 2 parameters at each head since we make $\{\pi_{j_1}, \pi_{j_2}\}$ for all $j = 1, \dots, N$ the same.

Soft-max H-head attention versus mixture of 2 Gaussian keys attention with H/2-head: Given the results in (iii) and (iv), when compared to the baseline softmax H -head attention, our mixture of 2 Gaussian keys attention with $H/2$ -head saves $HDD_x + 0.5(HD)^2 - H$ parameters. When H and D are large, this saving is significant.

C. Proofs of main results

In this appendix, we provide proofs for the main results in the paper.

C.1. Proof of Theorem 1

To ease the presentation of the proof, for any probability distribution G , we denote

$$p_G(x) := \int f(x-\theta)dG(\theta) = \int \phi(x|\theta, \sigma^2\mathbf{I})dG(\theta),$$

for all $x \in \mathbb{R}^d$ where $f(x) = \frac{1}{(\sqrt{2\pi}\sigma)^d} \exp\left(-\frac{\|x\|^2}{2\sigma^2}\right)$ for given $\sigma > 0$. It means that p_G is the convolution of f and the probability distribution G . Since the space of Gaussian mixtures is dense in the space of continuous probability measures (Bacharoglou, 2010), it indicates that there exists probability distribution G_1 such that

$$\sup_{x \in \mathbb{R}^d} |p(x) - p_{G_1}(x)| \leq \frac{\epsilon}{2}. \quad (9)$$

Our next step is to prove that there exists a probability measure G_2 with at most K supports where $K \leq (C \log(1/\epsilon))^d$ for some universal constant C such that

$$\sup_{x \in \mathbb{R}^d} |p_{G_1}(x) - p_{G_2}(x)| \leq \frac{\epsilon}{2}. \quad (10)$$

Indeed, from Lemma A.1 in (Ghosal & van der Vaart, 2001), for any $k \geq 1$ there exists a probability distribution G_2 with at most $(2k - 2)^d$ supports such that

$$\int \theta^\alpha d(G_1 - G_2)(\theta) = 0, \quad (11)$$

for any $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{N}^d$ such that $0 \leq |\alpha| = \sum_{j=1}^d \alpha_j \leq 2k - 2$. Here, $\theta^\alpha = \prod_{j=1}^d \theta_j^{\alpha_j}$.

Now, for any $M \geq 2a\sqrt{d}$, we have $\|x - \theta\| \geq \|x\| - \|\theta\| > M - a\sqrt{d} > M/2$ as long as $\|x\| > M$ and $\theta \in [-a, a]^d$. It indicates that

$$\begin{aligned} \sup_{\|x\| > M} |p_{G_1}(x) - p_{G_2}(x)| &= \sup_{\|x\| > M} \left| \int f(x - \theta) d(G_1 - G_2)(\theta) \right| \\ &\leq \sup_{\|x\| > M} \int \frac{1}{(\sqrt{2\pi}\sigma)^d} \exp\left(-\frac{\|x - \theta\|^2}{2\sigma^2}\right) d(G_1 + G_2)(\theta) \\ &\leq \frac{2}{(\sqrt{2\pi}\sigma)^d} \exp\left(-\frac{M^2}{8\sigma^2}\right). \end{aligned} \quad (12)$$

On the other hand, for any $k \geq 1$ we also have that

$$\begin{aligned} \sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| &= \sup_{\|x\| \leq M} \left| \int f(x - \theta) d(G_1 - G_2)(\theta) \right| \\ &\leq \sup_{\|x\| \leq M} \left| \int \left(f(x - \theta) - \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} \right) d(G_1 - G_2)(\theta) \right| \\ &\quad + \sup_{\|x\| \leq M} \left| \int \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} d(G_1 - G_2)(\theta) \right| \\ &= \sup_{\|x\| \leq M} \left| \int \left(f(x - \theta) - \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} \right) d(G_1 - G_2)(\theta) \right|, \end{aligned} \quad (13)$$

where the final equality is stems from

$$\int \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} d(G_1 - G_2)(\theta) = 0,$$

which is due to Eqn. (11).

To further bound the right-hand-side (RHS) of Eqn. (13), we use the following inequality:

$$\left| \exp(y) - \sum_{j=0}^{k-1} \frac{y^j}{j!} \right| \leq |y|^k / k!$$

for any $y \in \mathbb{R}$. Since $k! \geq (k/e)^k$ for any $k \geq 1$, the above bound can be rewritten as

$$\left| \exp(y) - \sum_{j=0}^{k-1} \frac{y^j}{j!} \right| \leq \frac{|ye|^k}{k^k}. \quad (14)$$

Further simplification of Eqn. (13) leads to

$$\begin{aligned}
 \sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| &\leq \sup_{\|x\| \leq M} \int \left| f(x - \theta) - \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} \right| d(G_1 + G_2)(\theta) \\
 &\leq 2 \sup_{\|x\| \leq M, \theta \in [-a, a]^d} \left| f(x - \theta) - \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{(\sqrt{2\pi})^d \sigma^{d+2j} j!} \right| \\
 &= \sup_{\|x\| \leq M, \theta \in [-a, a]^d} \frac{2}{(\sqrt{2\pi}\sigma)^d} \left| \exp\left(-\frac{\|x - \theta\|^2}{2\sigma^2}\right) - \sum_{j=0}^{k-1} \frac{(-1)^j \|x - \theta\|^{2j}}{\sigma^{2j} j!} \right| \\
 &\leq \sup_{\|x\| \leq M, \theta \in [-a, a]^d} \frac{e^k \|x - \theta\|^{2k}}{\sigma^{2k} (2k)^k},
 \end{aligned}$$

where the final inequality is based on an application of inequality (14) with $y = -\|x - \theta\|^2/(2\sigma^2)$. For $\|x\| \leq M$ and $\theta \in [-a, a]^d$, we have $\|x - \theta\| \leq \|x\| + \|\theta\| \leq M + a\sqrt{d}$. Therefore, we further have

$$\sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| \leq \sup_{\|x\| \leq M, \theta \in [-a, a]^d} \frac{e^k \|x - \theta\|^{2k}}{\sigma^{2k} (2k)^k} \leq \frac{e^k (M + a\sqrt{d})^{2k}}{\sigma^{2k} (2k)^k}.$$

When $M \geq 2a\sqrt{d}$, we have $M + a\sqrt{d} \leq \frac{3M}{2}$ and the above bound leads to

$$\sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| \leq \frac{(9e)^k M^{2k}}{(8\sigma^2 k)^k}. \quad (15)$$

By choosing $M^2 = 8\sigma^2 \log(1/\epsilon')$ for some $\epsilon' > 0$, the bounds in Eqns. (12) and (15) become

$$\begin{aligned}
 \sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| &\leq \frac{2}{(\sqrt{2\pi}\sigma)^d} \epsilon', \\
 \sup_{\|x\| > M} |p_{G_1}(x) - p_{G_2}(x)| &\leq \frac{(9e)^k (\log(1/\epsilon'))^k}{k^k}.
 \end{aligned} \quad (16)$$

As long as we choose $k = 9e^2 \log(1/\epsilon')$ and $\epsilon' \leq 1$, we have

$$\sup_{\|x\| > M} |p_{G_1}(x) - p_{G_2}(x)| \leq e^{-k} = e^{-9e^2 \log(1/\epsilon')} = (\epsilon')^{9e^2} \leq \epsilon'. \quad (17)$$

By choosing $\epsilon' = \frac{\epsilon}{2 \max\{\frac{2}{(\sqrt{2\pi}\sigma)^d}, 1\}}$, the results from Eqns. (16) and (17) indicate that

$$\sup_{\|x\| \leq M} |p_{G_1}(x) - p_{G_2}(x)| \leq \frac{\epsilon}{2}, \quad \text{and} \quad \sup_{\|x\| > M} |p_{G_1}(x) - p_{G_2}(x)| \leq \frac{\epsilon}{2}.$$

Therefore, if we choose $M = 8\sigma^2 \log\left(\frac{2 \max\{\frac{2}{(\sqrt{2\pi}\sigma)^d}, 1\}}{\epsilon}\right)$ and $k = 9e^2 \log\left(\frac{2 \max\{\frac{2}{(\sqrt{2\pi}\sigma)^d}, 1\}}{\epsilon}\right)$, we have

$$\sup_{x \in \mathbb{R}^d} |p_{G_1}(x) - p_{G_2}(x)| \leq \frac{\epsilon}{2}.$$

It indicates that we obtain the conclusion of claim (10) by choosing $K = (2k - 2)^d \leq \left(18e^2 \log\left(\frac{2 \max\{\frac{2}{(\sqrt{2\pi}\sigma)^d}, 1\}}{\epsilon}\right)\right)^d$.

Combining the results from Eqns. (9) and (10), we have

$$\sup_{x \in \mathbb{R}^d} |p(x) - p_{G_2}(x)| \leq \sup_{x \in \mathbb{R}^d} |p(x) - p_{G_1}(x)| + \sup_{x \in \mathbb{R}^d} |p_{G_1}(x) - p_{G_2}(x)| \leq \epsilon.$$

As a consequence, we obtain the conclusion of the theorem.