

ON LEARNING THE TRANSFORMER KERNEL

Anonymous authors

Paper under double-blind review

ABSTRACT

In this work we introduce **KERNELIZED TRANSFORMER**, a generic, scalable, data driven framework for learning the kernel function in Transformers. Our framework approximates the Transformer kernel as a dot product between spectral feature maps and learns the kernel by learning the spectral distribution. This not only helps in learning a generic kernel end-to-end, but also reduces the time and space complexity of Transformers from quadratic to linear. We show that **KERNELIZED TRANSFORMERS** achieve performance comparable to existing efficient Transformer architectures, both in terms of accuracy as well as computational efficiency. Our study also demonstrates that the choice of the kernel has a substantial impact on performance, and kernel learning variants are competitive alternatives to fixed kernel Transformers, both in long as well as short sequence tasks.

1 INTRODUCTION

Transformer models (Vaswani et al., 2017) have demonstrated impressive results on a variety of tasks dealing with language understanding (Devlin et al., 2019; Radford et al., 2018; Raffel et al., 2020; Brown et al., 2020), image processing (Parmar et al., 2018; Carion et al., 2020; Lu et al., 2019), as well as biomedical informatics (Rives et al., 2020; Ingraham et al., 2019; Madani et al., 2020). Albeit powerful, due to the global receptive field of self-attention, the time and memory complexity of Softmax Transformer models scale quadratically with respect to the sequence length. As a result, the application of Transformers to domains with long contexts is rather limited. This limitation has spawned several efficient Transformer designs (Liu et al., 2018; Parmar et al., 2018; Child et al., 2019; Zaheer et al., 2020; Beltagy et al., 2020; Roy et al., 2020; Tay et al., 2020a; Kitaev et al., 2020). Kernelization offers one such design. The use of kernel feature maps allows to reformulate the computation of attention in a way that avoids the explicit computation of the full attention matrix which is the key bottleneck for Softmax Transformer. This also opens up new directions for more generic attention mechanisms.

Tsai et al. (2019) first proposed a kernel-based formulation of the attention mechanism. However, the time and memory complexity of their approach remains quadratic with respect to the sequence length. To address this limitation, Katharopoulos et al. (2020) expressed self-attention as the inner product of kernel feature maps and made use of the associative property to reduce the complexity from quadratic to linear. For their experiments, they used the arbitrarily chosen LinearElu feature map $f(x) = \max(x + 1, e^x)$. Performer (Choromanski et al., 2021) replaces this with feature maps that can directly approximate the softmax kernel, thereby allowing the use of pre-trained Softmax Transformer weights in a linear time model. Concurrently with them, Peng et al. (2021) proposed a linear space and time method that added causal and recency based features to random Fourier methods. While the aforementioned approaches provide a significant reduction in computational and memory requirements, this often comes at the cost of performance, as can be seen from Fig. 1. In this work, we posit that this is partly due to the fact that the similarity functions/kernels, including scaled-dot-product, were hand picked and not learnt from data. Thus, we explore whether kernel learning can help to bridge this gap in performance while retaining the scalability of efficient Transformers.

Although, to the best of our knowledge, kernel learning has never been explored within the framework of Transformers, kernel learning methods have been an ubiquitous tool in machine learning. The most notable among them is Random Kitchen Sinks (RKS; Rahimi & Recht, 2007), a data-independent framework for approximating shift-invariant kernels using an explicit feature map. In RKS, the kernel is approximated by $\kappa(x, y) \approx \langle \phi(x), \phi(y) \rangle$, where the explicit feature map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^s$ is obtained by sampling from a spectral distribution defined by the inverse Fourier

transform of the kernel function κ . Wilson & Adams (2013) modeled the spectral density as a mixture of Gaussians, A la Carte (Yang et al., 2015) proposed an optimization based framework for learning the spectral distribution, BaNK (Oliva et al., 2016) modeled the spectral distribution using an infinite mixture of Gaussians, while Fang et al. (2020) implicitly modeled it using deep generative models. We build on these advances and incorporate them into the Transformer framework.

Contributions: In this work, we propose **KERNELIZED TRANSFORMER**, a scalable data driven framework for learning the kernel of Transformers and investigate whether a fully learnable kernel can help to improve the performance of linear, fixed kernel Transformers. Thus, we introduce Transformers with learnable similarity functions, which happen to retain the linear complexity in terms of the sequence length. We motivate our learning method using RKS and learn the kernel by learning the corresponding spectral distribution. In §2.1 we first propose to learn a generic Transformer kernel by explicitly approximating the spectral distribution using a Mixture of Gaussians (GMM) and propose modifications to scale it further. In an attempt to further explore the trade off between computational complexity and accuracy we also propose to model the spectral frequency distribution of Transformer kernels implicitly by using deep generative models (Goodfellow et al., 2014). Finally, we also propose a novel method to learn the spectral distribution of positive random feature (PRF) maps, which provides a better approximation of the softmax kernel (Choromanski et al., 2021).

We analyse the expressivity of our proposed models (§2.2) and show that the proposed GMM with positional encodings is Turing-complete (Pérez et al., 2019). We experimentally evaluate our models on *LRA* (tasks with long context) and *GLUE* (tasks with short context) and analyze the performance of our models (§3). In our experiments, we find that learnt kernels improve performance in some of our long-context tasks, while staying competitive to the Softmax Transformer of the same size in short-context tasks. We also benchmark the computational efficiency of **KERNELIZED TRANSFORMERS** and find that each of our proposed **KERNELIZED TRANSFORMERS** scales linearly with respect to the sequence length.

2 KERNEL LEARNING IN TRANSFORMERS

We begin with the generalized formulation of self-attention proposed by Tsai et al. (2019). Given a non-negative kernel function $\kappa(\cdot, \cdot) : \mathbb{R}^{d_k} \times \mathbb{R}^{d_k} \rightarrow \mathbb{R}_+$, the output of the generalized self-attention mechanism at index i operating on an input sequence $X = (x_1, \dots, x_L) \in \mathbb{R}^{L \times d}$ is defined as

$$\text{ATTN}(X)_i = \sum_{j=1}^L \frac{\kappa(q_i, k_j)}{\sum_{j'=1}^L \kappa(q_i, k_{j'})} v_j. \quad (1)$$

where $k_i = x_i W^K$, $q_i = x_i W^Q$, $v_i = x_i W^V$ are linear transformations of the input sequence into keys, queries and values of dimension $d_q = d_k$ and d_v respectively and $W^K \in \mathbb{R}^{d \times d_k}$, $W^Q \in \mathbb{R}^{d \times d_q}$, $W^V \in \mathbb{R}^{d \times d_v}$. While the formulation in Eq. (1) is more generic and defines a larger space of attention functions, it suffers from a quadratic time and memory complexity. To reduce the quadratic time and memory, we briefly review the method of random Fourier features for the approximation of kernels (Rahimi & Recht, 2007). The details of the method will help motivate and explain our models.

Random Fourier Features for Kernels: At the heart of this method lies the theorem of Bochner (Rudin, 1990) which states that a continuous shift invariant kernel $\kappa(q, k) = \tilde{\kappa}(q - k)$ over arbitrary variables q and k is a positive definite function if and only if $\tilde{\kappa}(\delta)$ is the Fourier transform of a non-negative measure $\rho(\omega)$. Moreover, if $\tilde{\kappa}(0) = 1$, then Bochner’s theorem guarantees




Figure 1: Peak memory (y-axis), average performance (x-axis) and speed (denoted by area of circle) for various efficient Transformer models (i.e. bigger circles in the bottom right corner are better) across four task introduced in *LRA* (Tay et al., 2021b). All values except for “This Paper” are taken from Tay et al. (2021b).

that $\rho(\omega)$ is a normalized density function, i.e.

$$\tilde{\kappa}(q - k) = \int_{\mathbb{R}^d} \rho(\omega) \exp(i\omega^T(q - k)) d\omega = \mathbb{E}_{\omega \sim \rho} [\exp(i\omega^T q) \exp(i\omega^T k)^*]. \quad (2)$$

Rahimi & Recht (2007) proposed to sample from the spectral density $\rho(\omega)$ for a Monte Carlo approximation to the integral in Eq. (2). Specifically, for real valued kernels, they define $\kappa(q, k) \approx \phi(q)^T \phi(k)$, where $\omega_i \sim \rho(\omega)$ and

$$\phi(x) := RKS(x, \Omega = (\omega_1, \dots, \omega_M)) := \frac{1}{\sqrt{M}} [\cos(\omega_1^T x), \dots, \cos(\omega_M^T x), \sin(\omega_1^T x), \dots, \sin(\omega_M^T x)] \quad (3)$$

To learn a kernel, we can either learn a parametric function $\kappa(\cdot, \cdot)$ or learn the corresponding parameterized feature map $\phi(\cdot)$ directly, which corresponds to learning the spectral density $\rho(\omega)$ (Wilson & Adams, 2013; Yang et al., 2015; Oliva et al., 2016). In this paper, we focus on the latter because this helps us in keeping the computational complexity linear in the sequence length L . This can be achieved by rewriting Eq. (1) as $\text{ATTN}(X)_i = \frac{\phi(q_i)^T \sum_{j=1}^L \phi(k_j) v_j^T}{\phi(q_i)^T \sum_{j'=1}^L \phi(k_{j'})}$. To the best of our knowledge this is the first attempt at learning the kernel of the generalized self-attention mechanism (Eq. 1).

2.1 LEARNING KERNELS IN SPECTRAL DOMAIN

GMM-RKS: Our objective is to enable learning of any translation invariant kernel. This is realizable if we can learn the spectral distribution. Gaussian Mixture Models (GMMs) are known universal approximators of densities and hence may approximate any spectral distribution. GMMs have been shown to be useful for kernel learning for regression and classification tasks (Wilson & Adams, 2013; Oliva et al., 2016). Thus, to learn the kernel of the generalized self-attention mechanism (Eq. 1), we model the spectral distribution of the kernel as a parameterized GMM, i.e.

$$\rho(\omega) = \sum_{c=1}^C \pi_c \mathcal{N}(\mu_c, \Sigma_c) \iff \kappa(q, k) = \sum_{c=1}^C \pi_c \exp(i\mu_c^T(q - k) - \frac{1}{2}(q - k)^T \Sigma_c (q - k)). \quad (4)$$

Here $\{\mu_c \in \mathbb{R}^d, \Sigma_c \in \mathbb{R}^{d^2}\}_{c=1}^C$ are the learnable parameters of the feature map and C is the number of components in the Gaussian mixture. It can be shown using Plancherel’s Theorem that $\rho(\omega)$ can approximate any shift invariant kernel (Silverman, 1986). Since we are working with only real valued kernels, the corresponding kernel reduces to $\kappa(q, k) = \sum_{c=1}^C \pi_c \exp(-\frac{1}{2}(q - k)^T \Sigma_c (q - k)) \cos(\mu_c^T(q - k))$.

To speedup learning, we assume that $\pi_c = \frac{1}{C}$ and parameterize the feature map as:

$$\phi_{\text{GMM-RKS}}(x) := RKS(x, \Omega = (\omega_{c,1}, \dots, \omega_{C,M})), \quad \omega_{c,m} = \Sigma_c n_m + \mu_c, \quad n_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (5)$$

This allows us to sample $n_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and learn the parameters of the feature map, $(\{\mu_c \in \mathbb{R}^{d_q}, \Sigma_c \in \mathbb{R}^{d_q^2}\}_{c=1}^C)$ end-to-end along with the other parameters of the Transformer.

FASTFOOD-RKS: GMM-RKS removes the quadratic dependency on context length, but we still need to calculate $\Omega^T \mathbf{Q}$ and $\Omega^T \mathbf{K}$ (where $\Omega = [\omega_1, \omega_2, \dots, \omega_M]$) which takes $\mathcal{O}(Md_q L)$ time and $\mathcal{O}(Md_q + d_q L + ML)$ space, which can be too much if M is large. For further gains in scalability, we approximate the spectral frequency matrix Ω , using the product of Hadamard matrices (FastFood; Le et al., 2013), such that the computation can be done in time log-linear in M , i.e.:

$$\phi_{\text{FASTFOOD-RKS}}(x) := RKS(x, V), \quad \text{where } V = \frac{1}{\sigma \sqrt{d_q}} SHG\Pi HB. \quad (6)$$

Here, $\Pi \in \{0, 1\}^{d_q \times d_q}$ is a permutation matrix, H is the Walsh-Hadamard matrix, B is a diagonal random matrix with $\{\pm 1\}$ entries, G is a diagonal matrix with Gaussian entries drawn from $\mathcal{N}(0, 1)$ and finally S is a random diagonal scaling matrix that makes the row lengths non-uniform. The entire multiplication can be carried out in logarithmic time, and the space requirement is reduced by storing diagonal matrices instead of full matrices. For $M > d_q$ we use multiple blocks, and the only restriction is that we need $M|d_q$. In order to make this learnable, Yang et al. (2015) proposed making S and optionally G and B learnable. For the main paper, we keep all three learnable (the case where only S is learnable is discussed in Appendix C).

GENERATIVE-RKS: If we increase the number of components (C) in GMM-RKS, the computation and space complexity increases dramatically. Instead, to learn a more generic kernel, without blowing up the computational complexity, we use deep generative models (DGMs). DGMs have achieved impressive results in density estimation (Goodfellow et al., 2014; Kingma & Welling, 2014; Richardson & Weiss, 2018; Ruthotto & Haber, 2021) and end-to-end kernel learning in the spectral domain for classification (Fang et al., 2020).

We adapt DGMs and incorporate them within the Transformer architecture in order to jointly learn the kernel (implicitly) in an end-to-end fashion. In GENERATIVE-RKS, the generator network acts as a sampler. It implicitly learns a latent density function $\rho_g(\omega)$ and generates samples from it. Given samples $\eta_1, \dots, \eta_M \in \mathbb{R}^{d_q}$ drawn from an arbitrary noise distribution $\rho_o(\cdot)$ (we use a standard Gaussian in our experiments), the generator $g(\cdot)$ takes them as input and outputs a set of samples $\omega_1, \dots, \omega_M \in \mathbb{R}^{d_q}$ which are distributed according to $\rho_g(\omega)$:

$$\phi_{\text{GENERATIVE-RKS}}(x) := \text{RKS}(x, \Omega = (\omega_1, \dots, \omega_M)), \quad \omega_m = g(n_m), \quad n_m \sim \rho_o(\cdot) \quad (7)$$

We experimented with and learnt the parameters of the generator network which consisted of 4 fully connected layers with batch normalisation and LeakyReLU activation, followed by a single fully connected layer with tanh activation.




Figure 2: Generalized self-attention with deep generative RKS. Q , V and K are linear transformations of input, X . The generator generates spectral frequency (Ω) from an implicit spectral distribution. Using RKS (Eq. 2) we create the feature map ϕ_{gen} (Eq. 7). The numerator of the output is calculated as $\phi_{gen}(Q)(\phi_{gen}(K)V)$ while the denominator is $\phi_{gen}(q_i)^T \sum_{j'=1}^L \phi_{gen}(k_{j'})$ making attention linear in sequence length L .

Positive Random Features (PRF): Until now, we focused on feature maps defined using RKS. While our formulation is very general, recently it was shown that positive random features provide a better approximation to both Gaussian and Softmax kernels (see Lemma 2 in Choromanski et al. 2021). In particular they showed that $\kappa(q, k) = \exp(q^T k) = \mathbb{E}_{\omega \sim \mathcal{N}(0, I)}[\exp(\omega^T q - \frac{\|q\|^2}{2}) \exp(\omega^T k - \frac{\|k\|^2}{2})]$ and demonstrated that Monte Carlo approximation to this expectation leads to a low variance estimate of the softmax kernel. Moreover, the presence of only positive values within the randomized feature map ensures that kernel estimates remain strictly non-negative. To incorporate this prior knowledge, we propose a novel kernel learning framework in which we learn the spectral density while using the feature map corresponding to the above expectation. For instance, when we model the spectral distribution using GMM we have that:

$$\kappa(q, k) := \mathbb{E}_{\omega \sim \sum_{c=1}^C \pi_c \mathcal{N}(\mu_c, \Sigma_c)}[\exp(\omega^T q - \frac{\|q\|^2}{2}) \exp(\omega^T k - \frac{\|k\|^2}{2})] \quad (8)$$

$$\text{PRF}(x, (\omega_1, \dots, \omega_M)) := \frac{\exp(-\frac{\|x\|^2}{2})}{\sqrt{M}} [\exp(\omega_1^T x), \dots, \exp(\omega_M^T x)] \quad (9)$$

$$\phi_{\text{GMM-PRF}}(x) := \text{PRF}(x, (\omega_{c,1}, \dots, \omega_{c,M})), \quad \omega_{c,m} = \Sigma_c n_m + \mu_c, \quad n_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (10)$$

Similarly we redefine the other two methods as:

$$\phi_{\text{GENERATIVE-PRF}}(x) := \text{PRF}(x, (\omega_1, \dots, \omega_M)), \quad \omega_m = g(n_m), \quad n_m \sim \rho_o(\cdot) \quad (11)$$

$$\phi_{\text{FASTFOOD-PRF}}(x) := \text{PRF}(x, V), \quad \text{where } V = \frac{1}{\sigma \sqrt{d_q}} \text{SHGIIHB}. \quad (12)$$

To the best of our knowledge we are the first to explore kernel learning with positive random features.

Model	Space Complexity	Time Complexity
Softmax Transformer	$\mathcal{O}(L^2(1 + d_q/L))$	$\mathcal{O}(L^2 d_q)$
Performer	$\mathcal{O}(L(d_q + M + M d_q/L))$	$\mathcal{O}(L M d_q)$
LinearElu	$\mathcal{O}(L(d_q + d_q^2/L))$	$\mathcal{O}(L d_q^2)$
GMM-RKS	$\mathcal{O}(L(d_q + C(d_q^2/L + M d_q/L + M)))$	$\mathcal{O}(M C(d_q^2 + L d_q))$
GMM-PRF	$\mathcal{O}(L(d_q + C M d_q/L + C M))$	$\mathcal{O}(L C M d_q)$
FASTFOOD(RKS/PRF)	$\mathcal{O}(L(d_q + M + M d_q/L))$	$\mathcal{O}(L M d_q)$
GENERATIVE(RKS/PRF)	$\mathcal{O}(L(d_q + d_q^2/L + M d_q/L + M))$	$\mathcal{O}(M(d_q^2 + L d_q))$

Table 1: Space and time complexity of self-attention kernel of **KERNELIZED TRANSFORMERS** compared with Softmax Transformer (Vaswani et al., 2017), Performer (Choromanski et al., 2021), and LinearElu (Katharopoulos et al., 2020). L refers to length of context, $d_q=d_v$ is the query/key/value dimension, while M is the number of samples (where applicable).

2.2 ANALYSIS

In this section, we explore what can be said about the expressivity of the proposed linear **KERNELIZED TRANSFORMERS**. While our understanding of the properties of Transformers is rudimentary, we would still like to know whether the known properties extend to our models. For example, it has been shown that Softmax Transformers and its sparse counterparts are Turing complete (Pérez et al., 2019; Zaheer et al., 2020). This raises the question as to whether the proposed linear **KERNELIZED TRANSFORMERS** are also Turing complete?

It is easy to see that the generalized kernel self-attention (Eq. 1) includes the softmax kernel and hence should satisfy the properties of Softmax Transformer. Interestingly, we can also show that this property holds for GMM-RKS Transformers with number of components $C = 1$, (for a more systematic definition, see Section A.1). More formally,

Theorem 1: *The class of GMM-RKS Transformers with positional embeddings is Turing complete.*

For a detailed proof, see Appendix A.

Complexity Analysis: While all of our models have linear complexity with respect to the context length L , differences still exist amongst the various methods. Notable, GMM-RKS and **GENERATIVE** have quadratic time and space complexity in the query size d_q . Both the **FASTFOOD** methods avoid this approximation, whereas GMM-PRF avoids this by the use of a diagonal covariance matrix. The complexities are listed in Table 1.

Another factor that controls timing is the sampling of Ω . Sampling too frequently can lead to significant slowdowns whereas sampling too few times can lead to biased learning. For our experiments, we resample every 100 training iterations, although this can be changed. A detailed list of all hyperparameters along with implementation details are provided in Appendix B.

3 EXPERIMENTS

In this section, we evaluate experimentally the performance of the methods proposed in Section 2.1. Our experimental study is designed in order to answer the following questions:

Q1: How do **KERNELIZED TRANSFORMERS** perform in comparison to other efficient Transformer architectures in domains that require a long context? In particular, does learning the kernel function within the attention mechanism improve the performance compared to approaches that use fixed kernels, e.g. LinearElu, Performer?

Q2: What is the trade-off between accuracy and efficiency for **KERNELIZED TRANSFORMER** models?

Model	Complexity	ListOps	Text	Retrieval	Image	Avg.
Random Predictor	NA	10.00	50.00	50.00	10.00	30.00
Baseline Models						
Softmax Trans.(Vaswani et al., 2017)	$\mathcal{O}(L^2)$	36.38	64.27	57.46	42.44	50.13
Synthesizer(Tay et al., 2021a)	$\mathcal{O}(L^2)$	36.50	61.68	54.67	41.61	48.62
Sinkhorn(Tay et al., 2020a)	$\mathcal{O}((L/B)^2)$	34.20	61.20	53.83	41.23	47.62
Sparse Trans.(Child et al., 2019)	$\mathcal{O}(L\sqrt{L})$	35.78	63.58	59.59	44.24	50.79
Reformer(Kitaev et al., 2020)	$\mathcal{O}(L \log L)$	36.30	56.10	53.40	38.07	45.97
Local Attention (Parmar et al., 2018)	$\mathcal{O}(LK)$	15.95	52.98	53.39	41.46	40.94
Longformer(Beltagy et al., 2020)	$\mathcal{O}(LK)$	36.03	62.85	56.89	42.22	49.50
Linformer(Wang et al., 2020)	$\mathcal{O}(L)$	35.49	53.49	52.27	38.56	44.95
Big Bird(Zaheer et al., 2020)	$\mathcal{O}(L)$	37.08	64.02	59.29	40.83	50.31
LinearElu(Katharopoulos et al., 2020)	$\mathcal{O}(L)$	17.15	65.90	53.09	42.34	44.62
Performer(Choromanski et al., 2021)	$\mathcal{O}(L)$	36.00	65.40	53.82	<u>42.77</u>	49.50
Kernelized Transformers						
GMM-RKS (Eq. 5)	$\mathcal{O}(L)$	18.15	<u>66.20</u>	58.74	42.33	46.36
FASTFOOD-RKS (Eq. 6)	$\mathcal{O}(L)$	18.2	65.91	57.47	36.74	44.52
GENERATIVE-RKS (Eq. 7)	$\mathcal{O}(L)$	17.80	66.37	59.02	39.84	45.76
GMM-PRF (Eqs. 9, 10)	$\mathcal{O}(L)$	36.95	62.70	59.64	39.94	49.81
FASTFOOD-PRF (Eqs. 9, 12)	$\mathcal{O}(L)$	37.25	64.69	67.90	38.31	52.04
GENERATIVE-PRF (Eqs. 9, 11)	$\mathcal{O}(L)$	<u>37.10</u>	62.39	<u>67.18</u>	40.01	<u>51.67</u>

Table 2: Experimental results on the *LRA* benchmark. We report accuracy on the test set. The best model is in boldface and the second best is underlined. Accuracy scores for all baseline models are due to Tay et al. (2021b). Furthermore, L generally refers to the sequence length, K refers to the size of a local window and $B \ll L$ is a model specific parameter.

Q3: How do KERNELIZED TRANSFORMERS perform against Softmax Transformer (Vaswani et al., 2017) on short sequencing modelling tasks? Is there a degradation in performance when learning the kernel in linear time within self-attention?

3.1 DOES KERNEL LEARNING IMPROVE PERFORMANCE OF FIXED KERNEL METHODS ON LONGER SEQUENCES?

Long Range Arena (*LRA*; Tay et al. 2021b) is a diverse benchmark for the purpose of evaluating the ability of sequence models to reason under long-context scenarios. It includes tasks that vary both in terms of the context length (ranging from $1K$ to $4K$ tokens) as well as the data modalities (including text, mathematical expressions and natural images). We evaluate the KERNELIZED TRANSFORMER architectures introduced in Section 2.1 on various *LRA* tasks and compare their performance against well-established efficient Transformer architectures that are already included in the benchmark.

Datasets We consider the following datasets from the *LRA* benchmark¹:

- *ListOps*: A longer version of the *ListOps* task introduced by Nangia & Bowman (2018). This dataset is used to assess the ability of a model to handle hierarchically structured data. The data is synthetically generated using a sequence length of $2K$ (Tay et al., 2021b).
- *Text*: This task benchmarks the ability of each model to process and classify long documents. The IMDB movie reviews dataset (Maas et al., 2011) is used. In order to simulate a longer context, a byte/character-level setup is considered, resulting in a sequence length of $4K$. This setting further benchmarks the ability of the model to effectively compose multiple characters into higher-level phrases.
- *Retrieval*: This is a byte/character-level document matching task, where the model infers whether there is a citation link between two papers using the ACL Anthology Network (AAN; Radev et al.,

¹*Pathfinder* dataset was not considered as all baselines and our models do not outperform random prediction.




Figure 3: We compare the peak memory consumption (y-axis), performance (x-axis) and speed (denoted by area of circle) for the various KERNELIZED TRANSFORMER architectures on the two *LRA* tasks with sequence length equal to $4K$. Memory usage refers to per device memory usage across each GPU, whereas speed (steps per second) for each model is reported relative to the speed of Softmax Transformer (larger circles denote faster models).

2009) dataset. A sequence length of $4K$ is used and similar to Tay et al. (2021b) all models are deliberately prevented from using cross-attention between documents.

- *Image*: An $N \times N$ image is flattened into a sequence of N^2 pixels which is then provided as input to the model. The gray-scaled CIFAR10 image classification dataset (Krizhevsky, 2009) is used, resulting in a sequence length of 1024.

Setup: To ensure a fair comparison, we closely follow the same data preprocessing, data split, model size and training procedure as in (Tay et al., 2021b). Within each task, a common configuration is used across all KERNELIZED TRANSFORMER models based on the configuration specified in the *LRA* code repository². We outline the hyperparameters for all tasks in Table 4 in the Appendix.

Results: The results across all *LRA* tasks are summarized in Table 2. With the only exception of the *Image* classification task, KERNELIZED TRANSFORMER variants that learn the kernel function directly from the data in an end-to-end manner outperform the baseline models by occupying both best and second-best performances in the remaining of the *LRA* tasks. We find that KERNELIZED TRANSFORMERS based on PRFs tend to outperform their RKS counterparts which is also reflected on the average *LRA* score, with FASTFOOD-PRF being the best-performing model. This is in line with Choromanski et al. (2021) who show that PRF reduces the expected variance in Softmax as well as Gaussian kernels, stabilizing the training of Transformer.

3.2 TRADE-OFF BETWEEN ACCURACY AND EFFICIENCY

We benchmark the efficiency of each KERNELIZED TRANSFORMER in terms of peak memory usage and training speed and compare it against three baseline models from the *LRA* benchmark. Specifically, we compare against other efficient Transformer architectures that employ fixed kernel feature maps (e.g. LinearElu and Performer) as well as the Softmax Transformer which is one of the strongest baseline models (see Table 2). We conduct efficiency benchmarks on the two *LRA* tasks with sequence length equal to $4K$ in order to assess the efficiency of these methods in modelling tasks that require a long context (results for the other two datasets are included in the Appendix). Speed measurements (steps per second) refer to wall-clock training time (including overheads). In both cases experiments are conducted on 8 NVIDIA TITAN RTX GPUs.

The comparison is illustrated in Figure 3. On the *Text* task, GMM-RKS and LinearElu provide the best trade-off in terms accuracy and memory consumption. Specifically, GMM-RKS outperforms LinearElu in terms of accuracy, whereas LinearElu consumes less memory and is faster. The best performing model (GENERATIVE-RKS) consumes more memory than the rest of KERNELIZED

²Available at: <https://github.com/google-research/long-range-arena>

Model	SST2 (acc)	MRPC (acc/f1)	QQP (acc/f1)	MNLI- m/mm (acc/acc)	QNLI (acc)	WNLI (acc)	RTE (acc)	CoLA (MCor)
Softmax Trans.	0.81	0.70/0.82	0.83/0.76	0.64/0.64	0.68	0.56	0.6	0.18
FASTFOOD-RKS	0.83	0.71/0.82	0.81/0.74	0.57/0.57	0.64	0.59	0.56	0.13
GMM-RKS	0.80	0.70/0.82	0.77/0.69	0.47/0.48	0.60	0.61	0.57	0.07
GENERATIVE-RKS	0.81	0.70/0.82	0.81/0.73	0.59/0.58	0.63	0.62	0.58	0.16
FASTFOOD-PRF	0.81	0.71/0.82	0.81/0.74	0.56/0.57	0.64	0.59	0.58	0.12
GENERATIVE-PRF	0.80	0.71/0.82	0.80/0.74	0.56/0.56	0.61	0.60	0.55	0.10
GMM-PRF	0.82	0.71/0.82	0.81/0.74	0.56/0.56	0.64	0.59	0.59	0.21

Table 3: Results on the GLUE benchmark after fine-tuning on respective tasks. KERNELIZED TRANSFORMERS continue to be competitive to Transformers even in short context problems.

TRANSFORMER architectures, but it is still more efficient than the Softmax Transformer. In *Retrieval* the situation is much clearer, with FASTFOOD-PRF and GENERATIVE-PRF outperforming significantly other models in terms of accuracy while having very low memory consumption.

Lastly, in Figure 4, we report the peak memory consumption as the sequence length changes from 1K to 4K on the *Text* dataset. As expected, all our models have a linear increase in memory consumption with increasing sequence length, as opposed to the Softmax Transformer which has dramatic increase in memory consumption. Furthermore, Figure 5 in the Appendix reports the memory usage of each KERNELIZED TRANSFORMER across all datasets. We find that FASTFOOD-PRF and GENERATIVE-PRF are not only our best performing models on average, but they also consume the least memory among various KERNELIZED TRANSFORMERS across all datasets. Thus, among the models proposed in this paper, we can recommend FASTFOOD-PRF and GENERATIVE-PRF as the model that achieves the best accuracy with the least memory consumption.




Figure 4: Memory consumption vs sequence length

3.3 HOW DO KERNELIZED TRANSFORMERS PERFORM AS COMPARED TO SOFTMAX TRANSFORMER ON SHORT SEQUENCE TASKS?

We compare the KERNELIZED TRANSFORMERS and Softmax Transformer in a common transfer learning setting. We adopt the setting of BERT-like models (Devlin et al., 2019; Liu et al., 2019), except that we have fewer layers and heads (see Table 6 for details) and pre-train all models (including Softmax Transformer) on the WikiText-103 dataset (Merity et al., 2016) using non-contextual WordPiece embeddings (Wu et al., 2016). Pre-trained models are then fine-tuned on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019), a collection of resources for training and evaluating language understanding systems. All tasks in GLUE consist of either one or two sentences as input and can therefore be easily captured with a context of size 512. Since the context length is rather short, the difference between training and inference time across the various models is minimal. Instead, the main goal of this task is to assess how do KERNELIZED TRANSFORMERS compare against Softmax Transformers on a set of tasks where the later have been established as the *de-facto* architecture.

The results on all downstream GLUE tasks are shown in Table 3. Crucially, we demonstrate that there is no significant loss in performance compared to Softmax Transformers when kernelized variants are used on short language modelling tasks. As illustrated in Table 3, KERNELIZED TRANSFORMERS perform on par with the Softmax Transformer.

4 RELATED WORK

Efficient Transformers: A wide variety of approaches belong to the class of efficient Transformers (Tay et al., 2020b). Memory Compressed Transformer (Liu et al., 2018), one of the earliest attempts to alleviate the quadratic complexity of self-attention uses a convolution kernel (of size and strides K) to sub-sample keys and values reducing the time and memory complexity to $\mathcal{O}(L^2 K^{-1})$. Inspired by the notion of sparsity within the rather dense attention mechanism, Child et al. (2019) introduced sparse factorizations of the attention matrix to reduce the overall complexity to $\mathcal{O}(L\sqrt{L})$. Roy et al. (2020) proposed the Routing Transformer, which employs K -means clustering to learn dynamic sparse attention regions. The optimal number of K was found to be \sqrt{L} , making the overall complexity $\mathcal{O}(L\sqrt{L})$.

As a further improvement, Kitaev et al. (2020) proposed the Reformer, which reduces complexity to $\mathcal{O}(L \log L)$ by using locality-sensitive hashing (LSH) to group together symbols with similar embeddings. However, LSH constraints the keys to be identical to the queries and as a result this method cannot be used for decoding tasks where the keys and the queries differ. Ye et al. (2019) also propose a $\mathcal{O}(L \log L)$ algorithm, using binary partitions of data. There also exist works that mainly focus on memory reduction (Liu et al., 2018; Tay et al., 2020a). While these methods are also faster than Softmax Transformers, the asymptotic time complexity stays quadratic. Parmar et al. (2018) was one of the first models to use local attention to have $\mathcal{O}(L)$ complexity in both time and space. Beltagy et al. (2020) also proposed a $\mathcal{O}(L)$ method by using local sliding windows as well as global attention components. Zaheer et al. (2020) then proposed Big Bird, another sparse attention mechanism which combines global attention, random attention and local attention to reduce the complexity from $\mathcal{O}(L^2)$ to $\mathcal{O}(L)$. Concurrently, a similar construction was previously also used by Ainslie et al. (2020).

Kernel Learning: While kernel methods have long been used to solve non-linear statistical problems (Vapnik et al., 1997; Cortes & Vapnik, 1995; Schölkopf et al., 1998; Schölkopf & Smola, 2001), they traditionally scaled poorly with the number of data points thereby limiting their applicability on large datasets (Vapnik et al., 1997; Cortes & Vapnik, 1995; Schölkopf et al., 1998; Schölkopf & Smola, 2001; Hofmann et al., 2008). Prior to RKS, there have been other attempts including greedy basis selection techniques (Smola & Schölkopf, 2000), divide-and-conquer approaches (Hsieh et al., 2014; Zhang et al., 2013; Liu et al., 2020) as well as Nyström methods (Williams & Seeger, 2001), which provide data-dependent approximations to the kernel function. The RKS method itself has been revisited by several papers that either tried to improve the approximation quality (Yu et al., 2016; Choromanski et al., 2017; Li, 2017; Avron et al., 2016; Lyu, 2017), reduce the time and memory complexity of the algorithm (Le et al. 2013; Choromanski & Sindhvani 2016; Feng et al. 2015; Dao et al. 2017; e.g. Le et al. 2013 forms the base of FASTFOOD) or analyze theoretically the risk and generalization properties of the RKS mechanism (Sutherland & Schneider, 2015; Sun et al., 2018; Li et al., 2019b). A systematic survey of the algorithms, theory and applications of random feature methods for approximating kernel functions can be found in (Liu et al., 2021).

In the space of making RKS learnable, beyond the single stage approaches used in this paper, two-stage procedures have also been proposed (Sinha & Duchi, 2016; Li et al., 2019a; Bullins et al., 2018; Shen et al., 2019). Two-stage approaches involve an intermediate step in which the problem of kernel-alignment is solved (Cristianini et al., 2002). Kernel alignment refers to approximating a known target function with the kernel, which needs labelled data for the target function. Since we lack such data, we did not explore these methods.

5 CONCLUSION

In this paper, we bridged the gap between advances in kernel learning and efficient Transformers. In the process, we proposed several kernel learning methods for Transformers that increase the expressiveness of Transformers while keeping the computational complexity linear in sequence length. We showed that our proposed KERNELIZED TRANSFORMER are Turing-complete. Experimentally we demonstrated that our proposed models perform on par with, and possibly exceed the performance of existing efficient transformer architectures on long context tasks without falling behind on short context tasks. We also found that for some datasets such as *ListOps*, RKS based models tend to fall short of their PRF counterparts. Our experiments further demonstrate that the memory consumption of our models scales linearly with the sequence length.

REPRODUCIBILITY STATEMENT

We plan to open source the entire code of the `KERNELIZED TRANSFORMER` framework (including the implementation of all models as well as the code for replicating all of our experiments) before the camera ready version of the paper. As part of this submission, we include code for all the methods proposed by us along with instructions on how to reproduce results. A detailed description of all hyperparameters (for both *LRA* as well as *GLUE* benchmarks) has been included in Appendix B. Finally, regarding our theoretical contributions, we present a detailed theoretical analysis in Appendix A.

ETHICS STATEMENT

Our work is on making Transformers computationally efficient without losing expressiveness. Our models were evaluated on publicly available benchmark datasets. The datasets used in our work do not contain sensitive information to the best of our knowledge. We have read the ICLR Code of Ethics and adhere to it.

REFERENCES

- Milton Abramowitz and Irene A Stegun. Handbook of mathematical functions with formulas, graphs, and mathematical tables. national bureau of standards applied mathematics series 55. tenth printing, 1972.
- Joshua Ainslie, Santiago Ontañón, Chris Alberti, Philip Pham, Anirudh Ravula, and Sumit Sanghai. ETC: encoding long and structured data in transformers. *CoRR*, abs/2004.08483, 2020. URL <https://arxiv.org/abs/2004.08483>.
- Haim Avron, Vikas Sindhwani, Jiyan Yang, and Michael W. Mahoney. Quasi-monte carlo feature maps for shift-invariant kernels. *Journal of Machine Learning Research*, 17(120):1–38, 2016.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer, 2020.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020.
- Brian Bullins, Cyril Zhang, and Yi Zhang. Not-so-random features. In *International Conference on Learning Representations*, 2018.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers, 2020.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers, 2019.
- Krzysztof Choromanski and Vikas Sindhwani. Recycling randomness with structure for sublinear time kernel expansions. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, pp. 2502–2510. JMLR.org, 2016.
- Krzysztof M Choromanski, Mark Rowland, and Adrian Weller. The unreasonable effectiveness of structured random orthogonal embeddings. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021.
- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Mach. Learn.*, 20(3):273–297, September 1995. ISSN 0885-6125. doi: 10.1023/A:1022627411411.
- Nello Cristianini, John Shawe-Taylor, André Elisseeff, and Jaz Kandola. On kernel-target alignment. In T. Dietterich, S. Becker, and Z. Ghahramani (eds.), *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2002.
- Tri Dao, Christopher M De Sa, and Christopher Ré. Gaussian quadrature for kernel features. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
- Kun Fang, Xiaolin Huang, Fanghui Liu, and Jie Yang. End-to-end kernel learning via generative random fourier features, 2020.
- Chang Feng, Qinghua Hu, and Shizhong Liao. Random feature mapping with signed circulant matrix projection. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI’15*, pp. 3490–3496. AAAI Press, 2015. ISBN 9781577357384.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. Kernel methods in machine learning. *The Annals of Statistics*, 36(3):1171 – 1220, 2008. doi: 10.1214/009053607000000677.
- Cho-Jui Hsieh, Si Si, and Inderjit S. Dhillon. A divide-and-conquer solver for kernel support vector machines. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML’14*, pp. I–566–I–574. JMLR.org, 2014.
- John Ingraham, Vikas Garg, Regina Barzilay, and Tommi Jaakkola. Generative models for graph-based protein design. In H. Wallach, H. Larochelle, A. Beygelzimer, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 15820–15831. Curran Associates, Inc., 2019.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are RNNs: Fast autoregressive transformers with linear attention. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5156–5165. PMLR, 13–18 Jul 2020.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2020.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Quoc Le, Tamas Sarlos, and Alexander Smola. Fastfood - computing hilbert space expansions in loglinear time. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 244–252, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

- Chun-Liang Li, Wei-Cheng Chang, Youssef Mroueh, Yiming Yang, and Barnabas Poczos. Implicit kernel learning. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2007–2016. PMLR, 16–18 Apr 2019a.
- Ping Li. Linearized gmm kernels and normalized random fourier features. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17*, pp. 315–324, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450348874. doi: 10.1145/3097983.3098081.
- Zhu Li, Jean-Francois Ton, Dino Oglic, and Dino Sejdinovic. Towards a unified analysis of random Fourier features. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3905–3914. PMLR, 09–15 Jun 2019b.
- Fanghui Liu, Xiaolin Huang, Chen Gong, Jie Yang, and Li Li. Learning data-adaptive non-parametric kernels. *Journal of Machine Learning Research*, 21(208):1–39, 2020.
- Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan A. K. Suykens. Random features for kernel approximation: A survey on algorithms, theory, and beyond, 2021.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences. In *International Conference on Learning Representations*, 2018.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In H. Wallach, H. Larochelle, A. Beygelzimer, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 13–23. Curran Associates, Inc., 2019.
- Yueming Lyu. Spherical structured feature maps for kernel approximation. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2256–2264. PMLR, 06–11 Aug 2017.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
- Ali Madani, Bryan McCann, Nikhil Naik, Nitish Shirish Keskar, Namrata Anand, Raphael R. Eguchi, Po-Ssu Huang, and Richard Socher. Progen: Language modeling for protein generation, 2020.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models, 2016.
- Nikita Nangia and Samuel Bowman. 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.
- Junier B. Oliva, Avinava Dubey, Andrew G. Wilson, Barnabas Poczos, Jeff Schneider, and Eric P. Xing. Bayesian nonparametric kernel-learning. In Arthur Gretton and Christian C. Robert (eds.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 1078–1086, Cadiz, Spain, 09–11 May 2016. PMLR.

- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4055–4064, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A. Smith, and Lingpeng Kong. Random feature attention. *CoRR*, abs/2103.02143, 2021. URL <https://arxiv.org/abs/2103.02143>.
- Jorge Pérez, Javier Marinkovic, and Pablo Barceló. On the turing completeness of modern neural network architectures. *CoRR*, abs/1901.03429, 2019. URL <http://arxiv.org/abs/1901.03429>.
- Jorge Pérez, Javier Marinković, and Pablo Barceló. On the turing completeness of modern neural network architectures. In *International Conference on Learning Representations*, 2019.
- Dragomir R. Radev, Pradeep Muthukrishnan, and Vahed Qazvinian. The ACL Anthology network. In *Proceedings of the 2009 Workshop on Text and Citation Analysis for Scholarly Digital Libraries (NLP4DL)*, pp. 54–61, Suntec City, Singapore, August 2009. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training, 2018.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Proceedings of the 20th International Conference on Neural Information Processing Systems, NIPS’07*, pp. 1177–1184, Red Hook, NY, USA, 2007. Curran Associates Inc. ISBN 9781605603520.
- Eitan Richardson and Yair Weiss. On gans and gmms. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *bioRxiv*, 2020. doi: 10.1101/622803.
- Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers, 2020.
- Walter Rudin. *Fourier Analysis on Groups*. John Wiley & Sons, Ltd, 1990.
- Lars Ruthotto and Eldad Haber. An introduction to deep generative modeling, 2021.
- Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001. ISBN 0262194759.
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319, 1998. doi: 10.1162/089976698300017467.

- Zheyang Shen, Markus Heinonen, and Samuel Kaski. Harmonizable mixture kernels with variational fourier features. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 3273–3282. PMLR, 16–18 Apr 2019.
- Bernard W Silverman. *Density estimation for statistics and data analysis*, volume 26. CRC press, 1986.
- Aman Sinha and John C Duchi. Learning kernels with random features. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Alex J. Smola and Bernhard Schölkopf. Sparse greedy matrix approximation for machine learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML ’00, pp. 911–918, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 1558607072.
- Yitong Sun, Anna Gilbert, and Ambuj Tewari. But how does it work in theory? linear svm with random features. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Dougal J. Sutherland and Jeff Schneider. On the error of random fourier features. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI’15, pp. 862–871, Arlington, Virginia, USA, 2015. AUAI Press. ISBN 9780996643108.
- Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. Sparse sinkhorn attention. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9438–9447. PMLR, 13–18 Jul 2020a.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey, 2020b.
- Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. Synthesizer: Rethinking self-attention in transformer models, 2021a.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*, 2021b.
- Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhutdinov. Transformer dissection: An unified understanding for transformer’s attention via the lens of kernel. 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. 4344–4353, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1443.
- Vladimir Vapnik, Steven Golowich, and Alex Smola. Support vector method for function approximation, regression estimation and signal processing. In M. C. Mozer, M. Jordan, and T. Petsche (eds.), *Advances in Neural Information Processing Systems*, volume 9. MIT Press, 1997.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008. Curran Associates, Inc., 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*, 2019.
- Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, 2020.

- Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In T. Leen, T. Dietterich, and V. Tresp (eds.), *Advances in Neural Information Processing Systems*, volume 13. MIT Press, 2001.
- Andrew Wilson and Ryan Adams. Gaussian process kernels for pattern discovery and extrapolation. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1067–1075, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation, 2016.
- Zichao Yang, Andrew Wilson, Alex Smola, and Le Song. A la Carte – Learning Fast Kernels. In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pp. 1098–1106, San Diego, California, USA, 09–12 May 2015. PMLR.
- Zihao Ye, Qipeng Guo, Quan Gan, Xipeng Qiu, and Zheng Zhang. Bp-transformer: Modelling long-range context via binary partitioning, 2019.
- Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-Rice, and Sanjiv Kumar. Orthogonal random features. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for longer sequences. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 17283–17297. Curran Associates, Inc., 2020.
- Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression. In Shai Shalev-Shwartz and Ingo Steinwart (eds.), *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pp. 592–617, Princeton, NJ, USA, 12–14 Jun 2013. PMLR.

On Learning The Transformer Kernel – Appendix

A DETAILED PROOF OF THEOREM 1

A.1 DEFINITIONS

Transformer: A transformer consists of an Encoder and a Decoder, which in turn consist of several encoder and decoder layers respectively. A single encoder layer consists of an attention layer (Att) and a 2 layer feed-forward neural network (O) :

$$\mathbf{a}_i = Att(W_q \mathbf{x}_i, W_k \mathbf{X}, W_v \mathbf{X}) + \mathbf{x}_i \quad (13)$$

$$\mathbf{z}_i = O(\mathbf{a}_i) + \mathbf{a}_i \quad (14)$$

In our case, the feed-forward neural network uses perceptron activations (ie $f_{perc}(x) = 1$ iff $x > 0$ and 0 otherwise) and the attention is gaussian (discussed in detail later). The final layer of the encoder is followed by a couple of two layer output neural networks, which produce the Encoder Key (\mathbf{K}_e) and Encoder Value (\mathbf{V}_e) to be used by the decoder. In our proof, we assume these to have ReLU activation ($f_{ReLU}(x) = \max(0, x)$)

The decoder layers are similar to the encoder layer except for an additional cross attention layer which attends to the encoder output:

$$\mathbf{p}_i = Att(W_q \mathbf{y}_i, W_k \mathbf{Y}, W_v \mathbf{Y}) + \mathbf{y}_i \quad (15)$$

$$\mathbf{a}_i = Att(W'_q \mathbf{p}_i, \mathbf{K}_e, \mathbf{V}_e) + \mathbf{y}_i \quad (16)$$

$$\mathbf{z}_i = O(\mathbf{a}_i) + \mathbf{a}_i \quad (17)$$

Unlike the encoder, the decoder self attention in Eq. 15 can only attend to previous position. After the final layer we have a two layer feed-forward neural network with ReLU activation to produce the output. The decoder is initialised with a special *seed vector* and is repeatedly applied with the right shifted output of the last application as the input of the current application, until some termination condition is fulfilled.

Both the encoder and decoder can further use *position embeddings*, which have the same dimension as the output of each layer, and are added to the input prior to the first layer. These help in establishing the order of the input

Since the output of any unit of a layer is independent of values to its right, these do not change with time and can be cached. The output of the final layer of the rightmost cell can therefore be regarded as the model state encoding (v)

Turing Machine: A Turing Machine is an abstract construct which consists of a right infinite tape and a read-write head. Each cell of the tape can hold one of many symbols from a predefined alphabet Σ which includes a special blank symbol b . Additionally, the read-write head can be in one of many possible states within the state-space Q which includes a special initial state q_{init} and a subset of final states F .

Initially, the tape contains the input followed by an infinite number of blank symbols, while the head starts off in the last non-blank cell. In each step, the head executes in accordance with a transition function $T(s^{(i)}, q^{(i)}) = (v^{(i)}, q^{(i+1)}, m^{(i)})$, ie, based on the symbol currently under the head and the current state, it decides the symbol it wants to overwrite the current symbol with, the state it will be in the next step and the direction it wants to move, which can be either left (-1) or right (1). We assume that the transition function already makes sure that the head never moves left from the leftmost cell.

For the purpose of our proof, we additionally define $c^{(i)}$ as the index of the cell to which the head currently points, $\ell(i)$, which represents the step number when the head last pointed to the current cell, ie $\ell(i) = \max\{j | c^{(j)} = c^{(i)}\}$. In the special case where the current cell is being visited for the first time, we have $\ell(i) = i - 1$

A.2 THE PROOF

Our proof is based on the similar proof in Pérez et al. (2019). Any symbols not explicitly defined have same meanings from that paper. We begin the proof by defining our model encoding(\mathbf{v}):

$$\begin{aligned} \mathbf{v} = & [\mathbf{q}_1, \mathbf{s}_1, x_1, x_2 \\ & \mathbf{q}_2, \mathbf{s}_2, x_3, x_4, x_5, x_6, \\ & x_7, \mathbf{s}_4, x_8, \\ & x_9, x_{10}, x_{11}, x_{12}] \end{aligned} \quad (18)$$

where $q_i \in \mathbb{Q}^{|Q|}$, $s_i \in \mathbb{Q}^{|S|}$, and $x_i \in \mathbb{Q}$, giving a total model size of $2|Q| + 3|S| + 14$. Hereafter, $\llbracket x \rrbracket$ represents the one-hot encoding for the state x or symbol x depending on the position it is being used in. $\mathbf{0}_q$ represents all 0's in a state field, and represents the q_{copy} state discussed later, while $\mathbf{0}_s$ represents all 0's in a symbol field, and represents the blank symbol. Further, $\beta^{(i)} = \min(i, n)$ where n is the size of the encoder and $\alpha^{(i)}$ represents the symbol at position $\beta^{(i)}$ in the encoder. We assume that atleast the last cell of the encoder contains a blank symbol.

This differs from Pérez et al. (2019) in the addition of a fourth scalar in the first group, in which we intend to store the current position $c^{(i)}$ of the head.

Our invariant is that \mathbf{y}_i the output from the decoder at timestep i , stores:

1. The current state of the Turing Machine($q^{(i)}$)
2. The symbol under the head($s^{(i)}$)
3. The direction of movement of the head in the previous timestep ($m^{(i-1)}$)
4. The current position of the head($c^{(i)}$)

In all, we get $\mathbf{y}_i = [\llbracket q^{(i)} \rrbracket, \llbracket s^{(i)} \rrbracket, m^{(i-1)}, c^{(i)}, 0, \dots, 0]$

Positional Embeddings: The last group($x_9, x_{10}, x_{11}, x_{12}$) is dedicated to the positional embeddings, which are given as $(1, i, \frac{1}{i}, \frac{1}{i^2})$. These same embeddings are added on both the Encoder and Decoder side.

Encoder: The encoder consists of a single layer. It gets as input the symbol at position i and the positional embeddings, ie $input_i = [\mathbf{0}_q, \mathbf{0}_s, 0, 0, \mathbf{0}_q, \llbracket s^{(i)} \rrbracket, 0, 0, 0, 0, i, \mathbf{0}_s, 0, 1, i, \frac{1}{i}, \frac{1}{i^2}]$ which has a trivial attention layer(ie, one that outputs all zeroes) and a feed forward layer which separates the positional embeddings from the symbols, giving $\mathbf{k}_i^e = [0, \dots, 0, i, -1, 0, 0]$ and $\mathbf{v}_i^e = [\mathbf{0}_q, \mathbf{0}_s, 0, 0, \mathbf{0}_q, \llbracket s^{(i)} \rrbracket, 0, 0, 0, 0, i, \mathbf{0}_s, 0, 0, 0, 0, 0]$.

Decoder Layer 1: The first layer of the decoder calculates the next state, the symbol to be written and the direction of movement of the head. This includes 2 cases:

1. Initially, the Decoder starts off with in the state q_{copy} . While the state is still q_{copy} , the head writes the symbol at the i^{th} position in the encoder and moves right, until a blank symbol is seen. Once a blank symbol is reached, the tape rewrites the blank symbol, moves left and the state changes to q_{init} .
2. Once we move into q_{init} , the output is fully defined by the current state and symbol under the head.

To facilitate the first case, we make use of the cross attention layer, to get

$$\begin{aligned} Att(\mathbf{q}, \mathbf{K}^e, \mathbf{V}^e) &= [0, \dots, 0, \\ & \mathbf{0}_q, \llbracket \alpha^{(i)} \rrbracket, 0, 0, 0, 0, \\ & \beta^{(i)}, \mathbf{0}_s, 0, \\ & 0, 0, 0, 0] \\ &= \mathbf{v}_{\beta^{(i)}}^e \end{aligned} \quad (19)$$

The details of this process are explained in lemma S.1(see sec. A.4) Adding in the residual connection, we have:

$$\begin{aligned} \mathbf{a}_i^1 = & [\llbracket q^{(i)} \rrbracket, \llbracket s^{(i)} \rrbracket, m^{(i-1)}, c^{(i)} \\ & \mathbf{0}_q, \llbracket \alpha^{(i)} \rrbracket, 0, 0, 0, 0, \\ & \beta^{(i)}, \mathbf{0}_s, 0, \\ & 1, i+1, \frac{1}{(i+1)}, \frac{1}{(i+1)^2}] \end{aligned} \quad (20)$$

Hereafter, we make use of the feed-forward layer to get:

$$\begin{aligned} O(\mathbf{a}_i^1) = & [-\llbracket q^{(i)} \rrbracket, -\llbracket s^{(i)} \rrbracket, -m^{(i-1)}, m^{(i)}, \\ & \llbracket q^{(i+1)} \rrbracket, \llbracket \bar{v}^{(i)} \rrbracket, m^{(i)}, m^{(i-1)}, 0, 0, \\ & 0, \dots, 0 \\ & 0, \dots, 0] \end{aligned} \quad (21)$$

If the state is q_{init} then we set $\llbracket \bar{v}^{(i)} \rrbracket = \mathbf{0}_s$, else we have $\llbracket \bar{v}^{(i)} \rrbracket = \llbracket v^{(i)} \rrbracket$. Note that this gives us $\llbracket \bar{v}^{(i)} \rrbracket + \llbracket \alpha^{(i)} \rrbracket = \llbracket v^{(i)} \rrbracket$.

To get all the required values, we first project $\llbracket q^{(i)} \rrbracket$ and $\llbracket s^{(i)} \rrbracket$ to a one-hot encoding of $Q \times \Sigma$. from there, we can calculate all the required values in a look-up table fashion. if the state is q_{init} then we set $\llbracket v^{(i)} \rrbracket = \mathbf{0}_s$

The final output of this layer is then:

$$\begin{aligned} \mathbf{z}_i^1 = & [0, \dots, 0, c^{(i+1)} \\ & \llbracket q^{(i+1)} \rrbracket, \llbracket v^{(i)} \rrbracket, m^{(i)}, m^{(i-1)}, 0, 0, \\ & \beta^{(i)}, \mathbf{0}_s, 0, \\ & 1, i+1, \frac{1}{(i+1)}, \frac{1}{(i+1)^2}] \end{aligned} \quad (22)$$

Decoder Layer 2: In this layer we calculate the symbol under the head in the next timestep. In order to do so, we first use the self attention layer to calculate $\llbracket v^{(\ell(i+1))} \rrbracket$ and $(\ell(i+1))$ (For details, see sec A.7.):

$$\begin{aligned} Att(W_q^2 \mathbf{z}_i^2, W_k^2 \mathbf{Z}^2, W_v^2 \mathbf{Z}^2) = & [0, \dots, 0, \\ & 0, \dots, 0, \\ & 0, \llbracket v^{(\ell(i+1))} \rrbracket, (\ell(i+1)), \\ & 0, 0, 0, 0] \end{aligned} \quad (23)$$

Adding the residual layer, we have

$$\begin{aligned} \mathbf{a}_i^2 = & [0, \dots, 0, c^{(i+1)} \\ & \llbracket q^{(i+1)} \rrbracket, \llbracket v^{(i)} \rrbracket, m^{(i)}, m^{(i-1)}, 0, 0, \\ & \beta^{(i)}, \llbracket v^{(\ell(i+1))} \rrbracket, (\ell(i+1)), \\ & 1, i+1, \frac{1}{(i+1)}, \frac{1}{(i+1)^2}] \end{aligned} \quad (24)$$

The feed-forward layer then gives $O_2(\mathbf{a}_i^2) = [\llbracket q^{(i+1)} \rrbracket, \llbracket v^{(\ell(i+1))} \rrbracket - f_{perc}((\ell(i+1) + 2 - (i+1))), m^{(i-1)}, 0, -M, \dots, -M]$ where M is a large negative value. The perceptron function in the \mathbf{s}_1 is added positionwise, and is 0 unless $\ell(i+1) = i$. In this special case, it makes \mathbf{s}_1 contain only 0 or -1 which is converted into $\mathbf{0}_s$ by the ReLU activation in the output MLP. The same is also true for every field after the first 4, where we add a large negative value to make the ReLU output 0.

A.3 THE ATTENTION MECHANISM

The attention mechanism in the Gaussian kernel is defined as follows:

$$\Phi(\mathbf{x}) = \frac{2^{d_q/2}}{\sqrt{k}} [\sin(\omega_0^T \mathbf{x}), \cos(\omega_0^T \mathbf{x}), \dots, \sin(\omega_{k-1}^T \mathbf{x}), \cos(\omega_{k-1}^T \mathbf{x})], \quad (25)$$

$$\text{where } \omega_i \sim \mathcal{N}(\mu_{\mathbf{c}}, \Sigma_{\mathbf{c}}) \quad (26)$$

$$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{V}(\text{ColNorm}(\Phi(\mathbf{Q})^T \Phi(\mathbf{K}))) \quad (27)$$

However, for the proof construction, we use a hard version of this attention mechanism. To begin with, we replace the kernel with its equivalent dual, using lemma S.2(sec A.6). In our construction, we do not require learnable means and variances, so we fix them to be 0_{d_q} and \mathbb{I} hereafter:

$$Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \sum_{i=0}^{h-1} \mathbf{W}_O^i \mathbf{V} \left(\text{ColNorm} \left[\left[e^{-\frac{\|\mathbf{q}_l - \mathbf{k}_m\|^2}{2}} \right]_{l=0, m=0}^{d-1, d-1} \right] \right) \quad (28)$$

where $\llbracket f(l, m) \rrbracket_{l=0, m=0}^{\alpha, \beta}$ denotes an $\alpha \times \beta$ matrix whose $(l, m)^{th}$ entry is $f(l, m)$, $\text{ColNorm}(\mathbf{X})$ indicates the matrix \mathbf{X} with its columns normalised to and d is the dimension of the query/key vector.

While this definition does not seem to allow multiplying the exponent, one must remember that the query and key matrices are calculated using projection matrices, and any required scalar factor can be incorporated into them. Therefore, we define hard gaussian attention as:

$$\text{score}(\mathbf{u}, \mathbf{v}) = -\|\mathbf{u} - \mathbf{v}\|^2 \quad (29)$$

Hard attention is then computed as

$$Att(\mathbf{q}_i, \mathbf{K}, \mathbf{V}) = \frac{\sum_{j=0}^{n-1} \mathbb{I}[\text{score}(\mathbf{q}_i, \mathbf{k}_j) = (\max_{j'} \text{score}(\mathbf{q}_i, \mathbf{k}_{j'}))] \mathbf{v}_j}{\sum_{j=0}^{n-1} \mathbb{I}[\text{score}(\mathbf{q}_i, \mathbf{k}_j) = \max_{j'} (\text{score}(\mathbf{q}_i, \mathbf{k}_{j'}))]} \quad (30)$$

Here \mathbb{I} in the indicator function.

A.4 LEMMA S.1

A.5 STATEMENT

Given

$$\begin{aligned} \mathbf{q} &= [\dots, \dots, \dots, 1, i, \dots, \dots] \\ \mathbf{k}_j^e &= [0, \dots, 0, \\ &\quad 0, \dots, 0, \\ &\quad 0, \dots, 0, \\ &\quad j, -1, 0, 0] \\ \mathbf{v}_j^e &= [0, \dots, 0, \\ &\quad \mathbf{0}_q, \llbracket s^{(j)} \rrbracket, 0, 0, 0, 0, \\ &\quad j, \mathbf{0}_s, 0, \\ &\quad 0, 0, 0, 0] \end{aligned} \quad (31)$$

For $j \in \{0, \dots, n\}$, we need a construction that gives

$$\begin{aligned} Att(\mathbf{q}, \mathbf{K}^e, \mathbf{V}^e) &= [0, \dots, 0, \\ &\quad \mathbf{0}_q, \llbracket \alpha^{(i)} \rrbracket, 0, 0, 0, 0, \\ &\quad \beta^{(i)}, \mathbf{0}_s, 0, \\ &\quad 0, 0, 0, 0] \\ &= \mathbf{v}_{\beta^{(i)}}^e \end{aligned} \quad (32)$$

A.5.1 PROOF

Note that while the key and value comes from the encoder, and is therefore fixed, the query comes from the decoder and thus can be projected as we please. It is easy to construct a projection matrix that gives $W_Q \mathbf{q} = [0, \dots, i, -1, 0, 0]$. Then we have $score(\mathbf{q}, \mathbf{k}_{j'}) = -||i = j||^2 = -(i - j)^2$, whose maxima on j' is unique and occurs at $i = \beta^{(j)}$. Thus, we have $Att(\mathbf{q}, \mathbf{K}^e, \mathbf{V}^e) = \mathbf{v}_{\beta^{(j)}}^e$, which is exactly what we wanted.

A.6 LEMMA S.2

A.6.1 STATEMENT

Let

$$\phi(x) := \frac{1}{\sqrt{M}} [\cos(\omega_1^T x), \dots, \cos(\omega_k^T x), \sin(\omega_1^T x), \dots, \sin(\omega_k^T x)] \quad (2)$$

We want to show that if $\omega \sim \mathcal{N}(0, \Sigma)$ then the kernel ϕ as defined above corresponds to the *GMM kernel*, ie.

$$\Phi(\mathbf{X})\Phi(\mathbf{Y}) \approx e^{-\frac{(\mathbf{x}-\mathbf{y})^T \Sigma'^{-1} (\mathbf{x}-\mathbf{y})}{2}} \quad (33)$$

where $\Sigma' = 2\Sigma^{-1}$ and $\Phi = 2^{m/4} \sqrt{\frac{M}{k}} \phi$. The proportionality is to avoid clutter, and gets nullified by normalisation.

A.6.2 PROOF

We start with Eq. 7.4.6 in Abramowitz & Stegun (1972) and extend it to vectors. We have,

$$\begin{aligned} & \int_{\mathbb{R}^m} e^{-\mathbf{t}^T \mathbf{A} \mathbf{t}} \cos(2\mathbf{t}^T \mathbf{x}) d\mathbf{t} \\ &= \int_{\mathbb{R}^m} e^{-(a_0 t_0^2 + \sum_{i=1}^{m-1} a_i t_i^2)} \cos\left(2x_0 t_0 + 2 \sum_{i=1}^{m-1} t_i x_i\right) dt_0 dt_1 \dots dt_{m-1} \\ &= \int_{\mathbb{R}^m} e^{-(a_0 t_0^2 + \sum_{i=1}^{m-1} a_i t_i^2)} \cos(2x_0 t_0) \cos\left(2 \sum_{i=1}^{m-1} t_i x_i\right) dt_0 dt_1 \dots dt_{m-1} \\ &\quad - \int_{\mathbb{R}^m} e^{-(a_0 t_0^2 + \sum_{i=1}^{m-1} a_i t_i^2)} \sin(2x_0 t_0) \sin\left(2 \sum_{i=1}^{m-1} t_i x_i\right) dt_0 dt_1 \dots dt_{m-1} \end{aligned}$$

The second integral involving sin is odd and therefore evaluates to 0. That leaves us with:

$$\begin{aligned} &= \int_{\mathbb{R}^m} e^{-(a_0 t_0^2 + \sum_{i=1}^{m-1} a_i t_i^2)} \cos(2x_0 t_0) \cos\left(2 \sum_{i=1}^{m-1} t_i x_i\right) dt_0 dt_1 \dots dt_{m-1} \\ &= \int_{\mathbb{R}^{m-1}} \left(\int_{-\infty}^{\infty} e^{-a_0 t_0^2} \cos(2x_0 t_0) dt_0 \right) e^{-\sum_{i=1}^{m-1} a_i t_i^2} \cos\left(2 \sum_{i=1}^{m-1} t_i x_i\right) dt_1 \dots dt_{m-1} \\ &= \frac{1}{2} \sqrt{\frac{\pi}{a_0}} e^{-\frac{x_0^2}{a_0}} \int_{\mathbb{R}^{m-1}} e^{-\sum_{i=1}^{m-1} a_i t_i^2} \cos\left(2 \sum_{i=1}^{m-1} t_i x_i\right) dt_1 \dots dt_{m-1} \end{aligned}$$

This process can now be repeated for every dimension of t and x to finally give:

$$\int_{\mathbb{R}^m} e^{-\mathbf{t}^T \mathbf{A} \mathbf{t}} \cos(2\mathbf{t}^T \mathbf{x}) d\mathbf{t} = \frac{\pi^{m/2}}{2^m |\mathbf{A}|^{1/2}} e^{-\mathbf{x}^T \mathbf{A}^{-1} \mathbf{x}} \quad (34)$$

Here, \mathbf{A} is the diagonal covariance matrix, with diagonal entries a_0, a_1, \dots, a_{m-1} . In order to generalise to a full covariance matrix Σ , note that since a covariance matrix is symmetric by definition,

we can define $\Sigma = Q A Q^T$ where A is diagonal and Q is Orthonormal. Then, we have:

$$\begin{aligned}
\int_{\mathbb{R}^m} e^{-\mathbf{t}^T \Sigma \mathbf{t}} \cos(2\mathbf{t}^T \mathbf{x}) d\mathbf{t} &= \int_{\mathbb{R}^m} e^{-\mathbf{t}^T Q A Q^T \mathbf{t}} \cos(2\mathbf{t}^T (Q Q^T) \mathbf{x}) d\mathbf{t} \\
&= \int_{\mathbb{R}^m} e^{-(Q^T \mathbf{t})^T A (Q^T \mathbf{t})} \cos(2(Q^T \mathbf{t})^T (Q^T \mathbf{x})) d(Q^T \mathbf{t}) \\
&= \frac{\pi^{m/2}}{2^m |\mathbf{A}|^{1/2}} e^{-(Q^T \mathbf{x})^T A^{-1} (Q^T \mathbf{x})} \\
&= \frac{\pi^{m/2}}{2^m |\Sigma|^{1/2}} e^{-\mathbf{x}^T \Sigma^{-1} \mathbf{x}}
\end{aligned} \tag{35}$$

Here we repeatedly use the fact that $|Q| = 1$. Now, coming back to our Kernel function,

$$\begin{aligned}
\Phi(\mathbf{X})\Phi(\mathbf{Y}) &= \frac{2^{m/2}}{k} \sum_{i=0}^{k-1} (\cos(\omega_i^T \mathbf{x}) \cos(\omega_i^T \mathbf{y}) + \sin(\omega_i^T \mathbf{x}) \sin(\omega_i^T \mathbf{y})) \\
&= \frac{2^{m/2}}{k} \sum_{i=0}^{k-1} \cos(\omega_i(\mathbf{x} - \mathbf{y})) \\
&\approx 2^{m/2} \mathbb{E}_{\omega_i \sim \mathcal{N}(\mathbf{0}_m, \Sigma)} \cos(\omega(\mathbf{x} - \mathbf{y})) \\
&= 2^{m/2} \int_{\mathbb{R}^m} \frac{|\Sigma|^{-1/2}}{(2\pi)^{m/2}} e^{-\omega^T \Sigma^{-1} \omega / 2} \cos(2\omega^T \frac{\mathbf{x} - \mathbf{y}}{2}) d\omega \\
&= e^{-\frac{(\mathbf{x} - \mathbf{y})^T \Sigma^{-1} (\mathbf{x} - \mathbf{y})}{2}}
\end{aligned} \tag{36}$$

The only imprecision stems from the third step where we convert the sum to an expectation, but this is bounded by a standard deviation or $\frac{\sigma}{\sqrt{k}}$ where σ is the standard deviation of the true distribution, which in turn can be bounded by the range of the support (2 for the cosine function). Therefore, by choosing a high enough k , we can make the estimate arbitrarily accurate.

A.7 LEMMA S.3

A.7.1 STATEMENT

Given,

$$\begin{aligned}
\mathbf{z}_i^1 &= [0, \dots, 0, c^{(i+1)}, \\
&\quad \llbracket q^{(i+1)} \rrbracket, \llbracket v^{(i)} \rrbracket, m^{(i)}, m^{(i-1)}, 0, 0, \\
&\quad \beta^{(i+1)}, \mathbf{0}_s, 0, \\
&\quad 1, (i+1), \frac{1}{(i+1)}, \frac{1}{(i+1)^2}]
\end{aligned} \tag{37}$$

we need a construction that gives

$$\begin{aligned}
Att(W_q^2 \mathbf{z}_i^2, W_k^2 \mathbf{z}^2, W_v^2 \mathbf{z}^2) &= [0, \dots, 0, \\
&\quad 0, \dots, 0, \\
&\quad 0, \llbracket v^{(\ell(i+1))} \rrbracket, (\ell(i+1)), \\
&\quad 0, 0, 0, 0]
\end{aligned} \tag{38}$$

A.7.2 PROOF

We set the weight matrices to get $\mathbf{q}_j = W_q^2 \mathbf{z}_j^2 = [0, \dots, 0, c^{(j+1)}, 0, 0]$, $\mathbf{k}_j = W_k^2 \mathbf{z}_j^2 = [0, \dots, 0, c^{(j)} = c^{(j+1)} - m^{(i)}, 0, \frac{1}{(j+1)}]$ and $\mathbf{v}_j = W_v^2 \mathbf{z}_j^2 = [0, \dots, 0, \llbracket v^{(j)} \rrbracket, j, 0, 0, 0, 0]$. All these are partial permutations and therefore can be done using appropriate binary matrices.

Note that the required output is exactly the value at $j = \ell(i + 1)$, so it is sufficient to show that the $score(\mathbf{q}_i, \mathbf{k}_j)$ is maximised in j for $j = \ell(i + 1)$, i.e.

$$j = \begin{cases} \max\{j' | c^{(j')} = c^{(i+1)}\}, & \text{if } \exists j' \text{ s.t. } c^{(j')} = c^{(i+1)} \\ i, & \text{otherwise} \end{cases} \quad (39)$$

Now we have $score(\mathbf{q}_i, \mathbf{k}_j) = -(c^{(i+1)} - c^{(j)})^2 - \frac{1}{(j+1)^2}$. For all j such that $c^{(i+1)} \neq c^{(j)}$, the $score$ is almost -1 since c is an integer. If there $\exists j' \text{ s.t. } c^{(j')} = c^{(i+1)}$ then the corresponding $score$ is greater than -1 , and the maxima is achieved at the highest such value of j . If such j does not exist however, then $\forall j < i$, $score(\mathbf{q}_i, \mathbf{k}_j) < -1 - \frac{1}{(i+1)^2}$ and therefore, the maxima is achieved at $j = i$.

B EXPERIMENTAL DETAILS

B.1 SOURCE CODE

We implemented KERNELIZED TRANSFORMERS in Python 3 and PyTorch (Paszke et al., 2019) and plan to open-source the code for reproducing all experiments upon acceptance.

B.2 HYPERPARAMETERS FOR LRA TASKS

Parameter	ListOps	Text	Retrieval	Image
Batch Size	32	32	32	256
Learning Rate	5×10^{-3}	5×10^{-2}	5×10^{-2}	5×10^{-4}
Training Steps/Epochs	10K/NA	20K/NA	5K/NA	NA/200
Optimizer	Adam with Weight Decay ($\beta_1 = 0.9, \beta_2 = 0.98$)			
Weight Decay	0.1	0.1	0.1	0.0
Warmup Steps	1000	8000	8000	175
Scheduler	Sqrt Decay	Sqrt Decay	Sqrt Decay	Cosine Decay
Loss	Cross Entropy			
Sequence Length	2000	4000	4000	1024
Num. Layers	6	4	4	1
Num. Heads	8	4	4	8
Embedding Dim.	512	256	128	128
Key/Query/Value Dim.	64	64	32	8
Feedforward Dim.	2048	1024	512	128
Dropout Rate	0.1	0.1	0.1	0.3
Activation Function	Gelu	Gelu	Gelu	Gelu
Positional Encoding	Sinusoidal	Sinusoidal	Sinusoidal	Learnable
Pooling Mode	CLS	CLS	CLS	CLS

Table 4: Hyperparameters for *LRA* tasks.

Model	ListOps	Text	Retrieval	Image
GMM-RKS	256	128	64	128
FASTFOOD-RKS	64	64	32	8
GENERATIVE-RKS	256	256	128	128
GMM-PRF	256	256	128	128
FASTFOOD-PRF	64	64	32	8
GENERATIVE-PRF	128	128	64	8

Table 5: Number of random samples M used within each KERNELIZED TRANSFORMER.

Further Notes:

- To benchmark memory in Figure 3, we used a batch size of 32 for *Text* and a batch size of 2 for *Retrieval*.
- For GMM-RKS and GMM-PRF the number of components C in the mixture was set to $C = 2$.

B.3 HYPERPARAMETERS FOR GLUE TASKS

Parameter	Value(s)
Pre-Training Batch Size	64
Batch Size	64
Pre-Training Learning Rate (η_{pre})	5×10^{-4}
Pre-Training Learning Rate at Step i	$\min(\frac{i}{10000}, \frac{I_{\text{pre}} - i}{I_{\text{pre}} - 10000}) * \eta_{\text{pre}}$
Training Learning Rate (η_{train})	$\{2 \times 10^{-3}, 1 \times 10^{-4}, 5 \times 10^{-4}, 2 \times 10^{-5}, 5 \times 10^{-6}\}$
Training Learning Rate at Step i (η_{train})	$\min(\frac{10i}{I_{\text{tune}}}, \frac{I_{\text{tune}} - i}{0.9 * I_{\text{tune}}}) * \eta_{\text{train}}$
Pre-Training Epochs	5
Training Epochs	10
Optimizer	Adam with Weight Decay ($\beta_1 = 0.9, \beta_2 = 0.999$)
Weight Decay	0.01
Loss	Cross Entropy
Sequence Length	512
Num. Layers	3
Num. Heads	10
Embedding Dimension	300
Key/Query/Value Dimension	64
Transformer Feedforward Dimension	512
Classifier Feedforward Dimension	128
Dropout Rate	0.1
Transformer Activation Function	Gelu
Classifier Activation Function	Tanh
Positional Encoding	Sinusoidal
Pooling Mode	CLS
Num. of Samples from Distribution	$\{64, 128\}$

Table 6: Hyperparameters for *GLUE* tasks. Where multiple parameters were tried, they are listed in curly brackets. I_{pre} denotes the total number of pre-training steps, whereas I_{tune} denotes the total number of fine-tuning steps on each *GLUE* task.

Our model has significantly fewer number of parameters as compared to Devlin et al. (2019) and therefore we perform poorer on than them on all datasets. They use 24 layers with 16 heads each. If reported in the same order as the columns of Table 3, their numbers would look like: 97.5, 89.3/85.4, 72.1/89.3, 87.2/86.4, 92.7, 65.1, 70.1.

While we had to limit model sizes due to resource limitations, this handicaps all models equally, and therefore should not prevent comparison across various models reported in our paper.

B.4 FURTHER RESULTS ON EFFICIENCY BENCHMARKS




Figure 5: Peak memory used by KERNELIZED TRANSFORMERS across different datasets.




Figure 6: We demonstrate the peak memory consumption (y-axis) and performance (x-axis) of the various Kernelized Transformer architectures on the ListOps and Image dataset from LRA. Memory usage refers to per device memory usage across each GPU.

C ABLATION STUDIES

C.1 FASTFOOD ATTENTION

In the main paper, we use FastFood-SGB, which has all the diagonal matrices learnable. However, B and G matrices have a very special structure (their elements being drawn from $Bernoulli_{\{-1,1\}}(0.5)$ and $\mathcal{N}(0,1)$ respectively), which is lost if we make them learnable. Therefore, it makes sense to have *FastFood-S*, which only has S learnable. Finally, we can also have everything fixed, giving us the basic *FastFood* version. The results of these two versions, along with the original FastFood-SGB kernel on the *GLUE* benchmark are summarised in Table 7. As one can see, FastFood-SGB is either the best or close to it except for WNLI and CoLA, therefore we choose to use this version for our main analysis.

Dataset	SST2 (acc)	MRPC (acc)	MRPC (f1)	QQP (acc)	QQP (f1)	MNLI (mat)	MNLI (mis)	QNLI (acc)	WNLI (acc)	RTE (acc)	CoLA (MCorr)
FastFood	0.814	0.713	0.820	0.811	0.738	0.571	0.568	0.629	0.634	0.563	0.152
FastFood-S	0.807	0.706	0.822	0.810	0.741	0.571	0.571	0.642	0.606	0.570	0.101
FastFood-SGB	0.828	0.707	0.820	0.810	0.739	0.569	0.572	0.638	0.592	0.563	0.129

Table 7: Ablation studies using FastFood variants on the *GLUE* benchmark.