

SPECTRAFORMER: A UNIFIED RANDOM FEATURE FRAMEWORK FOR TRANSFORMER

Anonymous authors

Paper under double-blind review

ABSTRACT

Linearization of attention using various kernel approximation and kernel learning techniques has shown promise. Past methods use a subset of combinations of component functions and weight matrices within the random features paradigm. We identify the need for a systematic comparison of different combinations of weight matrices and component functions for attention learning in Transformer. In this work, we introduce *Spectraformer*, a unified framework for approximating and learning the kernel function in linearized attention of the Transformer. We experiment with broad classes of component functions and weight matrices for three textual tasks in the LRA benchmark. Our empirical findings indicate that different kernels are good at different tasks and that kernel choice is fundamental to performant models. Our code is available at: <https://anonymous.4open.science/r/spectraformer-8A97>.

1 INTRODUCTION

Transformer (Vaswani et al., 2017) has revolutionized the landscape of natural language processing (NLP) and forms the basis of almost all state-of-the-art language models. Its influence has reached beyond NLP into computer vision (Han et al., 2023), speech processing (Lin et al., 2022) and other fields. Compared to its precursors like the LSTMs, Transformer leverages parallelism since it is fully based on the attention mechanism. Therefore, the performance of Transformer is tied to the effective use of attention. Attention in the Original Transformer (referred to as ‘**OT**’ hereafter) uses the softmax function which is quadratic in time complexity. However, softmax is not fundamental to attention. The softmax serves the role of a **compatibility function**, which captures the degree of compatibility between two tokens (see Equations 1 and 2). Extensive experiments with other compatibility functions show that the softmax is not the only possibility, and that the ‘best’ **compatibility function** depends on the task (Tsai et al., 2019).

Replacing the softmax formulation with alternatives may allow us to reduce the time complexity of attention computation. In this paper, we focus on the subset of alternatives that employs kernel functions. Kernel functions have been historically used in machine learning algorithms to simplify the estimation of parameters (Hofmann et al., 2008). Tsai et al. (2019) show that attention can be formulated using kernels and then be linearized (Katharopoulos et al., 2020), given the feature map of the respective attention kernel. Random feature-based algorithms (Liu et al., 2022) are a family of algorithms, inspired by spectral analysis, which provides the feature map associated with a kernel. A random feature consists of three components: component function, weight matrix, and component weight (see Section 3). Due to their robustness and flexibility, random features give rise to ‘*random feature Transformers*’, the family of Transformers with linearized attention via random features (hereby abbreviated as RF), first with the work by Choromanski et al. (2021). This is seen to lead to three strands of research: (A) **Optimization of weight matrices** either by making the matrices have lower time or space complexity, or by incorporating beneficial properties like orthogonality to reduce variance (Yu et al., 2016; Choromanski et al., 2021; 2022; Reid et al., 2023); (B) **Enhancement of component functions** by either: (i) finding a function with a tighter variance or approximation bound, or (ii) be engineered to output with numerical stability and be bounded (Choromanski et al., 2021; Likhoshesterov et al., 2022; 2023); (C) Parameterization of weight matrices to perform kernel learning instead of kernel approximation (Chowdhury et al., 2022). (A) and (B) seek to improve approximation quality via approximation error and variance reduction. (C) dispenses with the approximation of the attention kernel (typically the softmax), and makes the kernel itself learnable, to attain a better

performance. However, the three have been explored separately, creating a situation where there are different ways to construct the random features. Additionally, past works only explore certain combinations of component functions and weight matrices, creating disjoint overlaps and gaps in the literature. Finally, the weight matrices are often those explicitly designed in the context of kernelized attention in the Transformer (Yu et al., 2016; Choromanski et al., 2021; 2022; Reid et al., 2023), ignoring those which are popular in the kernel methods literature (e.g., Liu et al. (2022)). These limitations necessitate a unified framework that allows the comparison of different combinations in search of the best-performing one.

We introduce *Spectraformer*, a unified framework of random feature-based attention in Transformer. It utilizes **spectral** analysis-inspired random features to construct linearized attention in **Transformers**. *Spectraformer* allows for the experimentation of various combinations of weight matrices and component functions. We will open-source our codebase upon acceptance. Through our benchmarking experiments on **three textual Long Range Arena (Tay et al., 2021) tasks** and **18 combinations** of weight matrices and component functions, *Spectraformer* enables novel combinations of component functions and weight matrices from the kernel literature to be experimented easily in the Transformer setting. This process makes possible the discovery of combinations that outperform existing random feature-based Transformers in terms of performance, training time, and peak memory consumption. Various novel Spectraformer combinations are more accurate than the previous SOTA models, SADERF-ORF (a.k.a., FAVOR# (Likhoshesterov et al., 2023)), OPRF-ORF (a.k.a., FAVOR++ (Likhoshesterov et al., 2022)), PosRF-ORF (a.k.a., FAVOR+ (Choromanski et al., 2021)), whilst significantly reducing training time and memory consumption over both. We hope that *Spectraformer* will enable the mechanism to benchmark more random feature-based Transformers for these and other tasks. Our contributions are:

1. Our novel framework, *Spectraformer*, which generalizes over all past works in linearized attention with kernel approximation and kernel learning.
2. We demonstrate how Spectraformer can discover novel SOTA models in the random feature Transformer family for the LRA benchmark and that these models are more performant than existing random feature Transformers.
3. We empirically identify the relationship between performant combinations and task characteristics, aiding the right kernel choice when tackling a task.
4. Our code is currently available for review, but they will be publicly released upon acceptance. We detail guide on how to incorporate future kernel works into Spectraformer. With this work, every future kernel work that is introduced can easily be incorporated into the transformer.

2 PRELIMINARIES

We first cover the theoretical foundations to this work. A list of commonly used acronyms is available in Table 1.

Table 1: Table of acronyms

Acronym	Description
OT (Vaswani et al., 2017)	Original Transformer
RF (Rahimi & Recht, 2007)	Random Feature method
RBF	Radial Basis Function kernel (also called the Gaussian kernel)
TrigRF (Rahimi & Recht, 2007)	Trigonometric RF (component function)
PosRF (Choromanski et al., 2021)	Positive RF (component function)
OPRF (Likhoshesterov et al., 2022)	Optimized positive RF (component function)
SADERF (Likhoshesterov et al., 2023)	Simplified Asymmetric Dense-Exponential RF (component function)
ORF (Yu et al., 2016)	Orthogonal RF (weight matrix)
SORF (Yu et al., 2016)	Structural ORF (weight matrix)
QMC (Avron et al., 2016)	Quasi-Monte Carlo (weight matrix)
MM (Shen et al., 2017) (Liu et al., 2022)	Moment Matching (weight matrix)
SGQ (Dao et al., 2017)	Sparse Grid Quadrature (weight matrix)
FastFoodL (Yang et al., 2015)	FastFood Learnable (weight matrix)

2.1 KERNELIZED ATTENTION

OT attention (Vaswani et al., 2017) can be defined using the unified attention model (Galassi et al., 2021) as follows. Given a collection of inputs and a learning objective, **attention** mechanism learns the attention matrix \mathbf{A} which captures the associative information between pairs of tokens in a sequence of length N , with each row \mathbf{A}_i being the representation of a token i with every other token $j \in [1, N]$. This is done by creating three learnable representations of each token i called **query** (\mathbf{q}_i), **key** (\mathbf{k}_i), and **value** (\mathbf{v}_i). The attention representation \mathbf{A}_i is calculated as a weighted sum of value \mathbf{v}_j given **attention weights** a_j . The general equation for \mathbf{A}_i is given as:

$$\mathbf{A}_i = \sum_{j=1}^N \frac{f(\mathbf{q}_i, \mathbf{k}_j)}{\sum_l f(\mathbf{q}_i, \mathbf{k}_l)} \mathbf{v}_j = \sum_{j=1}^N a_j \mathbf{v}_j; \quad a = g(e); \quad e = f(\mathbf{q}_i, \mathbf{k}_j) \quad (1)$$

Attention weights are generated using a **distribution function** g applied on an energy score e . The energy score is captured by applying a **compatibility function** on a query \mathbf{q}_i and a set of keys \mathbf{K} , in order to compute a score of a specific relationship between the pairs. The attention in the OT is called the ‘scaled dot-product’ which defines the compatibility function f and the distribution function g as follows:

$$f(\mathbf{q}_i, \mathbf{K}) = \frac{\mathbf{q}_i^T \mathbf{K}}{\sqrt{d_k}} \quad g(\mathbf{x}_j) = \frac{\exp(\mathbf{x}_j)}{\sum_l \exp(\mathbf{x}_l)} \quad (2)$$

We can then have the attention matrix (with $\exp(\cdot)$ being applied element-wise) as:

$$\mathbf{A}_i = \sum_{j=1}^N \frac{\exp(\mathbf{q}_i \mathbf{k}_j^T / \sqrt{d_k})}{\sum_l \exp(\mathbf{q}_i \mathbf{k}_l^T / \sqrt{d_k})} \mathbf{v}_j \quad \mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \quad (3)$$

with $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ being the matrix of queries, keys, and values respectively with each row being $\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i$.

We now discuss the intrinsic connection between kernel and attention. With the compatibility function f as a kernel $f(\mathbf{q}_i, \mathbf{k}_j) = \mathcal{K}(\mathbf{q}_i, \mathbf{k}_j)$, Equation 1 can be rewritten as:

$$\mathbf{A}_i = \sum_{j=1}^N \frac{\mathcal{K}(\mathbf{q}_i, \mathbf{k}_j)}{\sum_l \mathcal{K}(\mathbf{q}_i, \mathbf{k}_l)} \mathbf{v}_j \quad (4)$$

As shown in Tsai et al. (2019), this is indeed the Nadaraya-Watson Kernel Estimator, with $\mathbb{E}_{p(\mathbf{k}_j|\mathbf{q}_i)}[\mathbf{v}_j|X = \mathbf{k}_j] = \mathbf{A}_i$, $l_i(\mathbf{k}_j) = p(\mathbf{k}_j|\mathbf{q}_i)$, $Y_i = \mathbf{v}_j$:

$$\mathbb{E}_{l_i}[Y_i|X = \mathbf{u}] = \sum_{i=1}^N l_i(\mathbf{u}) Y_i; \quad l_i(\mathbf{u}) = \frac{\mathcal{K}(\frac{\mathbf{u}-\mathbf{q}_i}{h})}{\sum_{j=1}^n \mathcal{K}(\frac{\mathbf{u}-\mathbf{q}_j}{h})} \quad (5)$$

Specifically, the scaled dot-product attention in the OT in Equation 3 corresponds to the softmax kernel (see Equation 6).

$$\mathcal{K}_{\text{softmax}}(\mathbf{q}_i, \mathbf{k}_j) = \exp(\mathbf{q}_i^T \mathbf{k}_j / \sqrt{d_k}) \quad (6)$$

2.2 LINEARIZING KERNELIZED ATTENTION VIA RANDOM FEATURES

Since a kernel is the inner product of two vectors in some space with respect to some feature map ϕ , i.e., $\mathcal{K}(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x})\phi(\mathbf{y})^T$, then we can rewrite Equation 4 as:

$$\mathbf{A}_i = \frac{\sum_j \phi(\mathbf{q}_i)\phi(\mathbf{k}_j)^T \mathbf{v}_j}{\sum_l \phi(\mathbf{q}_i)\phi(\mathbf{k}_l)^T} = \frac{\phi(\mathbf{q}_i) \sum_j \phi(\mathbf{k}_j) \otimes \mathbf{v}_j}{\phi(\mathbf{q}_i) \sum_l \phi(\mathbf{k}_l)^T} \quad (7)$$

If we pre-compute $\sum_j \phi(\mathbf{k}_j) \otimes \mathbf{v}_j$ and $\sum_l \phi(\mathbf{k}_l)^T$, the entire term becomes $O(1)$ and we only need to compute Equation 7 N times for each i , resulting in an attention matrix calculated in linear time $O(N)$. We now refer to Equation 7 as **linearized attention** (Katharopoulos et al., 2020).

All that is left is to retrieve the feature map ϕ , the process of which is called **kernel approximation** (see Section 2.3), which corresponds to the kernels we specify. There are many popular kernels

162 from which we can choose, but typically, we want to approximate the softmax kernel in Equation 6.
 163 However, since the softmax relates to the $\mathcal{K}_{RBF}(\mathbf{x}, \mathbf{y}) = \exp(-\frac{\|\mathbf{x}-\mathbf{y}\|^2}{2})$ via the following equation:
 164

$$165 \mathcal{K}_{softmax}(\mathbf{x}, \mathbf{y}) = \exp(\frac{\|\mathbf{x}\|^2}{2})\mathcal{K}_{RBF}(\mathbf{x}, \mathbf{y}) \exp(\frac{\|\mathbf{y}\|^2}{2}) \quad (8)$$

166 Approximating the RBF is easier than the softmax directly in the method we introduce in Section 2.3,
 167 hence we approximate the softmax indirectly via the RBF in Equation 8.
 168

169 Oftentimes, however, we might not want to specify a kernel, since kernel choice is indeed a hyperparameter and the softmax is by no means essential or necessary as shown in Tsai et al. (2019).
 170 Then kernel learning would be a more suitable option, the detail of which is discussed in Section 2.4.
 171 Notwithstanding, we first introduce the main kernel approximation technique: random features.
 172
 173

174 2.3 RANDOM FEATURES FOR KERNEL APPROXIMATION

175 Of the kernel approximation family of methods, the most successful and appropriate for linearized
 176 attention in Transformer has been the random features approach (Rahimi & Recht, 2007). Estimating
 177 kernels using random features relies on a fundamental insight from harmonic analysis called Bochner’s
 178 theorem (Rudin, 1990) stated as follows:
 179

180 *A continuous shift-invariant kernel $\mathcal{K}(\mathbf{x}, \mathbf{y}) = \mathcal{K}(\mathbf{x} - \mathbf{y})$ on \mathbb{R}^d is positive definite if and only if $\mathcal{K}(\boldsymbol{\delta})$
 181 is the Fourier transform of a non-negative measure $p(\boldsymbol{\omega})$. If $\mathcal{K}(\boldsymbol{\delta})$ is properly scaled, that is $\mathcal{K}(0) = 1$
 182 then the measure $p(\boldsymbol{\omega})$ is a proper probability distribution.*
 183

$$184 \mathcal{K}(\boldsymbol{\delta}) = \int p(\boldsymbol{\omega}) \exp(i\boldsymbol{\omega}\boldsymbol{\delta})d\boldsymbol{\omega} \quad (9)$$

185 When approximating the integration in this equation via Monte Carlo sampling with s samples,
 186 $\boldsymbol{\omega} \sim p(\boldsymbol{\omega})$, we can obtain the feature map:
 187

$$188 \begin{aligned} \mathcal{K}(\boldsymbol{\delta}) &= \mathbb{E}[\phi_{\boldsymbol{\omega}}(\mathbf{x})^T \phi_{\boldsymbol{\omega}}(\mathbf{y})] \approx \langle \phi_{\boldsymbol{\omega}}(\mathbf{x}), \phi_{\boldsymbol{\omega}}(\mathbf{y}) \rangle \\ \phi_{\boldsymbol{\omega}}(\mathbf{x}) &= \frac{1}{\sqrt{s}} [\exp(-i\boldsymbol{\omega}_1^T \mathbf{x}), \dots, \exp(-i\boldsymbol{\omega}_s^T \mathbf{x})]^T \end{aligned} \quad (10)$$

189 2.4 RANDOM FEATURES FOR KERNEL LEARNING

190 Previous random feature techniques are robust in approximating kernels for specific learning problems.
 191 However, kernels may be chosen using heuristics and by convention from a popular subset. The
 192 kernel is very much a hyperparameter. This could pose a problem since there is no free lunch in
 193 learning as much as in kernel choice (Schölkopf & Smola, 2018). Tsai et al. (2019) has shown
 194 experimental results that this is also true in the context of attention in Transformer. Hence, instead
 195 of picking a kernel for the attention, we could learn the kernel via **kernel learning** an established
 196 kernel methods technique. Kernel learning has shown to be effective in general (Tompkins et al.,
 197 2019; Wilson & Adams, 2013; Oliva et al., 2016) and in the case of attention via Gaussian Mixture
 198 Model (GMM), learnable weight matrices (FastFood), deep generative model (DGM) (Chowdhury
 199 et al., 2022). Kernel learning is achieved by making the weight matrix \mathbf{W} learnable where $\boldsymbol{\omega}_i$ is the
 200 output of a function parameterizing $p(\boldsymbol{\omega})$.
 201

202 The kernelized attention via random features discussed here is a popular approach, however it is
 203 by no means the only one. A more detailed discussion is offered in related work in Section A.2.
 204 Specifically, Section A.2.1 provides alternative kernelized attention formulations and Section A.2.2
 205 provides alternative kernel approximations to random features.
 206

207 3 SPECTRAFORMER

208 *Spectraformer* is a unified random feature framework that allows for the combination of any weight
 209 matrix with any component function to be used as kernelized attention in Transformer. The general
 210 equation for random feature is shown in Equation 11, generalized from Equation 10, based on Liu
 211 et al. (2022); Choromanski et al. (2021).
 212

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

$$\begin{aligned} \phi_{\omega}(\mathbf{x}) &= \frac{1}{\sqrt{s}} [a_1 f_1(\omega_1, \mathbf{x}), \dots, a_s f_s(\omega_s, \mathbf{x}), \dots, \\ &\quad a_1 f_1(\omega_1, \mathbf{x}), \dots, a_s f_s(\omega_s, \mathbf{x})]^\top \\ \mathbf{W} &= [\omega_1, \dots, \omega_s]^\top \in \mathbb{R}^{s \times d} \end{aligned} \tag{11}$$

Spectraformer is shown in Figure 1. The attention of the OT (on the left-hand side) is replaced by the attention of the Spectraformer which pre-computes the Hadamard (element-wise) product of values v and transformed keys $\phi(k)$ scaled by $\phi(k)$, the product of which is multiplied with the transformed queries $\phi(q)$. The linear map ϕ is composed from weight matrices and component functions. The linear map has s components with a being the component weight (see Equation 11). The component weight a can either be fixed (typically at $a = 1$) or learnable as proposed by various kernel works (Liu et al., 2022). Due to the fact that the majority of random feature Transformers (e.g., Likhoshesterov et al. (2022)) have $a = 1$, we follow this setting for consistent comparison between combinations.

The figure and the equation together highlight the three strands of research as mentioned previously in Section 1: (A) Weight matrix approximator: constructing \mathbf{W} more effectively; (B) Weight matrix learner: parameterizing \mathbf{W} instead of sampling; (C) Component function: constructing f_j more effectively.

3.1 WEIGHT MATRICES

In Spectraformer, the weight matrix \mathbf{W} which is either an approximator (from the families of (e.g., ORF and SORF (Yu et al., 2016)), Unit-cube (e.g., QMC (Avron et al., 2016), MM (Shen et al., 2017)), Quadrature (e.g., SGQ (Dao et al., 2017)). or a learner (FastFood_L (Chowdhury et al., 2022)). We term the matrix \mathbf{W} as weight matrix instead of random matrix since although \mathbf{W} acts like a

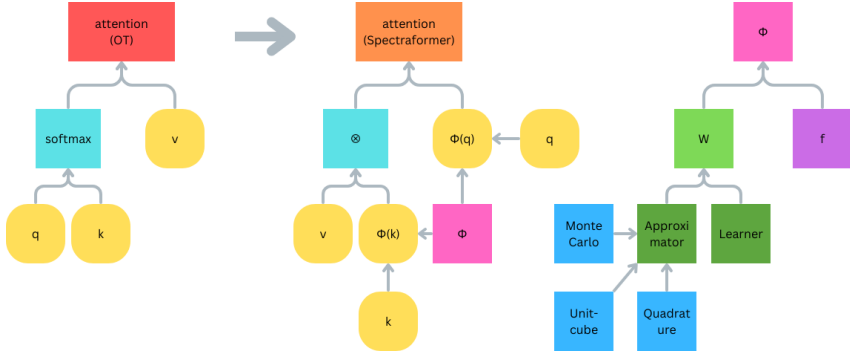


Figure 1: Spectraformer framework based on Equation 11

random matrix ($\omega_i \sim p(\cdot)$, see Equation 10), it is not guaranteed to be one. In the case of weight matrix approximator, they approximate a random matrix and in the case of weight matrix learner, they parameterize a weight matrix, thereby implicitly learning a distribution $p(\cdot)$ associated with such \mathbf{W} . The weight matrix is given in Definition 3.1.

Definition 3.1. Provided that a matrix can substitute for a random matrix in Equation 11, and the solution for this equation given such matrix and f as TrigRF being the solution for Equation 9, i.e., as an unbiased estimator of \mathcal{K} , and TrigRF being $f_1 = f_i = \exp(-i\omega^T \mathbf{x})$ (see Equation 10) then such a matrix is a weight matrix.

It follows from Definition 3.1 that for any weight matrix \mathbf{W} , $f = TrigRF$ is a valid component function. Spectraformer **enables weight matrix approximation** in terms of three families based on how they solve the intractable integral in Equation 9:

- **Monte Carlo sampling**-based weight matrix involves approximating the integral of Equation 9 using Monte Carlo sampling with the solution given by Equation 10 (Rahimi & Recht, 2007) with \mathbf{W} being the ‘Base’ random matrix of $p(\cdot)$, $\omega_i \sim p(\cdot)$. The random matrix \mathbf{W} can be constructed more efficiently (either reducing time or space complexity) with FastFood (Le et al., 2013). The random matrix \mathbf{W} can also be enforced to form geometrical couplings such as orthogonality to

reduce approximation error with ORF, SORF (Yu et al., 2016). Both these methods produce valid random matrices.

- **Unit-cube sampling**-based weight matrix transforms Equation 9 to an integral on the unit cube $[0, 1]^d$, then performs the approximation with uniform and independent points, thus reducing variance. The first approach is QMC (Avron et al., 2016), where $\mathbf{W} : \omega_i = \Phi^{-1}(t_i)$, Φ being the cumulative distribution function (CDF) associated with $p(\cdot)$ and t_i being a low discrepancy sequence. $\Phi^{-1}(t_i) \sim p(\cdot)$ whilst having a lower variance than direct sampling from $p(\cdot)$. MM (Shen et al., 2017; Liu et al., 2022) improves over QMC by replacing Φ^{-1} with a moment matching scheme $\tilde{\Phi}^{-1}$ on Φ^{-1} .
- **Quadrature**-based weight matrix uses quadrature rules with non-uniform deterministic weights. The main approach explored is SGQ (Dao et al., 2017) which uses one-dimensional Gaussian quadrature rule by assuming \mathcal{K} factorizes with respect to the dimensions. Smolyak rule is further implemented to alleviate the curse of dimensionality.

Liu et al. (2022) show that in various kernel benchmark datasets (among combination with TrigRF), QMC and ORF have consistently high performance. Due to the significant number of weight matrices in random features literature, we decide to cover only the most prominent ones. Appendix A.3 provides the technical details of the approximators discussed here and those that we left out.

Spectraformer also **allows to learn weight matrices** by learning $p(\cdot)$ via parameterizing \mathbf{W} . Yang et al. (2015) introduce several of these parametrization schemes, which are then adapted into the attention setting by Chowdhury et al. (2022). The best performing method is via making the constituent matrices of FastFood learnable, which we term FastFood_L. In general, a weight matrix learner is valid as long as it can implicitly model any distribution. In addition the TrigRF, PosRF has also been applied on FastFood_L. Technical details and other learners we have not covered are explored in Appendix A.4. We now explore existing works which have developed from Rahimi & Recht (2007) (see Equation 9, 10) and how they lead to the development of Spectraformer.

3.2 COMPONENT FUNCTIONS

Spectraformer incorporates component functions f such as PosRF, OPRF, SADERF. Component functions in Spectraformer combine each weight matrix row ω_i with the input \mathbf{x} . The base case of the component function is TrigRF. Component function is given in Definition 3.2

Definition 3.2. A function of \mathbf{x} with respect to ω is a valid component function if and only if $\mathcal{K} \approx \phi_\omega(\mathbf{x})^T \phi_\omega(\mathbf{y})$ given that $\omega_i \sim p(\cdot)$ in Equation 10.

It follows that all component functions have provided theoretical guarantees given that \mathbf{W} is a random matrix. In addition, PosRF, OPRF, and SADERF have also been proven theoretically to approximate \mathcal{K}_{RBF} given \mathbf{W} being the ORF.

While **TrigRF**(Rahimi & Recht, 2007) is the base component function and guarantees $\mathcal{K}(\delta)$ to be close to $\phi_{TrigRF}(\mathbf{x}) \cdot \phi_{TrigRF}(\mathbf{y})$ due to uniform convergence, it is typically evaluated using only the real component of $\exp(-i\omega_j^T \mathbf{x})$. This leads to cos and sin functions which can produce unstable behavior, especially when $K(\delta)$ is close to 0, where variance is infinite (Choromanski et al., 2021; Likhoshevstov et al., 2022). Therefore, TrigRF is avoided in Spectraformer.

- **PosRF** (Choromanski et al., 2021) is introduced in Spectraformer as a simple modification over TrigRF to counter this instability. PosRF guarantees positive function output and its variance approaching 0 as $K(\delta)$ approaches 0. Likhoshevstov et al. (2022) generalize TrigRF and PosRF to GERF.
- **OPRF** (Likhoshevstov et al., 2022) is one of the solution to the reduction of the variance equation of GERF. OPRF is both positive and bounded (whereas PosRF is positive and unbounded). Likhoshevstov et al. (2023) generalize GERF to dense matrix instead of scalar parameters leading to DERF.
- **SADERF** (Likhoshevstov et al., 2023) is one of the solutions to the reduction of the variance equation of DERF. Whilst DERF has many solutions, they all rely on operations which are incompatible with many deep learning libraries on GPU and TPU like SVD and eigen decompositions. SADERF is an extension of GERF with a tighter variance bound on DERF.

The alternatives for component functions allow Spectraformer to approximate the RBF with $p(\cdot)$ corresponding to the Gaussian. Likhoshesterov et al. (2023) and Likhoshesterov et al. (2022) show that SADERF-ORF is the SOTA combination followed by OPRF-ORF and PosRF-ORF. Technical details and other component function(s) not covered are explored in Appendix A.5.

Table 2: Summary of research gap in kernelized attention. Rows are the component functions, columns are the weight matrices. For each cell, the top item is the cited work for this combination for typical kernel tasks, while the bottom item is the cited work for Transformer-based tasks. \times indicates no prior work for the combination. Spectraformer explores all combinations listed here, excluding ‘Base’ due to exhaustive work as seen in the table.

		TrigRF	PosRF	OPRF	SADERF
Monte Carlo	Base	(Rahimi & Recht, 2007) (Choromanski et al., 2021)	(Choromanski et al., 2021)	(Likhoshesterov et al., 2022)	(Likhoshesterov et al., 2023)
	ORF	(Yu et al., 2016) (Choromanski et al., 2021)	(Choromanski et al., 2021)	(Likhoshesterov et al., 2022)	(Likhoshesterov et al., 2023)
	SORF	(Yu et al., 2016)	\times	\times	\times
Unit-cube	QMC	(Avron et al., 2016)	\times	\times	\times
		\times	\times	\times	\times
	MM	(Shen et al., 2017)	\times	\times	\times
Quadrature	SGQ	(Dao et al., 2017)	\times	\times	\times
		\times	\times	\times	\times
Learner	FastFood _L	(Yang et al., 2015) (Chowdhury et al., 2022)	\times (Chowdhury et al., 2022)	\times	\times

3.3 UTILITY

There has been an unsystematic attempt at combining component functions and weight matrices. On one hand, most component functions have had theoretical and experimental results in combining with Base and ORF. On the other hand, most weight matrices have had results in combining with TrigRF or PosRF. This means that a lot of potential combinations have not been studied previously. Table 2 shows the wide research gap across 14 combinations which have never been studied in either the kernel or Transformer setting. Spectraformer highlights this research gap. Excluding TrigRF due to its instability and Base due to its under-performance, we conduct experiments on 18 combinations.

Spectraformer is a generic random feature framework that allows for the combination of any weight matrix with any component function provided that the weight matrix satisfies Definition 3.1 and the component function satisfies Definition 3.2. We note that *Spectraformer* is an experimental rather than a theoretical framework. Due to the large number of weight matrices and component functions, it is infeasible for the time being to derive a general mathematical proof for the bounds and convergence of *Spectraformer*. Therefore, we seek to validate its feasibility and show its ability to discover novel and performant combination via experimental results.

3.4 ADDING NEW CODE TO SPECTRAFORMER

We now discuss how to add a new component function or weight matrix to the base code for further implementation. New component functions need to satisfy Definition 3.2 and new weight matrices need to satisfy Definition 3.1.

3.4.1 COMPONENT FUNCTION

1. Add the new component function `f` to `src/models/component_functions.py`, the arguments should include `data` (the input), and other optional parameters.
2. Import `f` and add a new entry to `FastAttention.comp_functions[f_name] = f` in `src/models/attention_performer.py` (line 176)

3.4.2 WEIGHT MATRIX

1. Add the new weight matrix `w` to `src/models/weight_matrix_approx.py` or `src/models/weight_matrix_learner.py`, the arguments should include `nb_rows` (number of rows), `nb_cols` (number of columns) and `device`.

2. Import `w` and add the `if` clause and `w` function call in `src/models/attention_performer.py` (line 208)

4 RESULTS

4.1 EXPERIMENT DETAILS

Table 3: Experimental results of Spectraformer variants on the LRA benchmark, over five seeds. We report mean accuracy (on the test set), mean time (training time (hour)), and mean memory (peak memory consumption (GB)). *L*: ListOps, *T*: Text, *R*: Retrieval, and μ for the average of the three. Entries are ranked by μ accuracy. The previous RF SOTAs (SADERF-ORF (Likhoshesterov et al., 2023) and OPRF-ORF (Likhoshesterov et al., 2022)) are in bold. We note that via Spectraformer, we find various variants outperforming these previous RF SOTAs in terms of accuracy, with different combinations offering different trade-off between accuracy, time, and memory. Entries are sorted by mean accuracy.

	Accuracy (%) \uparrow				Time (hour) \downarrow				Memory (GB) \downarrow			
	L	T	R	μ	L	T	R	μ	L	T	R	μ
OPRF-FastFoodL	37.55 (0.48)	64.41 (0.62)	77.70 (0.33)	59.89	1.07	2.07	2.12	1.75	0.86	1.72	1.68	1.42
OPRF-MM	38.08 (0.53)	60.40 (0.85)	81.09 (0.18)	59.86	0.68	1.26	1.24	1.06	1.36	2.71	2.56	2.21
PosRF-MM	37.06 (0.37)	61.87 (1.79)	80.58 (0.53)	59.84	0.56	1.05	1.06	0.89	1.17	2.31	2.10	1.86
OPRF-ORF	38.34 (0.22)	60.16 (0.79)	80.88 (0.17)	59.80	0.68	1.26	1.25	1.06	1.36	2.71	2.56	2.21
SADERF-QMC	37.37 (0.38)	61.14 (1.48)	80.84 (0.11)	59.78	0.68	1.25	1.29	1.07	1.44	2.86	2.69	2.33
PosRF-QMC	37.11 (0.09)	61.69 (0.96)	80.55 (0.13)	59.78	0.56	1.05	1.05	0.89	1.17	2.31	2.10	1.86
SADERF-MM	37.10 (0.22)	60.68 (1.88)	81.13 (0.17)	59.64	0.68	1.25	1.29	1.07	1.44	2.86	2.69	2.33
SADERF-ORF	37.10 (0.19)	60.39 (2.08)	81.05 (0.22)	59.51	0.68	1.25	1.28	1.07	1.44	2.86	2.69	2.33
OPRF-QMC	37.69 (0.62)	59.94 (0.59)	80.38 (0.49)	59.34	0.68	1.26	1.26	1.07	1.36	2.71	2.56	2.21
SADERF-SGQ	37.11 (0.21)	62.46 (0.54)	78.38 (0.25)	59.32	0.68	1.25	1.27	1.07	1.44	2.86	2.69	2.33
SADERF-FastFoodL	36.02 (1.38)	64.63 (0.18)	76.99 (0.61)	59.21	1.07	2.08	2.16	1.77	0.92	1.84	1.80	1.52
OPRF-SGQ	37.10 (0.23)	61.25 (0.54)	78.69 (0.54)	59.01	0.67	1.26	1.25	1.06	1.36	2.71	2.56	2.21
PosRF-ORF	34.35 (5.96)	60.30 (0.97)	80.45 (0.22)	58.37	0.56	1.05	1.06	0.89	1.17	2.31	2.10	1.86
PosRF-FastFoodL	33.46 (3.70)	64.65 (0.36)	76.95 (0.48)	58.35	1.02	1.98	2.03	1.68	0.79	1.57	1.53	1.30
SADERF-SORF	33.30 (0.98)	64.70 (0.36)	74.71 (1.90)	57.57	0.68	1.24	1.28	1.07	1.44	2.86	2.69	2.33
PosRF-SGQ	28.64 (7.54)	62.38 (0.53)	78.28 (0.20)	56.43	0.56	1.05	1.05	0.89	1.17	2.31	2.10	1.86
OPRF-SORF	27.91 (3.26)	64.76 (0.66)	75.92 (1.74)	56.20	0.67	1.26	1.24	1.06	1.36	2.71	2.56	2.21
PosRF-SORF	21.27 (6.65)	62.99 (0.40)	67.10 (1.11)	50.45	0.56	1.05	1.05	0.89	1.17	2.31	2.10	1.86

The LRA covers tasks of different sequence lengths, difficulty, and objective and it is designed to evaluate efficient Transformers, of which Spectraformer belongs to. Specifically, we evaluate the models on three textual LRA tasks: ListOps (Nangia & Bowman, 2018), Text (Byte-Level Text Classification, using the IMDb review dataset) (Howard & Ruder, 2018), and Retrieval (Byte-Level Document Retrieval, using the ACL Anthology Network dataset) (Radev et al., 2013). We discuss the task characteristics in greater details in Section 4.2. All our models are experimented on NVIDIA V100 and 12 CPU with 20GB of memory. Our codebase is based on and includes the adapted or original implementation from Chen et al. (2021); Chowdhury et al. (2022); Liu et al. (2022); Choromanski et al. (2021); Likhoshesterov et al. (2022, 2023); Wang et al. (2020); Xiong et al. (2021); Kitaev et al. (2020); Thomas et al. (2018); Zhou et al. (2021) (Apache License, Version 2.0).

We run our experiments on 5 different seeds due to computational limitation. The parameters are identical to Chen et al. (2021) and chosen to limit the parameter count to account for the marked training time of multiple models. All the parameters for Spectraformer variants are kept identical for fair comparison, with the only difference being their random feature component. Our code is implemented in Python 3.12 and Pytorch, we use the Transformer base from Chen et al. (2021). Our hyperparameters are available in Table 8. The model name convention is [component function]-[weight matrix] (e.g., OPRF-FastFood_L).

Our result in Table 3 shows that the top mean accuracy models, ranked in descending order, are: OPRF-FastFoodL, OPRF-MM, PosRF-MM, OPRF-ORF, SADERF-QMC. Our mean accuracy result indicates that novel Spectraformer combinations have outperformed the highlighted previous SOTA, i.e., OPRF-ORF and SADERF-ORF. However, we are aware that mean accuracy is insufficient in offering a comprehensive picture. Therefore, we first proceed with analyzing the statistics of the tasks before analyzing the significance of these results.

4.2 TASK CHARACTERISTICS

We first analyze the characteristics of the three textual classification LRA (Tay et al., 2021) tasks: ListOps (Nangia & Bowman, 2018), Text (Byte-Level Text Classification, using the IMDB review dataset) (Howard & Ruder, 2018), and Retrieval (Byte-Level Document Retrieval, using the ACL Anthology Network dataset) (Radev et al., 2013). We take into consideration three main characteristics:

- **Number of output classes.** ListOps is a task which inputs a sequence containing brackets and simple mathematical operators, with the output being an integer between 1 and 10, hence it has 10 classes. Text is a sentiment classification task on an IMDB review dataset with the class being either negative or positive, hence it has 2 classes. Retrieval inputs two sequences and the output is either False (meaning the sequences are not of the same document) or True (meaning the sequences are of the same document).
- **Maximum sequence size.** ListOps is 2K; Text is 4K; whilst Retrieval requires the processing of two sequences each of 4K length making the total maximum sequence size being 8K.
- **Required attention span.** This is the ‘mean distance between the query token and the attended tokens, scaled by attention weights’. The average values for required attention span calculated by Tay et al. (2021) is: ListOps being 0.8, Text being 0.3, and Retrieval being 1.3.

Table 4: Mean Pearson correlation between model performance statistics and each task characteristic across weight matrices and component functions. Model performance statistics include ‘acc’ (accuracy), ‘time’ (training time), ‘mem’ (peak memory consumption). Task characteristics are ‘size’ (maximum sequence size), ‘span’ (required attention span), and ‘classes’ (number of output classes). Correlation is calculated for each weight matrix or component function and the average value over weight matrix or component function respectively displayed. Left is weight matrix W , right is component function f .

W	acc	time	mem	f	acc	time	mem
size	0.93	0.74	0.64	size	0.93	0.56	0.60
span	0.37	0.02	-0.10	span	0.36	0.02	-0.10
classes	-0.91	-0.96	-0.94	classes	-0.91	-0.72	-0.89

We then calculate Pearson correlation of accuracy, training time, and memory consumption against these characteristics. Our result is shown in Table 4. The table indicates that: (1) contrary to intuition, longer sequence size gives higher accuracy but this is not because of increasing required attention span and they somewhat increase training time and peak memory consumption; (2) increasing the number of classes decreases accuracy, training time, and peak memory consumption.

4.3 ANALYSIS OF COMBINATIONS

Table 5: Mean statistics across weight matrices for each task. Task is denoted on the top left cell of each table. Mean statistics include ‘acc’ (accuracy (%)), ‘time’ (training time (hour)), and ‘mem’ (peak memory consumption (GB)).

ListOps	acc	time	mem	Text	acc	time	mem	Retrieval	acc	time	mem
SORF	27.49	0.64	1.32	ORF	60.28	1.19	2.63	SORF	72.58	1.19	2.45
SGQ	34.28	0.64	1.32	QMC	60.92	1.19	2.63	FastFoodL	77.22	2.10	1.67
FastFoodL	35.68	1.05	0.86	MM	60.98	1.19	2.63	SGQ	78.45	1.19	2.45
ORF	36.59	0.64	1.32	SGQ	62.03	1.19	2.63	QMC	80.59	1.20	2.45
QMC	37.39	0.64	1.32	SORF	64.15	1.18	2.63	ORF	80.79	1.20	2.45
MM	37.41	0.64	1.32	FastFoodL	64.56	2.04	1.71	MM	80.93	1.20	2.45

We now explore the suitability of weight matrices and component functions depending on task characteristics. Table 5 shows the mean statistics across weight matrices for each task. We notice that weight matrices form two different groups which align with the required attention span differences. ListOps and Retrieval are tasks with high required attention span, ListOps being 0.8 and Retrieval being 1.3. In both these tasks, we observe that ORF, QMC, and MM perform comparatively against each other and a lot better against SORF, SGQ, and FastFoodL. However, the Text task, which has

a much lower required attention span of 0.3, has much higher performance from the SORF-SGQ-FastFoodL group, whilst ORF, QMC, and MM perform much poorly on this task. Based on this observation, we can conclude that ORF, QMC, and MM are suited for high required attention span; whilst SORF, SGQ, and FastFoodL are suited for low required attention span. Analyzing mean statistics per task across component functions do not show any task-specific trend (refer to Table 9 for more information).

Table 6: Summary statistics of component functions f (top) and weight matrices W (bottom) in mean accuracy (on the test set), mean time (training time (hour)), and mean memory (peak memory consumption (GB)). f and W entries are sorted by mean accuracy respectively.

		Accuracy (%) \uparrow	Time (hour) \downarrow	Memory (GB) \downarrow
f	SADERF	59.17	1.19	2.19
f	OPRF	59.02	1.18	2.08
f	PosRF	57.20	1.02	1.77
W	MM	59.78	1.01	2.13
W	QMC	59.63	1.01	2.13
W	ORF	59.22	1.01	2.13
W	FastFoodL	59.15	1.73	1.41
W	SGQ	58.25	1.01	2.13
W	SORF	54.74	1.00	2.13

In analyzing efficiency, we follow the definition of efficient Transformers in Tay et al. (2022) in terms of memory and computation, the latter interpreted as computational cost. The overall trend of accuracy, time, and memory of component functions is that they increase linearly from PosRF, OPRF, to SADERF (see Table 6). Hence, when limited computational resource is a concern, PosRF is a better option. Training time and peak memory consumption relies more heavily on component functions than weight matrices, with the exception of FastFoodL. FastFoodL takes significantly less memory consumption at the expense of training time. The accuracy trend for weight matrices is as described above. Table 6 also shows that whilst FastFoodL generally does not work as well due to its low required attention span, its high performance in the task of Text which is low in required attention span and the anomalously performant OPRF-FastFoodL leads to a data skew in the final ranking in Table 3.

Comparing between combinations, we find that for ListOps, OPRF-ORF is the most accurate (38.34%), followed by OPRF-MM (38.08%) and OPRF-QMC (37.69%). OPRF-ORF, being the previous SOTA, is undoubtedly the best choice here since other models maintain the same training time and peak memory consumption with lower accuracy. For Text, the most accurate model is OPRF-SORF (64.76%), followed by SADERF-SORF (64.70%) and PosRF-FastFoodL (64.65%). These all outperform the previous SOTAs (OPRF-ORF and SADERF-ORF) by 4%+ whilst maintaining a diversity of training time and peak memory consumption options. Both OPRF-SORF and SADERF-SORF maintain the same level of training time and peak memory consumption as the SOTAs, PosRF-FastFoodL offers much less peak memory consumption at the expense of longer training time. For Retrieval, the most accurate model is SADERF-MM (81.13%), followed by OPRF-MM (81.09%) and SADERF-ORF (81.05%). Whilst these models only perform slightly better than the SOTA SADERF-ORF, OPRF-MM offers less peak memory consumption compared to SADERF-ORF.

5 CONCLUSION

We present Spectraformer, a framework for approximating and learning the attention kernel in the Transformer. Our paper generalizes past works and presents empirical findings on different component function and weight matrix combinations. We experiment with 18 combinations, of which 14 are novel. Novel combinations perform competitively against, if not more performant than the previous SOTA in the family for LRA textual tasks. We show empirically that the choice of performant weight matrix depends on the required attention span of the task. Our results and the extensibility of our work shows the viability of Spectraformer as a unifying random feature framework of Transformer.

REFERENCES

- 540
541
542 Haim Avron, Vikas Sindhwani, Jiyan Yang, and Michael W. Mahoney. Quasi-monte carlo feature
543 maps for shift-invariant kernels. *Journal of Machine Learning Research*, 17(120):1–38, 2016.
544 URL <http://jmlr.org/papers/v17/14-538.html>.
- 545 Francis R. Bach, Gert R. G. Lanckriet, and Michael I. Jordan. Multiple kernel learning, conic duality,
546 and the smo algorithm. In *Proceedings of the Twenty-First International Conference on Machine*
547 *Learning*, ICML ’04, pp. 6, New York, NY, USA, 2004. Association for Computing Machinery.
548 ISBN 1581138385. doi: 10.1145/1015330.1015424. URL [https://doi.org/10.1145/](https://doi.org/10.1145/1015330.1015424)
549 [1015330.1015424](https://doi.org/10.1145/1015330.1015424).
- 550 A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *2005 IEEE*
551 *Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 2,
552 pp. 60–65 vol. 2, 2005. doi: 10.1109/CVPR.2005.38.
- 553 Brian Bullins, Cyril Zhang, and Yi Zhang. Not-so-random features. In *International Confer-*
554 *ence on Learning Representations*, 2018. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=Hk8XMWgRb)
555 [Hk8XMWgRb](https://openreview.net/forum?id=Hk8XMWgRb).
- 556 Yifan Chen, Qi Zeng, Heng Ji, and Yun Yang. Skyformer: Remodel self-attention with gaussian kernel
557 and nyström method. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.),
558 *Advances in Neural Information Processing Systems*, 2021. URL [https://openreview.](https://openreview.net/forum?id=pZCYG7gjkKz)
559 [net/forum?id=pZCYG7gjkKz](https://openreview.net/forum?id=pZCYG7gjkKz).
- 560 Krzysztof Choromanski and Vikas Sindhwani. Recycling randomness with structure for sublinear
561 time kernel expansions. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of*
562 *The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine*
563 *Learning Research*, pp. 2502–2510, New York, New York, USA, 20–22 Jun 2016. PMLR. URL
564 <https://proceedings.mlr.press/v48/choromanski16.html>.
- 565 Krzysztof M Choromanski, Mark Rowland, and Adrian Weller. The unreasonable effective-
566 ness of structured random orthogonal embeddings. In I. Guyon, U. Von Luxburg,
567 S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Ad-*
568 *vances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,
569 2017. URL [https://proceedings.neurips.cc/paper_files/paper/2017/](https://proceedings.neurips.cc/paper_files/paper/2017/file/bf8229696f7a3bb4700cfddef19fa23f-Paper.pdf)
570 [file/bf8229696f7a3bb4700cfddef19fa23f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/bf8229696f7a3bb4700cfddef19fa23f-Paper.pdf).
- 571 Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea
572 Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser,
573 David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. Rethinking attention with
574 performers. In *International Conference on Learning Representations*, 2021. URL [https://](https://openreview.net/forum?id=Ua6zuk0WRH)
575 openreview.net/forum?id=Ua6zuk0WRH.
- 576 Krzysztof Marcin Choromanski, Han Lin, Haoxian Chen, Arijit Sehanobish, Yuanzhe Ma, Deepali
577 Jain, Jake Varley, Andy Zeng, Michael S Ryoo, Valerii Likhoshesterov, Dmitry Kalashnikov, Vikas
578 Sindhwani, and Adrian Weller. Hybrid random features. In *International Conference on Learning*
579 *Representations*, 2022. URL <https://openreview.net/forum?id=EMigfE6ZeS>.
- 580 Sankalan Pal Chowdhury, Adamos Solomou, Kumar Avinava Dubey, and Mrinmaya Sachan. Learning
581 the transformer kernel. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
582 URL <https://openreview.net/forum?id=tLIBAEYjcv>.
- 583 Tri Dao, Christopher M De Sa, and Christopher Ré. Gaussian quadrature for kernel features. In
584 I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Gar-
585 nett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Asso-
586 ciates, Inc., 2017. URL [https://proceedings.neurips.cc/paper_files/paper/](https://proceedings.neurips.cc/paper_files/paper/2017/file/62f91ce9b820a491ee78c108636db089-Paper.pdf)
587 [2017/file/62f91ce9b820a491ee78c108636db089-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/62f91ce9b820a491ee78c108636db089-Paper.pdf).
- 588 Chang Feng, Qinghua Hu, and Shizhong Liao. Random feature mapping with signed circulant
589 matrix projection. In *Proceedings of the 24th International Conference on Artificial Intelligence*,
590 *IJCAI’15*, pp. 3490–3496. AAAI Press, 2015. ISBN 9781577357384.
- 591
592
593

- 594 Andrea Galassi, Marco Lippi, and Paolo Torrioni. Attention in natural language processing. *IEEE*
595 *Transactions on Neural Networks and Learning Systems*, 32(10):4291–4308, 2021. doi: 10.1109/
596 TNNLS.2020.3019893.
- 597 Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang,
598 An Xiao, Chunjing Xu, Yixing Xu, Zhaohui Yang, Yiman Zhang, and Dacheng Tao. A survey
599 on vision transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):
600 87–110, 2023. doi: 10.1109/TPAMI.2022.3152247.
- 601 Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. Kernel methods in machine learning.
602 *The Annals of Statistics*, 36(3):1171 – 1220, 2008. doi: 10.1214/009053607000000677. URL
603 <https://doi.org/10.1214/009053607000000677>.
- 604 Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification.
605 In Iryna Gurevych and Yusuke Miyao (eds.), *Proceedings of the 56th Annual Meeting of the*
606 *Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 328–339, Melbourne,
607 Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1031.
608 URL <https://aclanthology.org/P18-1031>.
- 609 Cho-Jui Hsieh, Si Si, and Inderjit Dhillon. A divide-and-conquer solver for kernel support vector
610 machines. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st International Conference*
611 *on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 566–574,
612 Beijing, China, 22–24 Jun 2014. PMLR. URL [https://proceedings.mlr.press/v32/
613 hsieha14.html](https://proceedings.mlr.press/v32/hsieha14.html).
- 614 Pranav Jeevan and Amit Sethi. Resource-efficient hybrid x-formers for vision. In *2022 IEEE/CVF*
615 *Winter Conference on Applications of Computer Vision (WACV)*, pp. 3555–3563, 2022. doi:
616 10.1109/WACV51458.2022.00361.
- 617 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
618 RNNs: Fast autoregressive transformers with linear attention. In Hal Daumé III and Aarti Singh
619 (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of
620 *Proceedings of Machine Learning Research*, pp. 5156–5165. PMLR, 13–18 Jul 2020. URL
621 <https://proceedings.mlr.press/v119/katharopoulos20a.html>.
- 622 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.
623 URL <https://api.semanticscholar.org/CorpusID:216078090>.
- 624 Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In
625 *International Conference on Learning Representations*, 2020. URL [https://openreview.
626 net/forum?id=rkgNKkHtvB](https://openreview.net/forum?id=rkgNKkHtvB).
- 627 Quoc Le, Tamas Sarlos, and Alexander Smola. Fastfood - computing hilbert space expansions
628 in loglinear time. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th*
629 *International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning*
630 *Research*, pp. 244–252, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL [https://
631 proceedings.mlr.press/v28/le13.html](https://proceedings.mlr.press/v28/le13.html).
- 632 Valerii Likhoshesterov, Krzysztof M Choromanski, Kumar Avinava Dubey, Frederick Liu, Tamas
633 Sarlos, and Adrian Weller. Chefs'random tables: Non-trigonometric random features. In
634 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances*
635 *in Neural Information Processing Systems*, volume 35, pp. 34559–34573. Curran Asso-
636 ciates, Inc., 2022. URL [https://proceedings.neurips.cc/paper_files/paper/
637 2022/file/df2d62b96a4003203450cf89cd338bb7-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/df2d62b96a4003203450cf89cd338bb7-Paper-Conference.pdf).
- 638 Valerii Likhoshesterov, Krzysztof Marcin Choromanski, Kumar Avinava Dubey, Frederick Liu, Tamas
639 Sarlos, and Adrian Weller. Dense-exponential random features: Sharp positive estimators of the
640 gaussian kernel. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
641 URL <https://openreview.net/forum?id=S0xrBMFihS>.
- 642 Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. A survey of transformers. *AI Open*, 3:
643 111–132, 2022. ISSN 2666-6510. doi: <https://doi.org/10.1016/j.aiopen.2022.10.001>. URL [https://
644 www.sciencedirect.com/science/article/pii/S2666651022000146](https://www.sciencedirect.com/science/article/pii/S2666651022000146).

- 648 Tal Linzen, Grzegorz Chrupała, and Afra Alishahi (eds.). *GLUE: A Multi-Task Benchmark and*
649 *Analysis Platform for Natural Language Understanding*, Brussels, Belgium, November 2018.
650 Association for Computational Linguistics. doi: 10.18653/v1/W18-5446. URL [https://](https://aclanthology.org/W18-5446)
651 aclanthology.org/W18-5446.
- 652 Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan A. K. Suykens. Random features for kernel
653 approximation: A survey on algorithms, theory, and beyond. *IEEE Transactions on Pattern*
654 *Analysis and Machine Intelligence*, 44(10):7128–7148, 2022. doi: 10.1109/TPAMI.2021.3097011.
655
- 656 Yueming Lyu. Spherical structured feature maps for kernel approximation. In Doina Precup and
657 Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*,
658 volume 70 of *Proceedings of Machine Learning Research*, pp. 2256–2264. PMLR, 06–11 Aug
659 2017. URL <https://proceedings.mlr.press/v70/lyu17a.html>.
- 660 Yueming Lyu, Yuan Yuan, and Ivor Tsang. Subgroup-based rank-1 lattice quasi-monte
661 carlo. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Ad-*
662 *vances in Neural Information Processing Systems*, volume 33, pp. 6269–6280. Curran Asso-
663 ciates, Inc., 2020. URL [https://proceedings.neurips.cc/paper_files/paper/](https://proceedings.neurips.cc/paper_files/paper/2020/file/456048afb7253926e1fbb7486e699180-Paper.pdf)
664 [2020/file/456048afb7253926e1fbb7486e699180-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/456048afb7253926e1fbb7486e699180-Paper.pdf).
- 665 Marina Munkhoeva, Yermek Kapushev, Evgeny Burnaev, and Ivan Oseledets. Quadrature-based
666 features for kernel approximation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-
667 Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31.
668 Curran Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper_](https://proceedings.neurips.cc/paper_files/paper/2018/file/6e923226e43cd6fac7cfe1e13ad000ac-Paper.pdf)
669 [files/paper/2018/file/6e923226e43cd6fac7cfe1e13ad000ac-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/6e923226e43cd6fac7cfe1e13ad000ac-Paper.pdf).
- 670 Nikita Nangia and Samuel Bowman. ListOps: A diagnostic dataset for latent tree learning. In Sil-
671 vio Ricardo Cordeiro, Shereen Oraby, Umashanthi Pavalanathan, and Kyeongmin Rim (eds.),
672 *Proceedings of the 2018 Conference of the North American Chapter of the Association for*
673 *Computational Linguistics: Student Research Workshop*, pp. 92–99, New Orleans, Louisiana,
674 USA, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-4013. URL
675 <https://aclanthology.org/N18-4013>.
676
- 677 Tan Nguyen, Minh Pham, Tam Nguyen, Khai Nguyen, Stanley Osher, and Nhat Ho. Fouri-
678 erformer: Transformer meets generalized fourier integral theorem. In S. Koyejo, S. Mo-
679 hamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural In-*
680 *formation Processing Systems*, volume 35, pp. 29319–29335. Curran Associates, Inc.,
681 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/](https://proceedings.neurips.cc/paper_files/paper/2022/file/bc968adbdf4a2551649d464b83f264a-Paper-Conference.pdf)
682 [file/bc968adbdf4a2551649d464b83f264a-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/bc968adbdf4a2551649d464b83f264a-Paper-Conference.pdf).
- 683 Junier B. Oliva, Avinava Dubey, Andrew G. Wilson, Barnabas Poczos, Jeff Schneider, and Eric P.
684 Xing. Bayesian nonparametric kernel-learning. In Arthur Gretton and Christian C. Robert (eds.),
685 *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51
686 of *Proceedings of Machine Learning Research*, pp. 1078–1086, Cadiz, Spain, 09–11 May 2016.
687 PMLR. URL <https://proceedings.mlr.press/v51/oliva16.html>.
688
- 689 Dragomir R. Radev, Pradeep Muthukrishnan, Vahed Qazvinian, and Amjad Abu-Jbara. The acl
690 anthology network corpus. *Language Resources and Evaluation*, 47(4):919–944, 2013. ISSN
691 1574020X, 15728412. URL <http://www.jstor.org/stable/42636383>.
- 692 Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines.
693 In J. Platt, D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural In-*
694 *formation Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL
695 [https://proceedings.neurips.cc/paper_files/paper/2007/file/](https://proceedings.neurips.cc/paper_files/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf)
696 [013a006f03dbc5392effeb8f18fda755-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf).
- 697 Isaac Reid, Krzysztof Marcin Choromanski, Valerii Likhoshesterov, and Adrian Weller. Simplex
698 random features. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan
699 Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine*
700 *Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28864–28888. PMLR,
701 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/reid23a.html>.

- 702 Walter Rudin. *The Basic Theorems of Fourier Analysis*, chapter 1, pp. 1–34. John Wiley & Sons, Ltd,
703 1990. ISBN 9781118165621. doi: <https://doi.org/10.1002/9781118165621.ch1>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118165621.ch1>.
- 704
705 Bernhard Schölkopf and Alexander J Smola. *Kernels*. In *Learning with Kernels: Support Vector Ma-*
706 *chines, Regularization, Optimization, and Beyond*. The MIT Press, 06 2018. ISBN 9780262256933.
707 doi: 10.7551/mitpress/4175.003.0005. URL [https://doi.org/10.7551/mitpress/](https://doi.org/10.7551/mitpress/4175.003.0005)
708 [4175.003.0005](https://doi.org/10.7551/mitpress/4175.003.0005).
- 709
710 Weiwei Shen, Zhihui Yang, and Jun Wang. Random features for shift-invariant kernels with moment
711 matching. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*,
712 pp. 2520–2526. AAAI Press, 2017.
- 713
714 Alex J. Smola and Bernhard Schölkopf. Sparse greedy matrix approximation for machine learning. In
715 *Proceedings of the Seventeenth International Conference on Machine Learning, ICML ’00*, pp.
716 911–918, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 1558607072.
- 717
718 Kyungwoo Song, Yohan Jung, Dongjun Kim, and Il-Chul Moon. Implicit kernel attention. *Pro-*
719 *ceedings of the AAAI Conference on Artificial Intelligence*, 35(11):9713–9721, May 2021.
720 doi: 10.1609/aaai.v35i11.17168. URL [https://ojs.aaai.org/index.php/AAAI/](https://ojs.aaai.org/index.php/AAAI/article/view/17168)
[article/view/17168](https://ojs.aaai.org/index.php/AAAI/article/view/17168).
- 721
722 Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao,
723 Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena : A benchmark for efficient
724 transformers. In *International Conference on Learning Representations, 2021*. URL <https://openreview.net/forum?id=qVyeW-grC2k>.
- 725
726 Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey.
727 *ACM Comput. Surv.*, 55(6), dec 2022. ISSN 0360-0300. doi: 10.1145/3530811. URL <https://doi.org/10.1145/3530811>.
- 728
729 Anna Thomas, Albert Gu, Tri Dao, Atri Rudra, and Christopher Ré. Learning compressed transforms
730 with low displacement rank. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi,
731 and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Cur-
732 ran Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper_files/](https://proceedings.neurips.cc/paper_files/paper/2018/file/8e621619d71d0ae5ef4e631ad586334f-Paper.pdf)
733 [paper/2018/file/8e621619d71d0ae5ef4e631ad586334f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/8e621619d71d0ae5ef4e631ad586334f-Paper.pdf).
- 734
735 Anthony Tompkins, Ransalu Senanayake, Philippe Morere, and Fabio Ramos. Black box quantiles
736 for kernel learning. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the*
737 *Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of
738 *Proceedings of Machine Learning Research*, pp. 1427–1437. PMLR, 16–18 Apr 2019. URL
<https://proceedings.mlr.press/v89/tompkins19a.html>.
- 739
740 Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhut-
741 dinov. Transformer dissection: An unified understanding for transformer’s attention via the lens of
742 kernel. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019*
743 *Conference on Empirical Methods in Natural Language Processing and the 9th International Joint*
744 *Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4344–4353, Hong Kong,
745 China, 11 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1443. URL
<https://aclanthology.org/D19-1443>.
- 746
747 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
748 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von
749 Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Ad-*
750 *vances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,
751 2017. URL [https://proceedings.neurips.cc/paper_files/paper/2017/](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf)
[file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf).
- 752
753 Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention
754 with linear complexity, 2020.
- 755
756 Tinghua Wang, Dongyan Zhao, and Shengfeng Tian. An overview of kernel alignment and its
757 applications. *Artificial Intelligence Review*, 43:179–192, 2015.

- 756 Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In
757 *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7794–7803, 2018.
758 doi: 10.1109/CVPR.2018.00813.
- 759 Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra
760 Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins,
761 Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks,
762 William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. Ethical and social risks of
763 harm from language models, 2021.
- 764 Christopher Williams and Matthias Seeger. Using the nystrom method to speed up kernel machines.
765 In T. Leen, T. Dietterich, and V. Tresp (eds.), *Advances in Neural Information Processing Sys-*
766 *tems*, volume 13. MIT Press, 2000. URL [https://proceedings.neurips.cc/paper_](https://proceedings.neurips.cc/paper_files/paper/2000/file/19de10adbaa1b2ee13f77f679fa1483a-Paper.pdf)
767 [files/paper/2000/file/19de10adbaa1b2ee13f77f679fa1483a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2000/file/19de10adbaa1b2ee13f77f679fa1483a-Paper.pdf).
- 769 Andrew Wilson and Ryan Adams. Gaussian process kernels for pattern discovery and extrapolation.
770 In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference*
771 *on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1067–
772 1075, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL [https://proceedings.mlr.](https://proceedings.mlr.press/v28/wilson13.html)
773 [press/v28/wilson13.html](https://proceedings.mlr.press/v28/wilson13.html).
- 774 Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and
775 Vikas Singh. Nystromformer: A nystrom-based algorithm for approximating self-attention. *Pro-*
776 *ceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14138–14148, May 2021.
777 doi: 10.1609/aaai.v35i16.17664. URL [https://ojs.aaai.org/index.php/AAAI/](https://ojs.aaai.org/index.php/AAAI/article/view/17664)
778 [article/view/17664](https://ojs.aaai.org/index.php/AAAI/article/view/17664).
- 779 Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. A temporal kernel
780 approach for deep learning with continuous-time information, 2021.
- 782 Zichao Yang, Andrew Wilson, Alex Smola, and Le Song. A la Carte – Learning Fast Kernels.
783 In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Proceedings of the Eighteenth International*
784 *Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning*
785 *Research*, pp. 1098–1106, San Diego, California, USA, 09–12 May 2015. PMLR. URL <https://proceedings.mlr.press/v38/yang15b.html>.
- 787 Junwen Yao, N. Benjamin Erichson, and Miles E. Lopes. Error estimation for random fourier
788 features. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings*
789 *of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of
790 *Proceedings of Machine Learning Research*, pp. 2348–2364. PMLR, 25–27 Apr 2023. URL
791 <https://proceedings.mlr.press/v206/yao23a.html>.
- 793 Felix X. Yu, Sanjiv Kumar, Henry Rowley, and Shih-Fu Chang. Compact nonlinear maps and
794 circulant extensions, 2015.
- 795 Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-Rice,
796 and Sanjiv Kumar. Orthogonal random features. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon,
797 and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Cur-
798 ran Associates, Inc., 2016. URL [https://proceedings.neurips.cc/paper_files/](https://proceedings.neurips.cc/paper_files/paper/2016/file/53adaf494dc89ef7196d73636eb2451b-Paper.pdf)
799 [paper/2016/file/53adaf494dc89ef7196d73636eb2451b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/53adaf494dc89ef7196d73636eb2451b-Paper.pdf).
- 800 Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression.
801 In Shai Shalev-Shwartz and Ingo Steinwart (eds.), *Proceedings of the 26th Annual Conference*
802 *on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pp. 592–617,
803 Princeton, NJ, USA, 12–14 Jun 2013. PMLR. URL [https://proceedings.mlr.press/](https://proceedings.mlr.press/v30/Zhang13.html)
804 [v30/Zhang13.html](https://proceedings.mlr.press/v30/Zhang13.html).
- 805 Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang.
806 Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proceedings*
807 *of the AAAI Conference on Artificial Intelligence*, 35(12):11106–11115, May 2021. doi: 10.
808 1609/aaai.v35i12.17325. URL [https://ojs.aaai.org/index.php/AAAI/article/](https://ojs.aaai.org/index.php/AAAI/article/view/17325)
809 [view/17325](https://ojs.aaai.org/index.php/AAAI/article/view/17325).

810 A APPENDIX

811 A.1 FUTURE WORK, LIMITATIONS, AND BROADER IMPACT

812 Spectraformer provides an experimental framework for random features in the Transformer. However,
813 we have not provided formal mathematical proofs for the bounds the these combinations due to the
814 number of combinations discussed. We hope Spectraformer inspires future work to generalizes the
815 proof in such a framework.

816 Spectraformer has covered a significant number of combinations, however, we have not covered all
817 weight matrices and component functions in the literature. Spectraformer can serve as a foundation
818 for further random feature Transformer research using the techiques not yet experimented with but
819 covered in Appendices A.3, A.4, A.5.

820 Due to the significant number of experimented models, we are unable to perform hyper-parameter
821 fine-tuning. Future work can benefit from extending the scope of the current research to other tasks,
822 i.e., other non-textual LRA tasks (image classification, path-finder, and path-x) and short-context
823 datasets/ tasks such as GLUE (Linzen et al., 2018). Transformer-based architecture can also benefit
824 from Spectraformer. Past random feature Transformer works (Likhoshesterov et al., 2022; 2023) have
825 experimented with conformer and vision Transformer with ViT. Jeevan & Sethi (2022) provides
826 a good foundation to extend our work. Spectraformer, although currently experimented in the
827 Transformer setting, can also extend to non-parametric kernel classification. It could also benefit
828 from ablation studies like the comparison variance in Chowdhury et al. (2022) or in Yao et al. (2023)

829 Given the complex nature of language models in the ethical, energy, and societal domains, our work
830 should be used responsibly (Weidinger et al., 2021).s

831 A.2 RELATED WORK

832 Spectraformer is based on several theoretical assumptions and choice of mathematical framework
833 and approach. This section discusses alternatives to the choices of Spectraformer. Specifically, Ap-
834 pendix A.2.1 discusses alternative kernelized formulations to Section 2.1; Appendix A.2.2 discusses
835 alternative kernel approximation techniques to random features in Section A.2.2. There are also some
836 component functions and weight matrices which are not covered in the experiments. We introduce
837 them in the related work section in Appendices A.3, A.4, A.5.

838 A.2.1 ALTERNATIVE KERNELIZED ATTENTIONS

839 In addition to the kernelized formulation of attention as shown in Section 2.1, based on our literature
840 survey, we have identified the following three alternative kernel-based formulations of attention.

841 **Nadaraya-Watson kernel estimator in integration form (Nguyen et al., 2022):** Given the Gaussian
842 kernel, where the probability of a single variable and the joint probability of two variables are
843 estimated using kernel density estimation with the Gaussian kernel, then we have Equation 12,
844 leading to the formulation of FourierFormer.

$$\begin{aligned}
 \phi(\boldsymbol{\delta}) &= \exp(-\|\boldsymbol{\delta}\|^2/2\sigma^2) \\
 \hat{r}_n(\mathbf{k}) &= \int \frac{vp(\mathbf{v}, \mathbf{k})}{p(\mathbf{k})} dv = \frac{\sum_{j=1}^N \hat{v}_j \phi(\mathbf{k} - \mathbf{k}_j)}{\sum_{j=1}^N \phi(\mathbf{k} - \mathbf{k}_j)} \\
 \hat{r}_n(\mathbf{q}_i) &= \sum_{j=1}^N \text{softmax}(\mathbf{q}_i^T \mathbf{k}_j / \sigma^2) \hat{v}_j
 \end{aligned} \tag{12}$$

845 **Gaussian decomposition (Song et al., 2021):** the attention mechanism can also be decomposed into
846 the Gaussian kernel directly from Equation 4 without appealing to the use of kernel estimator (Song
847 et al., 2021). Equation 13 leads to Implicit Kernel Attention (IKA).

$$A_i = \sum_{j=1}^N \frac{1}{Z_l(\mathbf{q}_i, \mathbf{K})} \exp\left(-\frac{\|\mathbf{q}_i - \mathbf{k}_j\|_2^2}{2\sqrt{d_k}}\right) \exp\left(\frac{\|\mathbf{q}_i\|_{p=2}^2 + \|\mathbf{k}_j\|_{p=2}^2}{2\sqrt{d_k}}\right) \mathbf{v}_j \tag{13}$$

848 **Non-local operation (Wang et al., 2018)** is popular decomposition of attention in computer vision.
849 This is the basis for Vision Transformer and related work. Non-local operation itself a generalisation

of the non-local means used in image denoising (Buades et al., 2005). Non-local operation is essentially an operation that operates on the entire input range, instead of a specific window which is characteristic of convolutional operations (see Equation 14). Indeed, the attention vector output \mathbf{y} can be defined using f which is the softmax (see Equation 15) in the non-local operation.

$$\mathbf{y}_i = \frac{1}{C(\mathbf{x})} \sum_{\mathbf{v}_j} f(\mathbf{x}_i, \mathbf{x}_j) g(\mathbf{x}_j) \tag{14}$$

$$\mathbf{y} = \text{softmax}(\mathbf{x}^T \mathbf{W}_\theta^T \mathbf{W}_\theta \mathbf{x}) g(\mathbf{x}) \tag{15}$$

We will show that the non-local operation in Equation 14 is a generalisation of the Nadaraya-Watson kernel estimator. Specifically, when defining $C(\mathbf{x}) = \sum_l \mathcal{K}(\mathbf{q}_i, \mathbf{k}_l)$, $f = \mathcal{K}(\mathbf{q}_i, \mathbf{k}_l)$, $g = \mathbf{v}_j$, then we obtain the original Nadaraya-Watson kernel estimator in Equation 4. The variety of kernel decompositions of attention demonstrate the robustness of the kernel interpretation of attention and hence opening up an alley in attention improvement through the lens of kernel.

A.2.2 KERNEL APPROXIMATION BEYOND RANDOM FEATURES

Kernel approximation is presented in Section 2.3 through the use of random features. However, there has been other kernel approximation techniques outside of random features. These include greedy basis selection (Smola & Schölkopf, 2000), divide and conquer (Hsieh et al., 2014; Zhang et al., 2013), and nystrom methods (Williams & Seeger, 2000) (which led to the development of Skyformer (Chen et al., 2021) and Nystromformer (Xiong et al., 2021)).

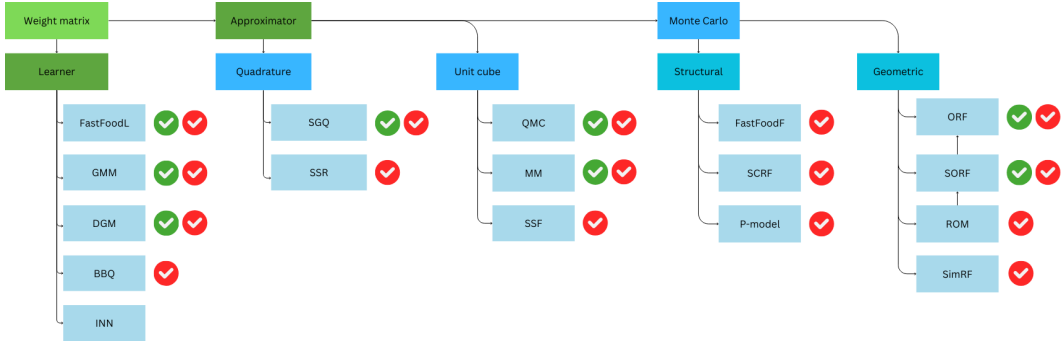


Figure 2: Spectraformer weight matrices. Red tick denotes techniques which have been explored in the literature. Green tick denotes techniques which are explored in this work. Please refer to Appendix A.3 for references.

A.3 TECHNICAL DETAILS OF WEIGHT APPROXIMATORS

Spectraformer weight matrices are shown in Figure 2. Weight matrices are either approximators or learners. There are three approaches to approximators: Monte Carlo sampling, Unit-cube sampling, and quadrature. We have discussed these approaches in detail in Section 3.1. Here we provide a more specific classification schemes for each method.

- **Monte Carlo sampling**-based weight matrices are improvements over Base and subdivided into two types: structural and geometric methods. Structural methods decompose the Base random matrices P into smaller matrices which reduce time and/ or space complexity. Structural matrices includes FastFood_F, SCRf and P-Model. Among these, we have explored the learner variant FastFood_L of FastFood. Geometric methods enforce certain geometrical couplings along some dimensions in the construction of random matrices to reduce the approximation error. Geometric methods include: ROM (ORF, SORF) and SimRF. Among these, we have explored the main ROM variants - ORF and SORF.
- **Unit-cube sampling**-based weight matrices reformulate the Bochner integral in Equation 9 as Equation A.3 with uniform and independent points. This approach includes QMC, MM, and SSF. We have not covered SSF in our experiments.

- **Quadrature**-based weight matrices estimates the intractable integral using non-uniform weights with deterministic rules. There are two approaches: SGQ and SSR. We have not covered SSR in our experiments.

We now explore these weight matrices in the order specified above.

Base (Rahimi & Recht, 2007) is the direct sampling from a distribution $P_{ij} \sim p$, corresponding to the kernel \mathcal{K} without any additional adjustment, and in approximating the RBF, $P = G$ where G is the normal distribution.

$$W_{Base} = \frac{1}{\sigma} \mathbf{P} \quad (16)$$

FastFood_F (Le et al., 2013) makes use of Hadamard and diagonal matrices to speed up Gaussian matrices G constructions for RFF in $O(n \log d)$ (100 times faster) with $O(n)$ space (1000 times less space) given $n \geq d$ against Base. However, FastFood_F increases variance and approximation error and decreases the concentration bound.

$$W_{FastFood_L} = \frac{1}{\sigma} \mathbf{S} \mathbf{H} \mathbf{G} \mathbf{\Pi} \mathbf{H} \mathbf{B} \quad (17)$$

where \mathbf{H} is the Walsh-Hadamard matrix, $\mathbf{\Pi} \in \{0, 1\}^{d \times d}$ is a permutation matrix, and $\mathbf{S}, \mathbf{G}, \mathbf{B}$ are diagonal random matrices, with the diagonal entries being $\{+ - 1\}$ entries on \mathbf{B} , random Gaussian entries on \mathbf{G} , and a random scaling matrix on \mathbf{S} , $S_{ii} = s_i \|\mathbf{G}\|_{Frob}^{-1/2}$, $s_i \sim (2\pi)^{-\frac{d}{2}} \mathbf{A}_{d-1}^{-1} r^{d-1} e^{-\frac{r^2}{2}}$.

SCRF (Signed Circulant Matrix Projection) (Feng et al., 2015) maintains a high matrix construction speed and $O(s)$ space ($s = t * d$) whilst maintaining the variance of Base, an improvement over FastFood_F. It does this by making use of circulant matrices.

$$\mathbf{W}_{SCRF} = [\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \dots, \mathbf{P}^{(t)}], t = d|D \quad (18)$$

where $\mathbf{P}^{(i)}$ is a signed circulant Gaussian matrix, with the first column of circulant random matrix $\mathbf{C}^{(i)[d]}$ of $\mathbf{P}^{(i)}$ drawn randomly from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}/\sigma^2)$. The circulant matrix via Discrete Fourier Transform can be defined as:

$$\begin{aligned} \mathbf{C} &= \frac{1}{m} \mathbf{F}^* \text{diag}(\mathbf{F} \mathbf{c}) \mathbf{F} \\ \mathbf{F} &= [e^{i \frac{2\pi}{m} kn}]_{k,n=0}^{m-1} \end{aligned} \quad (19)$$

P-Model (Choromanski & Sindhvani, 2016) generalizes over both circulant matrices (including SCRF), FastFood_F, along with other Gaussian and semi-Gaussian structured matrix methods.

$$\mathbf{W}_{P_{model}} = [\mathbf{g}^T \mathbf{P}_1, \mathbf{g}^T \mathbf{P}_2, \dots, \mathbf{g}^T \mathbf{P}_s]^T \quad (20)$$

where $\mathbf{g} \in \mathbb{R}^a$ is a Gaussian vector and a sequence of m $\mathbf{P}_i \in \mathbb{R}^{a \times d}$ matrices with L2 norm columns, where \mathbf{P}_i represents a P-model. Semi-Gaussian structured matrices including Toeplitz, Hankel, and other Toeplitz-like matrices are an open area of research for random matrices.

ORF (Orthogonal RF) (Yu et al., 2016) decreases the approximation error significantly compared to Base. This is done via replacing G with a properly scaled random orthogonal matrix. However, generating orthogonal matrices become costly quickly as the number of dimensions increases.

$$\mathbf{W}_{ORF} = \frac{1}{\sigma} \mathbf{S} \mathbf{Q} \quad (21)$$

where \mathbf{Q} is a uniformly distributed random orthogonal matrix (on the Stiefel manifold) obtained from the QR decomposition of \mathbf{G} , the set of rows of \mathbf{Q} forming a basis in \mathbb{R}^d , and \mathbf{S} is a diagonal matrix with entries sampled i.i.d from the χ -distribution with d degrees of freedom, thus making the rows of $\mathbf{S} \mathbf{Q}$ and \mathbf{G} identically distributed.

SORF (Structured ORF) (Yu et al., 2016) decreases time and space complexity of ORF (from $O(d^2)$

to $O(d \log d)$ with almost no extra memory cost) by imposing structure on the orthogonal matrices, inspired by structural methods. SORF is unbiased with large d . We replace \mathbf{S} with \sqrt{d} and \mathbf{Q} with structured matrix $\mathbf{H}\mathbf{D}_1\mathbf{H}\mathbf{D}_2\mathbf{H}\mathbf{D}_3$

$$\mathbf{W}_{SORF} = \frac{\sqrt{d}}{\sigma} \mathbf{H}\mathbf{D}_1\mathbf{H}\mathbf{D}_2\mathbf{H}\mathbf{D}_3 \quad (22)$$

where $\mathbf{D}_i \in \mathbb{R}^{d \times d}$, $i = 1, 2, 3$ is diagonal 'sign-flipping' matrices with diagonal entries sampled from the Rademacher distribution, \mathbf{H} is the normalized Walsh-Hadamard matrix.

ROM (Random Orthogonal Embeddings) (Choromanski et al., 2017) generalizes SORF ($t = 3$) to:

$$\mathbf{W}_{ROM} = \frac{\sqrt{d}}{\sigma} \prod_{i=1}^t \mathbf{S}\mathbf{D}_i \quad (23)$$

where \mathbf{S} is a class of L_2 -normalized versions of Kronecker product matrices, of which the Hadamard \mathbf{H} is a representative of, and \mathbf{D} being independent diagonal matrices. When $\mathbf{S} = \mathbf{H}$ and $t = 3$, we obtain SORF used in our experiments. Another ROM variant is S-Hybrid, which we do not experiment with, given as

$$\mathbf{W}_{SHybrid} = \mathbf{S}\mathbf{D}_k^{(U)} \prod_{i=1}^{k-1} \mathbf{S}\mathbf{D}_i^{(R)} \quad (24)$$

where $\mathbf{D}_k^{(U)}$ is a diagonal matrix with i.i.d. $Unif(S^1)$ on the unit circle of \mathbb{C} . When being used in the component function, instead of calculating $\phi(\mathbf{x})^T \phi(\mathbf{y})$ in Equation 10. We model the equation as $Real(\overline{\phi(\mathbf{x})}^T \phi(\mathbf{y}))$.

SimRF (Simplex RF) (Reid et al., 2023) is developed to be a positive RF with the most optimal solution, the MSE of which is the lowest in the geometric methods. This is an alternative to the ROM-based geometric methods. Two variants of SimRF are: weight-independent and weight-dependent. SimRF is defined as:

$$\mathbf{W}_{simp} = \mathbf{D}\mathbf{S}\mathbf{R} \quad (25)$$

where $\mathbf{D} \in \mathbb{R}^{d \times d} = diag(\omega_i)$, $\mathbf{R} \in \mathbb{R}^{d \times d}$ being a random orthogonal matrix drawn from Haar measure on $O(d)$ and $\mathbf{S} \in \mathbb{R}^{d \times d}$ with row \mathbf{s}_i :

$$\mathbf{s}_i = \begin{cases} \sqrt{\frac{d}{d-1}} \mathbf{e}_i - \frac{\sqrt{d+1}}{(d-1)^{3/2}} (1, \dots, 1, 0)^T & 1 \leq i < d \\ \frac{1}{\sqrt{d-1}} (1, 1, \dots, 1, 0)^T & i = d \end{cases} \quad (26)$$

being unit vectors which are manifestly normalized and subtend obtuse angles. SimRF variants are defined based on how ω_i is constructed. For SimRF_{indep}, $\omega_i \sim \chi_d$. For Sim_{dep}, we permit the random vector direction $\{\hat{\omega}_i\}$ to be correlated with norms $\{\omega_i\}$:

$$\omega_i = -\frac{\sum_{j \neq i} \omega_j}{\|\sum_{j \neq i} \omega_j\|_2} \omega_i, i = 1, \dots, d \quad (27)$$

QMC (Quasi-Monte Carlo) (Avron et al., 2016) evaluates on a low discrepancy sequence (e.g., Halton, Sobol' Faure, and Niederreiter) of points instead of random points in Monte Carlo. Although the approximation error is only reduced minimally, QMC has been shown to perform better than MC in high dimensions and does not have undesirable clustering effect. To calculate QMC, we first assume that $p(\mathbf{x}) = \prod_{j=1}^d p_j(\mathbf{x}_j)$ factorizes with respect to the dimensions with $p_j(\cdot)$ being a univariate density function. Then we define:

$$\Phi^{-1}(\mathbf{t}) = (\Phi_1^{-1}(t_1), \dots, \Phi_d^{-1}(t_d)) \in \mathbb{R}^d \quad (28)$$

where Φ_j being the cumulative distribution function of p_j , $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_s \in [0, 1]^d$ being a low discrepancy sequence, and $\omega_i = \Phi^{-1}(\mathbf{t}_i)$. We can thus transform the integral on \mathbb{R}^d in Equation 9 to an integral on the unit cube $[0, 1]^d$ as

$$\mathcal{K}(\mathbf{x} - \mathbf{x}') = \int_{[0,1]^d} \exp(i(\mathbf{x} - \mathbf{x}')^\top \Phi^{-1}(\mathbf{t})) d\mathbf{t}$$

The weight matrix then can be defined as:

$$\mathbf{W}_{\text{QMC}} = [\Phi^{-1}(\mathbf{t}_1), \Phi^{-1}(\mathbf{t}_2), \dots, \Phi^{-1}(\mathbf{t}_s)]^\top \in \mathbb{R}^{s \times d}.$$

QMC can be further improved by a sub-grouped based rank-one lattice construction which improved complexity (Lyu et al., 2020).

MM (Moment Matching) (Shen et al., 2017; Liu et al., 2022) improves over QMC by removing the undesirable clustering effect and having the same approximation error with less features. This is done by replacing Φ^{-1} with a moment matching scheme $\tilde{\Phi}^{-1}$:

$$\mathbf{W}_{\text{MM}} = [\tilde{\Phi}^{-1}(\mathbf{t}_1), \tilde{\Phi}^{-1}(\mathbf{t}_2), \dots, \tilde{\Phi}^{-1}(\mathbf{t}_s)]^\top \in \mathbb{R}^{s \times d}$$

where $\tilde{\Phi}^{-1}(\mathbf{t}_i) = \tilde{\mathbf{A}}^{-1}(\Phi^{-1}(\mathbf{t}_i) - \tilde{\boldsymbol{\mu}})$ can be constructed using moment matching with sample mean $\tilde{\boldsymbol{\mu}} = \frac{1}{s} \sum_{i=1}^s \Phi^{-1}(\mathbf{t}_i)$ and the square root of the sample covariance matrix $\tilde{\mathbf{A}}$ satisfying $\tilde{\mathbf{A}}\tilde{\mathbf{A}}^\top = \text{Cov}(\Phi^{-1}(\mathbf{t}_i) - \tilde{\boldsymbol{\mu}})$.

SSF (Spherical Structured Features) (Lyu, 2017) improves over QMC by including rotation-invariant kernel as well as in terms of complexity. Rotation-invariant property suggests the construction of feature maps using spherical equal weight approximation using Riesz s -energy on a d -dimensional sphere $\mathbb{S}^d := \{\mathbf{x} \in \mathbb{R}^{d+1} \mid \|\mathbf{x}\|_2 = 1\}$. Specifically, we construct $\{\mathbf{v}_i\}_{i=1}^s$ asymptotically uniformly distributed on the sphere and obtain $\mathbf{V} := [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_s] \in \mathbb{S}^{(d \times s)}$, that $\Phi^{-1}(t)$ uses the one-dimensional QMC point, we have the weight matrix:

$$\mathbf{W}_{\text{SSF}} = [\Phi^{-1}(t)\mathbf{v}_1, \Phi^{-1}(t)\mathbf{v}_2, \dots, \Phi^{-1}(t)\mathbf{v}_s]^\top \in \mathbb{R}^{s \times (d+1)},$$

SGQ (Sparse Grid Quadrature) (Dao et al., 2017) evolves from GQ. GQ (Gaussian Quadrature) (Dao et al., 2017) assumes that the kernel k factorizes with respect to the dimensions and thus can be approximated by a one-dimensional Gaussian quadrature rule. However the total number of points s scale exponentially with the dimensions. However GQ suffers from the curse of dimensionality. This is alleviated using Smolyak rule resulting in SGQ. Assuming the third-degree SGQ using symmetric univariate quadrature points $\{-\hat{p}_1, 0, \hat{p}_1\}$ with weights $(\hat{a}_1, \hat{a}_0, \hat{a}_1)$, then we have:

$$\mathbf{W}_{\text{SGQ}} = [\mathbf{0}_d, \hat{p}_1 \mathbf{e}_1, \dots, \hat{p}_1 \mathbf{e}_d, -\hat{p}_1 \mathbf{e}_1, \dots, -\hat{p}_1 \mathbf{e}_d]^\top \in \mathbb{R}^{(2d+1) \times d}$$

given that \mathbf{e}_i is the d -dimensional standard basis vector with the i^{th} element being 1.

SSR (Spherical-Radial Rules) (Munkhoeva et al., 2018) transforms Equation 10 into a double integral over a hyper-sphere and the real line. This leads to the following approximation:

$$\mathbf{W}_{\text{SSR}} = \boldsymbol{\vartheta} \otimes \begin{bmatrix} (\mathbf{Q}\mathbf{V})^\top \\ -(\mathbf{Q}\mathbf{V})^\top \end{bmatrix} \in \mathbb{R}^{(2d+1) \times d},$$

where $\boldsymbol{\vartheta} = [\vartheta_1, \vartheta_2, \dots, \vartheta_s]$, $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{d+1}]$, $\vartheta \sim \chi(d+2)$ and $\{\mathbf{v}_i\}_{i=1}^{d+1}$ being vertices of a unit regular d -simplex randomly rotated by \mathbf{Q} (a random orthogonal matrix).

A.4 TECHNICAL DETAILS OF WEIGHT LEARNERS

We give a non-mathematical definition of weight matrix learner in Section 3.1: weight matrix learner is any function which parameterize distributions. This allows for the discovery and future experimentation with other potential weight matrix learners in Spectraformer and random feature method in general. Weight matrix learner comes from a long line of work in kernel methods, namely, kernel learning via random features. Typically, a component function f is chosen and a weight matrix \mathbf{W} is produced via kernel approximation, providing us with the feature map ϕ . ϕ can then be combined with input \mathbf{x} and weight \mathbf{v} to fit the learning objective. Kernel learning, however, imposes an additional objective: learning the weight matrix \mathbf{W} .

The approach to the dual objectives separates the kernel learning methods into two-stage and one-stage methods (Liu et al., 2022). Two-stage methods (e.g., Yu et al. (2015); Wilson & Adams (2013);

Bullins et al. (2018)) solves the dual objective separately: \mathbf{W} is typically first learned via a kernel alignment scheme (Wang et al., 2015), then we solve for v . One-stage methods (e.g., Chowdhury et al. (2022); Tompkins et al. (2019)), on the other hand, solves the dual objective in parallel: $p(\cdot)$ corresponding to \mathbf{W} is parameterized then v is solved typically. It is not plausible to apply kernel alignment (Wang et al., 2015) on Transformer architectures. Therefore, only one-stage methods can be considered in the context of weight matrix learner. Two kernel learning methods, GMM and FastFood_L, have been experimented in the non-Transformer setting (Yang et al., 2015), then adapted into the Transformer setting (Chowdhury et al., 2022). DGM (Deep Generative Model) has also been effective in modeling distributions (Kingma & Welling, 2013). It too has been adapted by Chowdhury et al. (2022) for the Transformer setting. Due to the superior performance of FastFood_L with this technique, we have only experimented with this technique. We now explore them.

FastFood_L (see Equation 17) is proposed to be learnable by Yang et al. (2015) with two variations: FSARD (scaling matrix \mathbf{S} previously sampled from chi-squared is made learnable) and FSGARD (optimizes the marginal likelihood with respect to the diagonal matrices \mathbf{G} and \mathbf{B}). It is further adapted by Chowdhury et al. (2022) to make \mathbf{S} or $\mathbf{S}, \mathbf{G}, \mathbf{B}$ learnable parameters. In our experiments, we maintain this set up.

GMM (Gaussian Mixture Model) (Wilson & Adams, 2013; Oliva et al., 2016) is a known universal approximator of probability distribution. Specifically, we can consider \mathbf{W}_{GMM} being composed of s components, with each component c (with C total components) being a Gaussian with mean vector μ_c and covariance matrix Σ_c and component weight π_c . $p(\omega)$ can then be approximated as $p(\omega) = \sum_{c=1}^C \pi_c \mathcal{N}(\mu_c, \Sigma_c)$ (Chowdhury et al., 2022). Assuming $\pi_c = \frac{1}{C}$, the weight matrix with learnable weights Σ and μ can be parameterized as:

$$\mathbf{W}_{GMM} = [\Sigma_1 n + \mu_1, \dots, \Sigma_C n + \mu_C]^T, n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (29)$$

Another formulation of GMM is proposed in Yang et al. (2015) in combination with FastFood, albeit it extends beyond constructing the weight matrix.

DGM (Deep Generative Model) (Kingma & Welling, 2013) allows for a generic distribution modeling using neural network. DGM is initially proposed to make differentiating sampling processes possible via the reparametrization trick, where the generator g creates samples from the target distribution $\omega_m = g(\mathbf{n}_m)$ via transforming from a noise distribution $\mathbf{n}_m \sim p_0(\cdot)$.

In addition to the techniques discussed above, there are also some other prominent one-stage methods (e.g., INN, BBQ). INN (Invertible Neural Network) (Xu et al., 2021) is another universal approximator of probability distribution, consisting of invertible operations transforming samples from a known distribution to a different more complex distribution. BBQ (Black Box Quantile) (Tompkins et al., 2019) follows the QMC scheme and models Φ^{-1} using parameterized quantile function.

A.5 TECHNICAL DETAILS OF COMPONENT FUNCTIONS

Component functions discussed in Section 3.2 all belong to CDERF. However, there are other component functions including DIRF and HRF (see Figure 3) which we do not explore in our experiments. We now investigate the theoretical details of these component functions, the CDERF family is introduced in their order of discovery.

CDERF (Complex dense exponential random feature) is given in Equation 30, where $\omega, \mathbf{x} \in \mathbb{R}^d$, $p = \mathcal{N}(0, 1)^d$, $k \in \{1, 2\}$ and \mathbb{S}_{d_c} being a set of $d * d$ complex symmetric matrices. The parameters and constraints of different CDERF functions are specified in Table 7:

$$f_{DE}^{(k)}(\omega, \mathbf{x}) = \text{Real}(D \exp(\omega^T \mathbf{A} \omega + \omega^T \mathbf{B}^{(k)} \mathbf{x} + \mathbf{x}^T \mathbf{C}^{(k)} \mathbf{x})) \quad (30)$$

TrigRF (Trigonometric RF) (Rahimi & Recht, 2007) is the base RFF implementation.

$$\begin{aligned} f_{TrigRF}^{(1)}(\omega, \mathbf{x}) &= \sqrt{2} \cos(\omega^T \mathbf{x} + \mathbf{b}) \\ f_{TrigRF}^{(2)}(\omega, \mathbf{y}) &= \sqrt{2} \sin(\omega^T \mathbf{y} + \mathbf{b}) \\ \mathbf{b} &\sim \text{Uniform}(0, 2\pi) \end{aligned} \quad (31)$$

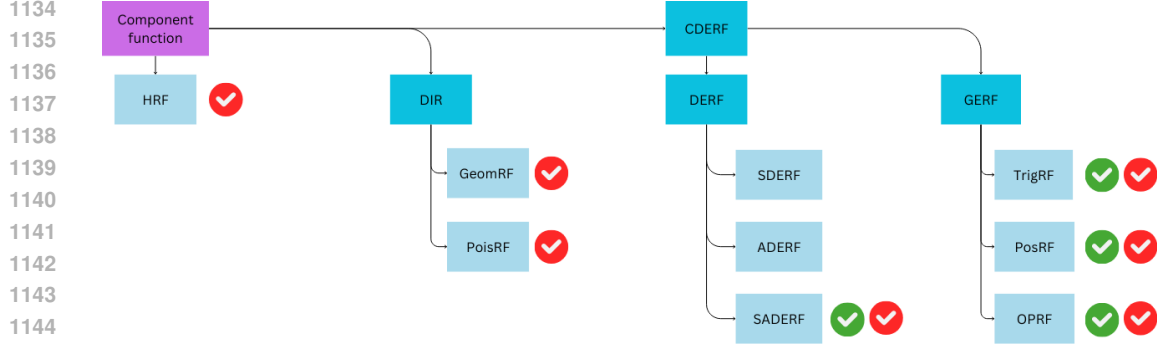


Figure 3: Spectraformer component functions. Red tick denotes techniques which have been explored in the literature. Green tick denotes techniques which are explored in this work. Please refer to Appendix A.5 for references.

Table 7: CDERF component functions

	A	$B^{(1)}$	$B^{(2)}$	$C^{(1)}$	$C^{(2)}$	D	x	y
CDERF	S_{d_C}	$\mathbb{C}^{d \times d}$		$\mathbb{C}^{d \times d}$		\mathbb{R}	x	y
DERF (Likhoshesterov et al., 2023)	S_d	$\mathbb{R}^{d \times d}$		$\mathbb{R}^{d \times d}$		\mathbb{R}	x	y
SADERF (Likhoshesterov et al., 2023)	$A_{GE} \mathbf{I}_d$	$B_{GE} \mathbf{I}_d$	$B_{GE} \mathbf{I}_d$	$C_{GE} \mathbf{I}_d$		D_{GE}	Ψx	$\Psi^{-1} y$
GERF (Likhoshesterov et al., 2022)	$A_{GE} \mathbf{I}_d$	$B_{GE}^{(1)} \mathbf{I}_d$	$B_{GE}^{(2)} \mathbf{I}_d$	$C_{GE} \mathbf{I}_d$		D_{GE}	x	y
TrigRF (Rahimi & Recht, 2007)	$0 \mathbf{I}_d$	$i \mathbf{I}_d$	$-i \mathbf{I}_d$	$0 \mathbf{I}_d$		1	x	y
PosRF (Choromanski et al., 2021)	$0 \mathbf{I}_d$		$1 \mathbf{I}_d$	$-1 \mathbf{I}_d$		1	x	y
OPRF (Likhoshesterov et al., 2022)	$0 \mathbf{I}_d$		$1 \mathbf{I}_d$	$-1 \mathbf{I}_d$		1	x	y

The use of the trigonometric sine and cosine functions leads to unstable behavior when the inputs have negative dimension-values. This can be further exacerbated when the values TrigRF try to approximate are close to 0 (since most values are of low significance). This causes the variance to approach infinity (Choromanski et al., 2021; Likhoshesterov et al., 2022). Therefore, we do not want to use TrigRF to perform kernel approximation in attention.

PosRF (Positive RF) (Choromanski et al., 2021) fixes the problem of TrigRF by enforcing positive component function output in the softmax. The variance of PosRF (in contrast to the variance of TrigRF approaching infinity) approaches 0 as the approximated value of the Softmax kernel approaches 0. PosRF has two forms: Positive RF Base (PosRF-B) (see Equation 32) and Positive RF Hyperbolic (PosRF-Hyp), which is multi-component, i.e., Equation 11: $l = 2$, (see Equation 33).

$$f_{PosRF_B}(\omega, \mathbf{x}) = \exp(\omega^T \mathbf{x} - \frac{\|\mathbf{x}\|^2}{2}) \quad (32)$$

$$\begin{aligned} f_{1,PosRF_{Hyp}}(\omega, \mathbf{x}) &= \exp(\omega^T \mathbf{x} - \|\mathbf{x}\|^2) \\ f_{2,PosRF_{Hyp}}(\omega, \mathbf{x}) &= \exp(-\omega^T \mathbf{x} - \|\mathbf{x}\|^2) \end{aligned} \quad (33)$$

We only use PosRF-B in our experiments.

GERF (Generalized exponential RF) (Likhoshesterov et al., 2022) generalizes both TrigRF and PosRF with Equation 34.

$$\begin{aligned} f_{GERF}^{(1)}(\omega, \mathbf{x}) &= D \exp(A \|\omega\|^2 + B \omega^T \mathbf{x} + C \|\mathbf{x}\|^2) \\ f_{GERF}^{(2)}(\omega, \mathbf{y}) &= D \exp(A \|\omega\|^2 + s B \omega^T \mathbf{y} + C \|\mathbf{y}\|^2) \end{aligned} \quad (34)$$

$$\begin{aligned} \text{Re}(1 - 4A) &> 0, \mathbf{B} = \sqrt{s(1 - 4A)} \\ \mathbf{C} &= -(s + 1)/2, D = (\sqrt[4]{1 - 4A})^d \\ \mathbf{A} &\in \mathbb{C}, s \in \{-1, +1\} \end{aligned}$$

where $\sqrt{\cdot}$ and $\sqrt[4]{\cdot}$ denoting a principal root with a complex argument

1188

1189

1190

1191

OPRF (Optimized positive RF) (Likhoshevstov et al., 2022) is the solution to the minimization of the variance of GERF. Specifically it is defined as Equation 34 with $s = +1$, $\|\mathbf{x} + \mathbf{y}\|^2 > 0$, and $\mathbf{A} \in \mathbb{R}$ defined in terms of p^* as:

1192

1193

1194

1195

$$p^* = \frac{\mathbf{A} = (1 - 1/p^*)/8}{(\sqrt{(2\|\mathbf{x} + \mathbf{y}\|^2 + d)^2 + 8d\|\mathbf{x} + \mathbf{y}\|^2} - 2\|\mathbf{x} + \mathbf{y}\|^2 - d) / (4\|\mathbf{x} + \mathbf{y}\|^2)} \quad (35)$$

1196

1197

1198

Whilst PosRF is not bounded, OPRF is. OPRF can provide $e^{60} \times$ variance reduction in estimating the Softmax compared to TrigRF.

1199

1200

1201

1202

1203

1204

DERF (Dense-exponential random features) (Likhoshevstov et al., 2023) extends GERF and replace $\mathbf{A}, \mathbf{B}, \mathbf{C}$ with dense matrices. DERF is CDERF when \mathbf{B}, \mathbf{C} are in the real instead of the complex plane. Minimizing the variance of DERF leads to two approaches: ADERF and SDERF. However both these approaches rely on SVD and eigen decompositions which are not extensively supported on GPU and deep learning libraries. Therefore SADERF is proposed.

1205

1206

1207

1208

1209

1210

SADERF (simplified ADERF) (Likhoshevstov et al., 2023) is a special case of ADERF and extends GERF, requiring only basic unary operations in addition.

$$f_{SADERF}^{(1)}(\boldsymbol{\omega}, \mathbf{x}) = f_{GE}^{(1)}(\boldsymbol{\omega}, \Psi \mathbf{x}), f_{SADERF}^{(2)}(\boldsymbol{\omega}, \mathbf{y}) = f_{GE}^{(2)}(\boldsymbol{\omega}, \Psi^{-1} \mathbf{y}) \quad (36)$$

$$\Psi_{l,l}^* = (\sum_j (\mathbf{y}_l^{(j)})^2 / \sum_i (\mathbf{x}_l^{(i)})^2)^{1/4}$$

1211

1212

1213

1214

DIRF (Discretely-induced random features) (Likhoshevstov et al., 2022) are based on the assumption of $p(\boldsymbol{\omega})$ being a discrete distribution with $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_d$ being i.i.d., $\mathbb{P}(\boldsymbol{\omega}_l = k) = p_k$, $\sum_{k=0}^{\infty} p_k = 1$, $p_k > 0, k \in \{0\} \cup \mathbb{N}$. By making use of the Taylor series, we can define DIRF as:

1215

1216

1217

$$f_{DI}(\boldsymbol{\omega}, \mathbf{x}) = \exp(-\frac{\|\mathbf{x}\|^2}{2}) \prod_{l=1}^d \mathbf{x}_i^{\boldsymbol{\omega}_l} (\boldsymbol{\omega}_l!)^{-\frac{1}{2}} p_{\boldsymbol{\omega}_l}^{-\frac{1}{2}} \quad (37)$$

1218

1219

1220

1221

DIRF has two variations: PoisRF (Poisson RF) which requires $p(\cdot)$ to be the Poisson distribution and GeomRF (Geometric RF) which requires $p(\cdot)$ to be the geometric distribution. Since the many weight matrices introduced in Section 3.1 do not readily approximate these distributions, we decide to leave DIRF out of the paper.

1222

1223

1224

1225

1226

1227

HRF (Hybrid RF) (Choromanski et al., 2022) combines multiple base estimators, similar to multiple kernel learning methods (Bach et al., 2004), in order to provide the most accurate approximation in regions of interest. Choromanski et al. on solving softmax approximation using HRF. HRF is constructed using a weighted combination ϵ of $p + 1$ base estimators $\hat{S}\hat{M}^k(\mathbf{x}, \mathbf{y})$ and λ -coefficients, Λ p binary estimators $\hat{\lambda}^k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$, both constructed independently:

1228

1229

1230

$$\hat{S}\hat{M}^{\epsilon, \Lambda} = \sum_{k=1}^p \hat{\lambda}^k(\mathbf{x}, \mathbf{y}) \hat{S}\hat{M}^k(\mathbf{x}, \mathbf{y}) + (1 - \sum_{k=1}^p \hat{\lambda}^k(\mathbf{x}, \mathbf{y})) \hat{S}\hat{M}^{p+1}(\mathbf{x}, \mathbf{y}) \quad (38)$$

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Since HRF requires the base estimators $SM(\mathbf{x}, \mathbf{y})$ to be unbiased, the easiest solution is the complex exponential $\hat{S}\hat{M}_m^{ceexp}$ which can be derived directly from the expectation of the softmax as follows:

$$\hat{S}\hat{M}_m^{ceexp}(\mathbf{x}, \mathbf{y}) = \psi_A^m(\mathbf{x})^T \psi_{(A^T)^{-1}}^m(\mathbf{y})$$

$$\psi_M^m(u) = \frac{1}{\sqrt{m}} \exp(-\frac{(M\mathbf{u})^2}{2}) \quad (39)$$

$$(\exp(\boldsymbol{\omega}_i^T M\mathbf{u}), \dots, \exp(\boldsymbol{\omega}_m^T M\mathbf{u}))^T, \boldsymbol{\omega}_i \sim \mathcal{N}(\mathbf{0}, I_d)$$

To construct λ -coefficients, three methods are proposed: simultaneous accurate approximation of both small and large values, Gaussian Lambda Coefficients, and adaption to data admitting clustering structure.

A.6 EXPERIMENTAL DETAILS IN SPECTRAFORMER

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

Table 8: Hyper-parameters for Spectraformer, following the setting from Chen et al. (2021) due to limited computational power. L stands for ListOps, T for Text, and R for Retrieval. The seed that can be used to reproduce our experiments is 42.

	L	T	R
Embedding dim.	64	64	64
Transformer dim	64	64	64
Hidden dim	128	128	128
Head dim	32	32	32
Num. heads	2	2	2
Num. layers	2	2	2
Vocabulary size	32	512	512
Sequence length	2000	4000	4000
Dropout rate	0.1	0.1	0.1
Att. dropout rate	0.1	0.1	0.1
Pooling mode		mean	
Batch size	32	32	16
Learning rate	0.0001	0.0001	0.0002
Warmup steps	1000	80	800
Learning rate decay		linear	
Weight decay	0	0	0
Evaluation freq.	500	500	1000
Num. epochs	50k	50K	50k
Num. init steps	1k	3K	3k
Num. eval steps	62	200	300
Patience	10	10	10
Num. features	128	128	128

Table 9: Mean statistics across component functions for each task. Task is denoted on the top left cell of each table. Mean statistics include ‘acc’ (accuracy), ‘time’ (training time), and ‘mem’ (peak memory consumption).

ListOps	acc	time	mem	Text	acc	time	mem	Retrieval	acc	time	mem
PosRF	31.98	0.63	1.10	OPRF	61.82	1.40	2.54	PosRF	77.32	1.22	2.01
OPRF	36.11	0.74	1.28	PosRF	62.31	1.21	2.19	SADERF	78.85	1.43	2.54
SADERF	36.33	0.75	1.35	SADERF	62.33	1.39	2.69	OPRF	79.11	1.39	2.41