
NeuralEF: Deconstructing Kernels by Deep Neural Networks

Zhijie Deng^{1,2} Jiaxin Shi³ Jun Zhu¹

Abstract

Learning the principal eigenfunctions of an integral operator defined by a kernel and a data distribution is at the core of many machine learning problems. Traditional nonparametric solutions based on the Nyström formula suffer from scalability issues. Recent work has resorted to a parametric approach, i.e., training neural networks to approximate the eigenfunctions. However, the existing method relies on an expensive orthogonalization step and is difficult to implement. We show that these problems can be fixed by using a new series of objective functions that generalizes the EigenGame (Gemp et al., 2020) to function space. We test our method on a variety of supervised and unsupervised learning problems and show it provides accurate approximations to the eigenfunctions of polynomial, radial basis, neural network Gaussian process, and neural tangent kernels. Finally, we demonstrate our method can scale up linearised Laplace approximation of deep neural networks to modern image classification datasets through approximating the Gauss-Newton matrix. Code is available at <https://github.com/thudzj/neuraleigenfunction>.

1. Introduction

Kernel methods (Muller et al., 2001; Schölkopf et al., 2002) are among the most important and fundamental tools in machine learning (ML) for processing nonlinear data. Like most nonparametric methods, kernel methods have limited applicability in large-scale scenarios, and kernel approximation methods like *Random Fourier features* (RFFs) (Rahimi & Recht, 2007) and *Nyström method* (Nyström, 1930) are

¹Dept. of Comp. Sci. & Tech., BNRist Center, Tsinghua-Bosch Joint Center for ML, Tsinghua University; Peng Cheng Laboratory ²Qingyuan Research, Shanghai Jiao Tong University ³Microsoft Research New England. Correspondence to: Jun Zhu <dcszj@tsinghua.edu.cn>.




Figure 1: NeuralEF embeds the structures inside a kernel into an ensemble of NNs under the principle of eigendecomposition. NeuralEF defines a new axis system for data, and is capable of recovering the kernel according to Mercer’s theorem.

frequently introduced as a remedy.

Even so, kernel methods still lag behind in handling complex high-dimensional data as classic local kernels suffer from the curse of dimensionality (Bengio et al., 2005). New developments such as neural network Gaussian process (NN-GP) kernels (Lee et al., 2017) and neural tangent kernels (NTKs) (Jacot et al., 2018) mitigate this issue by leveraging the inductive bias of neural networks (NNs). Yet, these modern kernels are expensive to evaluate and have poor compatibility with standard kernel approximation methods—an accurate approximation with RFFs requires many NN forward passes to construct an adequate number of features, while Nyström method involves costly kernel evaluation in the test phase.

We present a new kernel approximation method for addressing these challenges. It relies on deconstructing a kernel as a spectral series and approximating the eigenfunctions with neural networks. Because such a method can identify the principal components, we need much less NN forward passes than RFFs to accurately approximate NN-GP kernels and NTKs. Moreover, it gives rise to an unsupervised representation learning paradigm, where the pairwise similarity captured by kernels is embedded into NNs. The learned neural eigenfunctions can serve as plug-and-play feature extractors and be finetuned for downstream applications.

Training neural networks to approximate eigenfunctions is a difficult task due to the orthonormality constraints. Spectral Inference Networks (SpIN) (Pfau et al., 2018) made the first attempt by formulating this as an optimization problem. Nevertheless, SpIN’s objective function only allows recover-




Figure 2: Estimate the top-10 eigenfunctions $\{\hat{\psi}_j\}_{j=1}^{10}$ (we plot only the top-3 for visualization) of the polynomial kernel $\kappa(x, x') = (x^\top x' + 1.5)^4$ (Left) and the RBF kernel $\kappa(x, x') = \exp(-\|x - x'\|^2/2)$ (Middle) with the Nyström method (the ground truth), SpIN, and NeuralEF (our) given 64 data points. The results obtained with larger sample size are deferred to Appendix C.1. We also plot how the total training time varies w.r.t. sample size for the polynomial kernel (Right). The time complexity of Nyström method grows cubically with the sample size due to eigendecomposition. Both NeuralEF and SpIN consume nearly constant training time for different sample sizes, but SpIN produces less accurate approximations and is less efficient.

ing the subspace spanned by the top eigenfunctions¹ instead of the eigenfunctions themselves. To recover individual eigenfunctions, they resort to an expensive orthogonalization step that requires Cholesky decomposition and tracking the NN Jacobian matrix during training. As a result, SpIN suffer from inefficiency issues and non-trivial implementation challenges.

We resolve the problems of SpIN with a new formulation of kernel eigendecomposition as *simultaneously* solving a series of *asymmetric* optimization problems. Our work is inspired from Gemp et al. (2020) which uses a related objective to cast principal component analysis as a game, and can be viewed as an extension of it to the function space.

With this, we can naturally incorporate k NNs as function approximators to the top- k eigenfunctions of the kernel, and train them under a more amenable objective than SpIN. We leverage stochastic optimization strategies to comfortably train these neural eigenfunctions in the big data regime, and develop a novel NN layer to make the neural eigenfunctions fulfill certain constraints. We dub our method as *neural eigenfunctions*, or *NeuralEF* for short.

After training, NeuralEF approximates the kernel evaluation on novel data points with the outputs of the k NNs (see Figure 1). This brings great benefits when coping with modern kernels whose evaluation incurs heavy burden like NN-GP kernels and NTKs. Besides, NeuralEF forms a set of new orthogonal bases for the data and is compatible with various pattern recognition algorithms.

We empirically evaluate NeuralEF in a variety of scenarios. We first evidence that NeuralEF can perform as well as the

¹By saying top- k eigenpairs or eigenfunctions, we mean those associated with the top- k eigenvalues. Besides, we assume they are ranked in the descending order of the corresponding eigenvalues without loss of generality.

Nyström method and beat SpIN when handling classic kernels, yet with more sustainable resource consumption (see Figure 2). We then show that NeuralEF can learn the informative structures in NN-GP kernels and NTKs to achieve effective classification and clustering. To demonstrate the scalability of NeuralEF, we further conduct a large-scale study to scale up linearised Laplace approximation where the kernel matrix has size $500,000 \times 500,000$.

2. Background

2.1. Kernel Methods

Kernel methods project data into a high-dimensional feature space \mathcal{H} to enable linear manipulation of nonlinear data. The subtlety is that we can leverage the “kernel trick” to bypass the need of specifying \mathcal{H} explicitly—given a positive definite kernel $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, there exists a feature map $\varphi : \mathcal{X} \rightarrow \mathcal{H}$ such that $\kappa(\mathbf{x}, \mathbf{x}') = \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle_{\mathcal{H}}$.

There is a rich family of kernels. Classic kernels include the linear kernel, the polynomial kernel, the radial basis function (RBF) kernel, etc. However, they may easily fail when processing real-world data like images and texts due to inadequate expressive power and the curse of dimensionality. Thereby, various modern kernels which encapsulate the inductive biases of NN architectures have been developed (Wilson et al., 2016). The NN-GP kernels (Neal, 1996; Lee et al., 2017; Garriga-Alonso et al., 2018; Novak et al., 2018) and NTKs (Jacot et al., 2018; Arora et al., 2019; Du et al., 2019) are two representative modern *matrix-valued* kernels, as defined below:

$$\kappa_{\text{NN-GP}}(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\theta \sim p(\theta)} g(\mathbf{x}, \theta) g(\mathbf{x}', \theta)^\top, \quad (1)$$

$$\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\theta \sim p(\theta)} \partial_{\theta} g(\mathbf{x}, \theta) \partial_{\theta} g(\mathbf{x}', \theta)^\top, \quad (2)$$

where $g(\cdot, \theta) : \mathcal{X} \rightarrow \mathbb{R}^{N_{\text{out}}}$ denotes a function represented by an NN with weights θ , and $p(\theta)$ is a layerwise isotropic

Gaussian. When the NN specifying g has infinite width, $\kappa_{\text{NN-GP}}$ and κ_{NTK} have analytical formulae and can be computed recursively due to the hierarchical nature of NNs. The NTKs are valuable for the theoretical analysis of neural networks, but are used much less frequently in practice than the empirical NTKs defined below:

$$\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') = \partial_{\theta} g(\mathbf{x}, \boldsymbol{\theta}) \partial_{\theta} g(\mathbf{x}', \boldsymbol{\theta})^{\top}. \quad (3)$$

If there is no misleading, we refer to the empirical NTKs as NTKs in the following.

Despite the non-parametric flexibility, kernel methods suffer from inefficiency issues – the involved computations grow at least quadratically and usually cubically w.r.t. data size. Moreover, for NN-GP kernels and NTKs, writing down their detailed mathematical formulae is non-trivial (Arora et al., 2019) and evaluating them with recursion is both time and memory consuming.

2.2. Kernel Approximation

An idea to scale up kernel methods is to approximate the kernel with the inner product of some explicit vector representations of the data, i.e., $\kappa(\mathbf{x}, \mathbf{x}') \approx \nu(\mathbf{x})^{\top} \nu(\mathbf{x}')$, where $\nu: \mathcal{X} \rightarrow \mathbb{R}^k$ denotes a mapping function. *Random Fourier features* (RFFs) (Rahimi & Recht, 2007; 2008; Yu et al., 2016; Munkhoeva et al., 2018; Francis & Raimond, 2021) and *Nyström method* (Nyström, 1930; Williams & Seeger, 2001) are two popular approaches in this spirit. RFFs can easily handle shift-invariant kernels, but may face obstacles when being applied to the others. Besides, RFFs usually entail using a relatively large k .

Alternatively, Nyström method finds the eigenfunctions for kernel approximation according to Mercer’s theorem:

$$\kappa(\mathbf{x}, \mathbf{x}') = \sum_{j \geq 1} \mu_j \psi_j(\mathbf{x}) \psi_j(\mathbf{x}'), \quad (4)$$

where $\psi_j \in L^2(\mathcal{X}, q)^2$ denote the eigenfunctions of the kernel $\kappa(\mathbf{x}, \mathbf{x}')$ w.r.t. the probability measure q , and $\mu_j \geq 0$ refer to the corresponding eigenvalues.

Typically, the eigenfunctions obey the following equation:

$$\int \kappa(\mathbf{x}, \mathbf{x}') \psi_j(\mathbf{x}') q(\mathbf{x}') d\mathbf{x}' = \mu_j \psi_j(\mathbf{x}), \quad \forall j \geq 1. \quad (5)$$

And they are demanded to be orthonormal under q :

$$\int \psi_i(\mathbf{x}) \psi_j(\mathbf{x}) q(\mathbf{x}) d\mathbf{x} = \mathbb{1}[i = j], \quad \forall i, j \geq 1. \quad (6)$$

Given the access to a training set of i.i.d. samples $\mathbf{X}_{\text{tr}} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ from q , the Nyström method approximates the

² $L^2(\mathcal{X}, q)$ is the space containing all square-integrable functions w.r.t. q .

integration in Equation (5) by Monte Carlo (MC) estimation, which gives rise to

$$\frac{1}{N} \sum_{n'=1}^N \kappa(\mathbf{x}, \mathbf{x}_{n'}) \psi_j(\mathbf{x}_{n'}) = \mu_j \psi_j(\mathbf{x}), \quad \forall j \geq 1. \quad (7)$$

It is easy to see that we can get the approximate top- k eigenpairs $\{(\hat{\mu}_j, [\hat{\psi}_j(\mathbf{x}_1), \dots, \hat{\psi}_j(\mathbf{x}_N)]^{\top})\}_{j=1}^k$ of the kernel κ by eigendecomposing the kernel matrix $\kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}) \in \mathbb{R}^{N \times N}$. Nyström method plugs them back into Equation (7) for out-of-sample extension:

$$\hat{\psi}_j(\mathbf{x}) = \frac{1}{N \hat{\mu}_j} \sum_{n'=1}^N \kappa(\mathbf{x}, \mathbf{x}_{n'}) \hat{\psi}_j(\mathbf{x}_{n'}), \quad j \in [k]. \quad (8)$$

$[k]$ denotes set of integers from 1 to k . Then, we have $\kappa(\mathbf{x}, \mathbf{x}') \approx \sum_{j=1}^k \hat{\mu}_j \hat{\psi}_j(\mathbf{x}) \hat{\psi}_j(\mathbf{x}')$.

Unlike RFFs, Nyström method can be applied to approximate any positive-definite kernel. Yet, it is non-trivial to scale it up. On the one hand, the matrix eigendecomposition is costly for even medium sized training data. On the other hand, as shown in Equation (8), evaluating $\hat{\psi}_j$ on a new datum entails evaluating κ for N times, which is unaffordable when coping with the modern kernels specified by deep architectures.

2.3. Spectral Inference Networks (SpIN)

Kernel approximation with NNs has the potential to ameliorate these pathologies due to their universal approximation capability and parametric nature. SpIN (Pfau et al., 2018) is a pioneering work in this line. In a similar spirit to the Nyström method, SpIN recovers the top eigenfunctions for kernel approximation, yet with NNs. Concretely, SpIN introduces a vector-valued NN function $\Psi(\cdot, \mathbf{w}): \mathcal{X} \rightarrow \mathbb{R}^k$ and train it under the following eigendecomposition principle:

$$\begin{aligned} \max_{\mathbf{w}} \text{Tr} \left(\iint \Psi(\mathbf{x}, \mathbf{w}) \Psi(\mathbf{x}', \mathbf{w})^{\top} \kappa(\mathbf{x}, \mathbf{x}') q(\mathbf{x}) q(\mathbf{x}') d\mathbf{x} d\mathbf{x}' \right) \\ \text{s.t.: } \int \Psi(\mathbf{x}, \mathbf{w}) \Psi(\mathbf{x}, \mathbf{w})^{\top} q(\mathbf{x}) d\mathbf{x} = \mathbf{I}_k, \end{aligned} \quad (9)$$

where Tr computes the trace of a matrix and \mathbf{I}_k denotes the identity matrix of size $k \times k$; q is the empirical distribution of data.

However, the above objective makes Ψ recover the subspace spanned by the top- k eigenfunctions rather than the top- k eigenfunctions themselves (Pfau et al., 2018). Such a difference is subtle yet significant. In the extreme case of finding the top- N eigenfunctions given a training set \mathbf{X}_{tr} of size N , Equation (9) becomes:

$$\begin{aligned} \max_{\mathbf{w}} \text{Tr} \left(\Psi(\mathbf{X}_{\text{tr}}, \mathbf{w})^{\top} \kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}) \Psi(\mathbf{X}_{\text{tr}}, \mathbf{w}) \right) \\ \text{s.t.: } \Psi(\mathbf{X}_{\text{tr}}, \mathbf{w}) \in \mathbb{R}^{N \times N}, \Psi(\mathbf{X}_{\text{tr}}, \mathbf{w})^{\top} \Psi(\mathbf{X}_{\text{tr}}, \mathbf{w}) = \mathbf{I}_N, \end{aligned} \quad (10)$$

which equals to $\max_{\mathbf{w}} \text{Tr} (\Psi(\mathbf{X}_{\text{tr}}, \mathbf{w}) \Psi(\mathbf{X}_{\text{tr}}, \mathbf{w})^\top \kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}})) = \text{Tr} (\kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}))$. As the last term shows, the objective is independent of \mathbf{w} , thus it cannot be used to learn all N eigenfunctions.

To fix this issue, SpIN relies on a gradient masking trick to ensure that Ψ converges to ordered eigenfunctions. This solution is expensive as it involves a Cholesky decomposition per training iteration. Besides, to debias the stochastic optimization, SpIN involves tracking the exponential moving average (EMA) of the Jacobian $\partial_{\mathbf{w}} \Sigma_{\mathbf{w}}$. As a result, SpIN suffers from inefficiency issues (see Figure 2) and non-trivial implementation challenges. The explicit manipulation of Jacobian also precludes SpIN from utilizing very deep architectures.

3. Methodology

3.1. Eigendecomposition as Maximization Problems

NeuralEF builds upon a new representation of the eigenpairs of the kernel, which leads to a more amenable learning principle for kernel approximation than those of the Nyström method and SpIN:

Theorem 1 (Proof in Appendix A.1). *The eigenpairs of the kernel $\kappa(\mathbf{x}, \mathbf{x}')$ can be recovered by simultaneously solving the following asymmetric maximization problems:*

$$\max_{\hat{\psi}_j} R_{jj} \quad \text{s.t.}: C_j = 1, R_{ij} = 0, \forall j \geq 1, i \in [j-1], \quad (11)$$

where $\hat{\psi}_j \in L^2(\mathcal{X}, q)$ represent the introduced approximate eigenfunctions, and

$$R_{ij} := \iint \hat{\psi}_i(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') \hat{\psi}_j(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x}, \quad (12)$$

$$C_j := \int \hat{\psi}_j(\mathbf{x}) \hat{\psi}_j(\mathbf{x}) q(\mathbf{x}) d\mathbf{x}. \quad (13)$$

$(R_{jj}, \hat{\psi}_j)$ will converge to the eigenpair associated with j -th largest eigenvalue of κ .

Intuitively, $C_j = 1$ guarantees the normalization of $\hat{\psi}_j$. $R_{ij} = 0$ ($\forall i < j$) ensures that $\hat{\psi}_j$ and $\hat{\psi}_i$ are orthogonal, and induces a hierarchy among the problems – the solutions to the latter problems are restricted in the *orthogonal complement* of the subspace spanned by the solutions to the former ones. This way, we bypass the reliance on the troublesome Cholesky decomposition, gradient masking, and manipulation of Jacobian as required by SpIN.

Conceptually, R_{jj} can be thought of as an extension of the Rayleigh quotient. It is then interesting to see that our optimization objective makes an analogy with that of the EigenGame (Gemp et al., 2020) for finite-dimensional symmetric matrices. We clarify that our method is a generalization of EigenGame – by setting κ as the linear kernel

$\kappa(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\top \mathbf{x}'$, our objective forms a function-space generalization of that in EigenGame.

In the seek of tractability, we slack the constraints on orthogonality in Equation (11) as penalties and solve:

$$\max_{\hat{\psi}_j} R_{jj} - \sum_{i=1}^{j-1} \frac{R_{ij}^2}{R_{ii}} \quad \text{s.t.}: C_j = 1, \forall j \in [k], \quad (14)$$

where we set the penalty coefficients as $\frac{1}{R_{11}}, \dots, \frac{1}{R_{(j-1)(j-1)}}$ to make the two kinds of forces in the objective on a similar scale, as suggested by EigenGame (Gemp et al., 2020). We consider only the top- k eigenpairs to strike a balance between efficiency and effectiveness. $k \in \mathbb{N}^+$ is a tunable hyper-parameter. We will exposit how to handle the normalization constraints later.

Ideally, by Theorem 1, $\hat{\psi}_j$ would converge to the ground truth eigenfunction ψ_j , but in practice the exact convergence may be unattainable. Given this situation, we provide an analysis in Appendix A.2 to demonstrate that small errors in the solutions to the former problems would not significantly bias the solutions to the latter problems.

3.2. Neural Networks as Eigenfunctions

We opt to optimize over the parametric function class defined by NNs to scale up our method to large data. Concretely, we incorporate k NNs with the same architecture but dedicated weights into the optimization in Equation (14), i.e., $\hat{\psi}_j(\cdot) = \hat{\psi}(\cdot, \mathbf{w}_j) : \mathcal{X} \rightarrow \mathbb{R}, j \in [k]$.

Mini-batch training We learn the parameters via mini-batch training. Given $\mathbf{X} = \{\mathbf{x}_b\}_{b=1}^B \subseteq \mathbf{X}_{\text{tr}}$ at per iteration, we approximate R_{ij} by MC integration:

$$\begin{aligned} \tilde{R}_{ij} &= \sum_{b=1}^B \sum_{b'=1}^B \frac{1}{B^2} \hat{\psi}_i(\mathbf{x}_b) \kappa(\mathbf{x}_b, \mathbf{x}_{b'}) \hat{\psi}_j(\mathbf{x}_{b'}) \\ &= \frac{1}{B^2} \hat{\psi}_i^{\mathbf{X}} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_j^{\mathbf{X}}, \end{aligned} \quad (15)$$

where $\hat{\psi}_j^{\mathbf{X}} = [\hat{\psi}(\mathbf{x}_1, \mathbf{w}_j), \dots, \hat{\psi}(\mathbf{x}_B, \mathbf{w}_j)]^\top \in \mathbb{R}^B$ refer to the concatenation of scalar outputs and $\kappa^{\mathbf{X}, \mathbf{X}} = \kappa(\mathbf{X}, \mathbf{X}) \in \mathbb{R}^{B \times B}$ is the kernel matrix associated with the mini-batch of data. We track \tilde{R}_{jj} via EMA to get an estimate of the j -th largest eigenvalue μ_j .

L2 Batch normalization (L2BN) Then we settle the normalization constraints. When doing mini-batch training, we have $C_j \approx \frac{1}{B} \sum_{b=1}^B \hat{\psi}(\mathbf{x}_b, \mathbf{w}_j) \hat{\psi}(\mathbf{x}_b, \mathbf{w}_j)$. To make C_j equal to 1, we should guarantee the outcomes of each neural eigenfunction to have an L2 norm equal to \sqrt{B} . To realise this, we append an L2 batch normalization (L2BN) layer at the end of every involved NN, whose form is:

$$h_b^{\text{out}} = \frac{h_b^{\text{in}}}{\sigma}, \quad \text{with } \sigma = \sqrt{\frac{1}{B} \sum_{b=1}^B h_b^{\text{in}^2}}, \quad b \in [B]. \quad (16)$$

Algorithm 1 Find the top- k eigenpairs of a kernel by NeuralEF

- 1: **Input:** Training data \mathbf{X}_{tr} , kernel κ , batch size B .
- 2: Initialize NNs $\hat{\psi}_j(\cdot) = \hat{\psi}(\cdot, \mathbf{w}_j)$ and scalars $\hat{\mu}_j, j \in [k]$;
- 3: Compute the kernel matrix $\kappa^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}} = \kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}})$;
- 4: **for iteration do**
- 5: Draw a mini-batch $\mathbf{X} \subseteq \mathbf{X}_{\text{tr}}$; retrieve $\kappa^{\mathbf{X}, \mathbf{X}}$ from $\kappa^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}}$;
- 6: Do forward propagation $\hat{\psi}_j^{\mathbf{X}} = \hat{\psi}(\mathbf{X}, \mathbf{w}_j), j \in [k]$;
- 7: $\hat{\mu}_j \leftarrow \text{EMA}(\hat{\mu}_j, \frac{1}{B^2} \hat{\psi}_j^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_j^{\mathbf{X}}), j \in [k]$;
- 8: Compute $\nabla_{\mathbf{w}_j} \ell, j \in [k]$ by Equation (18) and do SGD;
- 9: **end for**

$\{h_b^{\text{in}} | h_b^{\text{in}} \in \mathbb{R}\}_{b=1}^B$ and $\{h_b^{\text{out}} | h_b^{\text{out}} \in \mathbb{R}\}_{b=1}^B$ denote the batched input and output. We apply the transformation in Equation (16) during training, while recording the EMA of σ for testing. Thereby, we sidestep the barrier of optimizing under explicit normalization constraints. Arguably, L2BN also ensures the neural eigenfunctions are square-integrable w.r.t. q .

The training loss and its gradients Given these setups, we arrive at the following training loss:

$$\min_{\mathbf{w}_1, \dots, \mathbf{w}_k} \ell = -\frac{1}{B^2} \sum_{j=1}^k \left(\hat{\psi}_j^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_j^{\mathbf{X}} - \sum_{i=1}^{j-1} \frac{\left(\text{sg}(\hat{\psi}_i^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_j^{\mathbf{X}}) \right)^2}{\text{sg}(\hat{\psi}_i^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_i^{\mathbf{X}})} \right), \quad (17)$$

where sg denotes the `stop_gradient` operation. The gradients of ℓ w.r.t. \mathbf{w}_j are:

$$\nabla_{\mathbf{w}_j} \ell = -\frac{2}{B^2} \kappa^{\mathbf{X}, \mathbf{X}} \left(\hat{\psi}_j^{\mathbf{X}} - \sum_{i=1}^{j-1} \frac{\hat{\psi}_i^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_j^{\mathbf{X}}}{\hat{\psi}_i^{\mathbf{X}^\top} \kappa^{\mathbf{X}, \mathbf{X}} \hat{\psi}_i^{\mathbf{X}}} \hat{\psi}_i^{\mathbf{X}} \right) \cdot \partial_{\mathbf{w}_j} \hat{\psi}_j^{\mathbf{X}}. \quad (18)$$

In practice, we analytically compute the first part of the RHS of Equation (18) and then perform vector-Jacobian product via reverse-mode autodiff.

Extension to matrix-valued kernels So far, we have figured out the manipulation of scalar kernels, and it is natural to generalize NeuralEF to handle *matrix-valued* kernels, whose outputs are $N_{\text{out}} \times N_{\text{out}}$ matrices. In this regime, a direct solution is to set $\hat{\psi}_j$ as NNs with N_{out} output neurons and modify the L2BN layer accordingly. Then, we have $\kappa^{\mathbf{X}, \mathbf{X}} \in \mathbb{R}^{B N_{\text{out}} \times B N_{\text{out}}}$ and $\hat{\psi}_j^{\mathbf{X}} = [\hat{\psi}(\mathbf{x}_1, \mathbf{w}_j)^\top, \dots, \hat{\psi}(\mathbf{x}_B, \mathbf{w}_j)^\top]^\top \in \mathbb{R}^{B N_{\text{out}}}$. The loss ℓ is still leveraged to guide training.

The algorithm Algorithm 1 shows the training procedure.³

3.3. Learning with NN-GP kernels and NTKs

The Nyström method, SpIN, and NeuralEF all entail computing the kernel matrix on training data $\kappa^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}}$, which

³It is not indispensable to precompute the big $\kappa^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}}$ – we can instead compute the small $\kappa^{\mathbf{X}, \mathbf{X}}$ at per training iteration.

may be frustratingly difficult for the NN-GP kernels and NTKs. As a workaround, we suggest approximately computing $\kappa^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}}$ with MC estimation based on *finitely wide* NNs for NN-GP kernels and NTKs instead of performing analytical evaluation (see also Novak et al. (2018)).

Specifically, for the NN-GP kernels, it is easy to see:

$$\kappa_{\text{NN-GP}}^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}} \approx \frac{1}{S} \sum_s g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}_s) g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}_s)^\top, \quad (19)$$

with $g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}_s) \in \mathbb{R}^{B N_{\text{out}}}$ as the concatenation of the vectorized outputs of g and $\boldsymbol{\theta}_s \sim p(\boldsymbol{\theta}), s \in [S]$. In other words, we evaluate a finitely wide NN g on the training data \mathbf{X}_{tr} for S times under various weight configurations to calculate the NN-GP training kernel matrix.

For the NTKs, it is hard to explicitly calculate and store the Jacobian matrix for modern NNs. Nevertheless, we notice that if there exists a distribution $p(\mathbf{v})$ satisfying $\mathbb{E}_{p(\mathbf{v})}[\mathbf{v} \mathbf{v}^\top] = \mathbf{I}_{\dim(\boldsymbol{\theta})}^4$, then

$$\kappa_{\text{NTK}}^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}} = \mathbb{E}_{\mathbf{v} \sim p(\mathbf{v})} [\partial_{\boldsymbol{\theta}} g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}) \mathbf{v}] [\partial_{\boldsymbol{\theta}} g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}) \mathbf{v}]^\top. \quad (20)$$

By Taylor series expansion, we have $\partial_{\boldsymbol{\theta}} g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta}) \mathbf{v} \approx (g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta} + \epsilon \mathbf{v}) - g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta})) / \epsilon$ with ϵ as a small scalar (e.g., 10^{-5}). Combining this with MC estimation, we get

$$\kappa_{\text{NTK}}^{\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}} \approx \frac{1}{S} \sum_{s=1}^S \left[\frac{g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta} + \epsilon \mathbf{v}_s) - g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta})}{\epsilon} \right] \left[\frac{g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta} + \epsilon \mathbf{v}_s) - g(\mathbf{X}_{\text{tr}}, \boldsymbol{\theta})}{\epsilon} \right]^\top$$

with $\mathbf{v}_s \sim p(\mathbf{v}), s \in [S]$. Namely, we perform forward propagation for $S+1$ times with the NN g to approximately estimate the NTK training kernel matrix. Popular examples of $p(\mathbf{v})$ fulfilling the above requirement include the multivariate standard normal $\mathcal{N}(0, \mathbf{I}_{\dim(\boldsymbol{\theta})})$ and the multivariate Rademacher distribution (the uniform distribution over $\{\pm 1\}^{\dim(\boldsymbol{\theta})}$). We use the latter in our experiments.

Discussion One may question that now that we can approximate the NN-GP kernels and NTKs via these *random feature* strategies, why do we still need to train a NeuralEF for kernel approximation? We clarify that the sample size S for these strategies has to be large enough to realize high-fidelity approximation (e.g., > 1000), which is unaffordable when testing. In contrast, NeuralEF captures the kernel by only several top eigenfunctions, which prominently accelerates the testing (see Section 4.2.2 for empirical evidence).

4. Experiments

We apply NeuralEF to a handful of kernel approximation scenarios to demonstrate its promise. We set batch size B as 256 and optimize with an Adam (Kingma & Ba, 2015) optimizer with 10^{-3} learning rate unless otherwise stated. The results of SpIN are obtained based on its official code⁵.

⁴ $\dim(\boldsymbol{\theta})$ denotes the dimensionality of $\boldsymbol{\theta}$.

⁵https://github.com/deepmind/spectral_inference_networks.




Figure 3: Project 2-D data to the 3-D space by the approximate top-3 eigenfunctions of the MLP-GP kernels. We use a 3-layer MLP with ReLU activations and a 3-layer MLP with Erf activations to specify the MLP-GP kernels for “two-moon” and “circles” respectively.

The detailed experiment settings are in Appendix B.

4.1. Deconstructing Classic Kernels

At first, we experiment on classic kernels. The baselines include SpIN and the Nyström method. The latter converges to the ground truth solutions as the number of samples increases. We set $q(x)$ as uniform distributions over $[-1, 1]$ and $[-2, 2]$ for the polynomial kernel $\kappa(x, x') = (x^\top x' + 1.5)^4$ and the RBF kernel $\kappa(x, x') = \exp(-\|x - x'\|^2/2)$, respectively. We use $k = 10$ multilayer perceptrons (MLPs) of width 32 to instantiate $\hat{\psi}_j$. We set batch size B as the dataset size N if $N < 256$ else 256, and train for 2000 iterations.

Figure 2 and Appendix C.1 display the recovered eigenfunctions under various scales of training data. As shown, NeuralEF delivers as good solutions as the Nyström method. Yet, as more training samples involved, *the training overhead of NeuralEF grows gracefully while the Nyström method suffers from inefficiency issues*. The results also corroborate the inefficiency issues of SpIN.

4.2. Deconstructing Modern Kernels

We then use NeuralEF to deconstruct the NN-GP kernels and the NTKs specified by NNs with a scalar output.

4.2.1. THE NN-GP KERNELS

The MLP-GP kernels We start from a simple problem of processing 1000 2-D data from the “two-moon” or “circles” dataset. We consider the NN-GP kernel associated with a 3-layer MLP architecture with ReLU or Erf activation, referred to as MLP-GP kernel for short. We utilize the strategies in Section 3.3 to comfortably train NeuralEF. We are interested in the top-3 eigenfunctions for data visualization. We instantiate the neural eigenfunctions as 3-layer MLPs of width 32. The optimization settings are matched with those in Section 4.1. We train SpIN models under the same settings for a fair comparison. After training, we use the approximate eigenfunctions to project the data into the 3-D space for visualization. As shown in Figure 3, NeuralEF yields more appealing outcomes than SpIN.

Table 1: The accuracy of Logistic regression classifiers trained on the projections of MNIST images given by the top-10 eigenfunctions for various kernels. Note that Nyström method does not apply to our CNN-GP kernel because no analytical form of the kernel is known when max-pooling layers are used. We have tuned the hyper-parameters for the polynomial and RBF kernels extensively but have not tuned those for the CNN-GP kernel.

Method	LR test accuracy
<i>Our method (CNN-GP kernel)</i>	84.98%
<i>Nyström (CNN-GP kernel)</i>	N/A
<i>Nyström (polynomial kernel)</i>	78.00%
<i>Nyström (RBF kernel)</i>	77.55%

The CNN-GP kernels We then consider using the NN-GP kernels specified by convolutional neural networks (CNNs) (dubbed as CNN-GP kernels) to process MNIST images. Without loss of generality, we use the following architecture Conv5-ReLU-MaxPool2-Conv5-ReLU-MaxPool2-Linear-ReLU-Linear with Conv5 being a 5×5 convolution and MaxPool2 being a 2×2 max-pooling. Note that the max-pooling operations make the analytical kernel evaluation highly nontrivial and no closed-form solution is known (Novak et al., 2019), thereby preventing the use of Nyström method. Yet, NeuralEF is applicable as it does not hinge on kernelized solutions and the kernel matrix on training data can be efficiently estimated. We use 10 CNNs also with the aforementioned CNN architecture to set up $\{\hat{\psi}_j\}_{j=1}^{10}$, and train them on MNIST training images for 20000 iterations. After that, we use them to project all MNIST images to the 10-D space. We report the test accuracy of the (multiclass) Logistic regression (LR) model trained on these projections in Table 1. We compare the results with those obtained by polynomial kernel and RBF kernel using Nyström method. We exclude SpIN from the comparison due to nontrivial implementation challenges involved in the application of SpIN to the CNN-GP kernel.

Table 1 shows that the CNN-GP kernels are superior to the classic kernels for processing images, and our approach can easily unleash their potential. This also demonstrates the efficacy of NeuralEF for dealing with *out-of-sample extension*. We plot the top-3 dimensions of the projections




Figure 4: NeuralEF v.s. random features for approximating the NTK of a ResNet-20 binary classifier. The evaluation is performed on 128 CIFAR-10 test images. We explicitly compute the Jacobian matrix to attain the ground truth kernel matrix. Testing a datum with NeuralEF entails $k = 10$ NNs forward passes. We use $S = 10$ MC samples in the random feature approach such that the two methods have similar overhead.




Figure 5: A plot of the eigenvalue against the index of the eigenvalue for the NTK corresponding to a ResNet-20 binary classifier. Our method can approximate the top-10 ground truth eigenvalues with small errors.

given by NeuralEF in Appendix C.2. We see the projections form class-conditional clusters, implying that NeuralEF can learn the discriminative structures in CNN-GP kernels.

4.2.2. THE NTKS

We next experiment on the empirical NTKs corresponding to *practically sized* NNs. Without loss of generality, we train a ResNet-20 classifier (He et al., 2016) to distinguish the airplane images from the automobile ones from CIFAR-10 (Krizhevsky et al., 2009), and target the NTK associated with the trained binary classifier. The MC estimation strategy in Section 3.3 is used to approximate the training kernel matrix. We train a NeuralEF to find the top-10 eigenpairs of the NTK. Since analytical evaluation of the NTK of ResNet-20 is very time-consuming, we do not include the Nyström method in this experiment.

Figure 5 displays the discovered top-10 eigenvalues and the ground truth ones, which are estimated by naively eigendecomposing the ground truth training kernel matrix. Besides, we use the found eigenfunctions to recover the kernel matrix on test data: $\kappa(\mathbf{x}, \mathbf{x}') \approx \sum_{j=1}^k \hat{\mu}_j \hat{\psi}_j(\mathbf{x}) \hat{\psi}_j(\mathbf{x}')$. The recovered matrix and the ground truth calculated from exact

Jacobian are displayed in Figure 4. We also plot the kernel matrix estimated by the random feature (MC estimation) strategy in Section 3.3 with $S = 10$ samples. We can see that: (i) NeuralEF can approximate the top-10 ground truth eigenvalues with small errors; (ii) The kernel matrix recovered by NeuralEF suffers from distortion, perhaps due to that the eigenspectrum of the NTK has a *long tail* and hence capturing only the top-10 eigenpairs is insufficient; (iii) Under identical resource consumption, the random feature approach underperforms NeuralEF in aspect of approximation quality, reflecting the importance of capturing eigenpairs; (iv) This NTK has *low-rank* structures.

Further, we perform K-means clustering on the projections yielded by NeuralEF for the test images, and get **97.5%** accuracy. As references, the K-means given the random features ($S = 10$) yields 81.3% accuracy, and the K-means given the raw pixels delivers 67.8% accuracy. This substantiates that *NTK can characterize the intrinsic similarities between data points* and NeuralEF can approximate NTK in a more efficient way than random features.

4.3. Advanced Application: Scale up Linearised Laplace Approximation

Laplace approximation (LA) (Mackay, 1992) is a canonical approach for approximate posterior inference, yet suffers from underfitting (Lawrence, 2001). Foong et al. (2019) proposed to linearise the output of the model about the maximum a posteriori (MAP) solution to alleviate this issue (see also (Khan et al., 2019)). Specifically, we consider learning an NN model $g(\cdot, \boldsymbol{\theta}) : \mathcal{X} \rightarrow N_{\text{out}}$ for data $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ under an isotropic Gaussian prior $p(\boldsymbol{\theta}) = \mathcal{N}(0, \sigma_0^2 \mathbf{I}_{\text{dim}(\boldsymbol{\theta})})$. The linearised Laplace approximation (LLA) first finds the maximum a posteriori (MAP) solution $\boldsymbol{\theta}_{\text{MAP}}$, and then uses the following Gaussian process to form a function-space approximate posterior:

$$\mathcal{GP}(f|g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}), \partial_{\boldsymbol{\theta}} g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}) \boldsymbol{\Sigma} \partial_{\boldsymbol{\theta}} g(\mathbf{x}', \boldsymbol{\theta}_{\text{MAP}})^{\top}), \quad (21)$$




Figure 6: The histograms of the uncertainty estimates for in-distribution CIFAR-10 test data and OOD SVHN test data. We experiment on ResNet-110 here. AUPR refers to the area under the precision-recall curve.

where Σ is the inversion of the Gauss-Newton matrix⁶: $\Sigma^{-1} = \sum_i \partial_{\theta} g(\mathbf{x}_i, \boldsymbol{\theta}_{\text{MAP}})^{\top} \Lambda_i \partial_{\theta} g(\mathbf{x}_i, \boldsymbol{\theta}_{\text{MAP}}) + 1/\sigma_0^2 \mathbf{I}_{\dim(\theta)}$ with $\Lambda_i := -\partial_{\mathbf{g}\mathbf{g}}^2 \log p(\mathbf{y}_i | \mathbf{g})|_{\mathbf{g}=g(\mathbf{x}_i, \boldsymbol{\theta}_{\text{MAP}})} \in \mathbb{R}^{N_{\text{out}} \times N_{\text{out}}}$.

However, like the vanilla Laplace approximation, LLA also suffers from the time-consuming matrix inversion on the matrix of size $\dim(\boldsymbol{\theta}) \times \dim(\boldsymbol{\theta})$. We next show that NeuralEF can be incorporated to overcome this issue.

The covariance kernel in Equation (21) is closely related to the NTK $\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') = \partial_{\theta} g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}) \partial_{\theta} g(\mathbf{x}', \boldsymbol{\theta}_{\text{MAP}})^{\top}$. Empowered by this observation, we prove that the above GP can be approximated as (see Appendix A.3 for the proof):

$$\mathcal{GP} \left(f|g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}), \tilde{\Psi}(\mathbf{x}) \left[\sum_i \tilde{\Psi}(\mathbf{x}_i)^{\top} \Lambda_i \tilde{\Psi}(\mathbf{x}_i) + \frac{1}{\sigma_0^2} \mathbf{I}_k \right]^{-1} \tilde{\Psi}(\mathbf{x}')^{\top} \right), \quad (22)$$

where $\tilde{\Psi}(\mathbf{x}) := [\sqrt{\hat{\mu}_1} \hat{\psi}_1(\mathbf{x}), \dots, \sqrt{\hat{\mu}_k} \hat{\psi}_k(\mathbf{x})] \in \mathbb{R}^{N_{\text{out}} \times k}$ with $\hat{\psi}_i: \mathcal{X} \rightarrow \mathbb{R}^{N_{\text{out}}}$ as the approximate multi-output eigenfunction (see Section 3.2) corresponding to the approximate i -th largest eigenvalues $\hat{\mu}_i$ of κ_{NTK} . With this, we only need to invert a matrix of size $k \times k$ to estimate the covariance.

Then, the whole pipeline for the refined LLA is: we first find $\boldsymbol{\theta}_{\text{MAP}}$, then use NeuralEF to find the top- k (we set $k = 10$ in the experiments) eigenpairs of the kernel $\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') = \partial_{\theta} g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}) \partial_{\theta} g(\mathbf{x}', \boldsymbol{\theta}_{\text{MAP}})^{\top}$, then iterate through the training set to compute $\sum_i \tilde{\Psi}(\mathbf{x}_i)^{\top} \Lambda_i \tilde{\Psi}(\mathbf{x}_i)$, and finally obtain the approximate posterior in Equation (22).

We verify on CIFAR-10 using ResNet architectures where $\dim(\boldsymbol{\theta}) > 10^5$. Naive LLA fails due to scalability issues, so we take MAP, Kronecker factored LLA (KFAC LLA), diagonal LLA (Diag LLA), and last-layer LLA as baselines. We implement the last three baselines based on the `laplace` library⁷ (Daxberger et al., 2021). The training kernel matrix for NeuralEF is of size $5 \cdot 10^5 \times 5 \cdot 10^5$, but the training only lasts for half a day. As these methods exhibit similar test accuracy, we only report the comparison on negative log-likelihood (NLL) and expected calibration error (ECE) (Guo et al., 2017) in Table 2. We also depict the histograms of the uncertainty estimates (measured by predictive entropy)

⁶The Gauss-Newton matrix serves as a workaround of the Hessian since that the Hessian is more difficult to estimate.

⁷<https://github.com/AlexImmer/Laplace>.

Table 2: Comparison on test NLL \downarrow and ECE \downarrow on CIFAR-10.

Method	ResNet-20		ResNet-56		ResNet-110	
	NLL	ECE	NLL	ECE	NLL	ECE
<i>Ours</i>	0.277	0.016	0.234	0.012	0.241	0.010
<i>MAP</i>	0.357	0.049	0.336	0.050	0.345	0.046
<i>KFAC LLA</i>	0.906	0.468	1.576	0.707	1.767	0.749
<i>Diag LLA</i>	0.934	0.480	1.606	0.712	1.797	0.754
<i>Last-layer LLA</i>	0.264	0.026	0.231	0.024	0.233	0.019

for in-distribution and out-of-distribution (OOD) data in Figure 6. As shown, LLA with NeuralEF is consistently better than the baselines in the aspect of model calibration. The NLLs and predictive uncertainty of LLA with NeuralEF are also better than or on par with the competitors.

See Appendix C.3 for one more application of NeuralEF where we use it to approximate the implicit kernel induced by the stochastic gradient descent (SGD) trajectory to perform Bayesian deep learning (BDL).

5. Related Work

Recently, there is ongoing effort to associate specific (Bayesian) NNs with kernel methods or Gaussian processes to gain insights for the theoretical understanding of NNs (Neal, 1996; Lee et al., 2017; Garriga-Alonso et al., 2018; Matthews et al., 2018; Novak et al., 2018; Jacot et al., 2018; Arora et al., 2019; Khan et al., 2019; Sun et al., 2020) or to enrich the family of kernels (Wilson et al., 2016). However, it has been rarely explored how to solve the equally important ‘‘inverse’’ problem—designing appropriate NN counterparts for the kernels of interest. We show in this work that approximating kernels with NNs can be the key to scaling up kernel methods to large data.

The Nyström method (Nyström, 1930; Williams & Seeger, 2001) is a classic kernel approximation method, and has been extended to enable the out-of-sample extension of spectral embedding methods by Bengio et al. (2004). But as discussed, the Nyström method faces inefficiency issues when handling big data and modern kernels like NTKs. Instead, deconstructing kernels by NNs has the potential to ameliorate these pathologies due to the expressiveness and

scalability of NNs. SpIN is the first work in this spirit (Pfau et al., 2018). It trains NNs to approximate the top eigenfunctions of the kernel for kernel approximation, but it suffers from an ill-posed objective function and thereby an involved learning procedure. It is difficult to extend SpIN to treat modern kernels and big data due to the requirement of Cholesky decomposition and manipulation of Jacobians. Relatedly, EigenGame (Gemp et al., 2020) identifies a common mistake in literature for interpreting PCA as an optimization problem and proposes ways to fix it. It turns out that the same spirit also applies to fixing the SpIN objective function. In fact, NeuralEF can be viewed as a function-space extension of EigenGame.

6. Conclusion

We propose NeuralEF for scalable kernel approximation in this paper. During the derivation of the method, we have deepened the connections between kernels and NNs. We show the efficacy of NeuralEF and further apply it in several interesting yet challenging scenarios in unsupervised and supervised learning. We discuss the limitations and possible future works of NeuralEF below.

Limitations Currently, we represent each eigenfunction with a dedicated NN, so we train k NNs to cover the top- k eigenpairs. This may become costly when we have to use a large k (i.e., the eigenspectrum is long-tail). Besides, we empirically observe that NeuralEF has difficulties capturing the eigenpairs with relatively small eigenvalues (e.g., 1% of the largest eigenvalue) perhaps due to issues in stochastic optimization or numerical errors. Finally, It is difficult to tune the hyper-parameters of the kernel while using NeuralEF approximations.

Future work To promote parameter efficiency and potentially support a very large k , we need to perform *weight-sharing* among the k neural eigenfunctions. More appealing applications such as using NeuralEF as unsupervised representation learners also deserve future investigation.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (Nos. 2020AAA0104304, 2017YFA0700904), NSFC Projects ((Nos. 62061136001, 62106122, 62076147, U19B2034, U1811461, U19A2081)), Tsinghua-Huawei Joint Research Program, a grant from Tsinghua Institute for Guo Qiang, and the High Performance Computing Center, Tsinghua University.

References

Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R., and Wang, R. On exact computation with an infinitely wide

neural net. *arXiv preprint arXiv:1904.11955*, 2019.

Bengio, Y., Delalleau, O., Roux, N. L., Paiement, J.-F., Vincent, P., and Ouimet, M. Learning eigenfunctions links spectral embedding and kernel PCA. *Neural Computation*, 16(10):2197–2219, 2004.

Bengio, Y., Delalleau, O., and Roux, N. The curse of highly variable functions for local kernel machines. *Advances in Neural Information Processing Systems*, 18, 2005.

Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., and Hennig, P. Laplace redux-effortless Bayesian deep learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Du, S. S., Hou, K., Salakhutdinov, R. R., Póczos, B., Wang, R., and Xu, K. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *Advances in Neural Information Processing Systems*, 32:5723–5733, 2019.

Foong, A. Y., Li, Y., Hernández-Lobato, J. M., and Turner, R. E. ‘in-between’ uncertainty in bayesian neural networks. *arXiv preprint arXiv:1906.11537*, 2019.

Francis, D. P. and Raimond, K. Major advancements in kernel function approximation. *Artificial Intelligence Review*, 54(2):843–876, 2021.

Garriga-Alonso, A., Rasmussen, C. E., and Aitchison, L. Deep convolutional networks as shallow gaussian processes. *arXiv preprint arXiv:1808.05587*, 2018.

Gemp, I., McWilliams, B., Vernade, C., and Graepel, T. Eigengame: Pca as a nash equilibrium. *arXiv preprint arXiv:2010.00554*, 2020.

Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. *arXiv preprint arXiv:1706.04599*, 2017.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. PMLR, 2015.

Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., and Wilson, A. G. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018.

- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.
- Khan, M. E. E., Immer, A., Abedi, E., and Korzepa, M. Approximate inference turns deep networks into Gaussian processes. *Advances in Neural Information Processing Systems*, 32, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- Lawrence, N. D. *Variational inference in probabilistic models*. PhD thesis, Citeseer, 2001.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- Mackay, D. J. C. *Bayesian methods for adaptive models*. PhD thesis, California Institute of Technology, 1992.
- Maddox, W. J., Izmailov, P., Garipov, T., Vetrov, D. P., and Wilson, A. G. A simple baseline for bayesian uncertainty in deep learning. In *Advances in Neural Information Processing Systems*, pp. 13153–13164, 2019.
- Mandt, S., Hoffman, M. D., and Blei, D. M. Stochastic gradient descent as approximate Bayesian inference. *Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- Matthews, A. G. d. G., Rowland, M., Hron, J., Turner, R. E., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018.
- Muller, K.-R., Mika, S., Ratsch, G., Tsuda, K., and Scholkopf, B. An introduction to kernel-based learning algorithms. *IEEE transactions on neural networks*, 12(2):181–201, 2001.
- Munkhoeva, M., Kapushev, Y., Burnaev, E., and Oseledets, I. Quadrature-based features for kernel approximation. *Advances in neural information processing systems*, 31, 2018.
- Neal, R. M. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pp. 29–53. Springer, 1996.
- Novak, R., Xiao, L., Lee, J., Bahri, Y., Yang, G., Hron, J., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Bayesian deep convolutional networks with many channels are gaussian processes. *arXiv preprint arXiv:1810.05148*, 2018.
- Novak, R., Xiao, L., Hron, J., Lee, J., Alemi, A. A., Sohl-Dickstein, J., and Schoenholz, S. S. Neural tangents: Fast and easy infinite neural networks in python. *arXiv preprint arXiv:1912.02803*, 2019.
- Nyström, E. J. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. *Acta Mathematica*, 54:185–204, 1930.
- Pfau, D., Petersen, S., Agarwal, A., Barrett, D. G., and Stachenfeld, K. L. Spectral inference networks: Unifying deep and spectral learning. In *International Conference on Learning Representations*, 2018.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- Rahimi, A. and Recht, B. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. *Advances in neural information processing systems*, 21, 2008.
- Schölkopf, B., Smola, A. J., Bach, F., et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Sun, S., Shi, J., and Grosse, R. B. Neural networks as inter-domain inducing points. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2020.
- Williams, C. and Seeger, M. Using the nyström method to speed up kernel machines. In *Proceedings of the 14th annual conference on neural information processing systems*, number CONF, pp. 682–688, 2001.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. In *Artificial intelligence and statistics*, pp. 370–378. PMLR, 2016.
- Woodbury, M. A. *Inverting modified matrices*. Statistical Research Group, 1950.
- Yu, F. X. X., Suresh, A. T., Choromanski, K. M., Holtmann-Rice, D. N., and Kumar, S. Orthogonal random features. *Advances in neural information processing systems*, 29, 2016.

A. Proof

A.1. Proof of Theorem 1

Theorem 1. *The eigenpairs of the kernel $\kappa(\mathbf{x}, \mathbf{x}')$ can be recovered by simultaneously solving the following asymmetric maximization problems:*

$$\max_{\hat{\psi}_j} R_{jj} \quad \text{s.t.: } C_j = 1, R_{ij} = 0, \forall j \geq 1, i \in [j-1], \quad (11)$$

where $\hat{\psi}_j \in L^2(\mathcal{X}, q)$ represent the introduced approximate eigenfunctions, and

$$R_{ij} := \iint \hat{\psi}_i(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') \hat{\psi}_j(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x}, \quad (12)$$

$$C_j := \int \hat{\psi}_j(\mathbf{x}) \hat{\psi}_j(\mathbf{x}) q(\mathbf{x}) d\mathbf{x}. \quad (13)$$

$(R_{jj}, \hat{\psi}_j)$ will converge to the eigenpair associated with j -th largest eigenvalue of κ .

Proof. First of all, it is easy to see that the functions in $L^2(\mathcal{X}, q)$ form an inner product space with inner product given by:

$$\langle \varphi, \varphi' \rangle = \int \varphi(\mathbf{x}) \varphi'(\mathbf{x}) q(\mathbf{x}) d\mathbf{x}, \quad \forall \varphi, \varphi' \in L^2(\mathcal{X}, q).$$

By Mercer's theorem, we have $\kappa(\mathbf{x}, \mathbf{x}') = \sum_{j \geq 1} \mu_j \psi_j(\mathbf{x}) \psi_j(\mathbf{x}')$ where (μ_j, ψ_j) refers to the j -th ground truth eigenpair of κ . Without loss of generality, we assume $\forall i < j: \mu_i > \mu_j$.

It is easy to reformulate the problems in Equation (11) as:

$$\begin{aligned} \max_{\hat{\psi}_1} R_{11} \quad \text{s.t.: } C_1 &= 1 \\ \max_{\hat{\psi}_2} R_{22} \quad \text{s.t.: } C_2 &= 1, R_{12} = 0 \\ \max_{\hat{\psi}_3} R_{33} \quad \text{s.t.: } C_3 &= 1, R_{13} = 0, R_{23} = 0 \\ &\dots \end{aligned}$$

Interestingly, when simultaneously solving these problems, the solution to the j -th problem only depends on the those to the preceding problems and is independent of those to the following problems. Therefore, we can derive the solutions to the problems in a sequential manner.

Specifically, we first consider the maximization objective in the first problem:

$$\begin{aligned} R_{11} &= \iint \hat{\psi}_1(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') \hat{\psi}_1(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x} \\ &= \iint \hat{\psi}_1(\mathbf{x}) \left(\sum_{j \geq 1} \mu_j \psi_j(\mathbf{x}) \psi_j(\mathbf{x}') \right) \hat{\psi}_1(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x} \\ &= \sum_{j \geq 1} \mu_j \iint \hat{\psi}_1(\mathbf{x}) \psi_j(\mathbf{x}) \psi_j(\mathbf{x}') \hat{\psi}_1(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x} \\ &= \sum_{j \geq 1} \mu_j \int \hat{\psi}_1(\mathbf{x}) \psi_j(\mathbf{x}) q(\mathbf{x}) d\mathbf{x} \int \hat{\psi}_1(\mathbf{x}') \psi_j(\mathbf{x}') q(\mathbf{x}') d\mathbf{x}' \\ &= \sum_{j \geq 1} \mu_j \langle \hat{\psi}_1, \psi_j \rangle^2. \end{aligned}$$

Given the definition that the eigenfunctions $\{\psi_j\}_{j \geq 1}$ are orthonormal (see Equation (6)), we know they form a set of orthonormal bases of the $L^2(\mathcal{X}, q)$ space. Thus, we can represent $\hat{\psi}_1$ in such a new axis system by coordinate (w_1, w_2, \dots) :

$$\hat{\psi}_1 = \sum_{i \geq 1} w_i \psi_i.$$

We can then rewrite the maximization objective as:

$$R_{11} = \sum_{j \geq 1} \mu_j \langle \hat{\psi}_1, \psi_j \rangle^2 = \sum_{j \geq 1} \mu_j \langle \sum_{i \geq 1} w_i \psi_i, \psi_j \rangle^2 = \sum_{j \geq 1} \mu_j w_j^2.$$

Recalling the constraint $C_1 = 1$, we have

$$\langle \hat{\psi}_1, \hat{\psi}_1 \rangle = \langle \sum_{i \geq 1} w_i \psi_i, \sum_{j \geq 1} w_j \psi_j \rangle = \sum_{i, j \geq 1} w_i w_j \langle \psi_i, \psi_j \rangle = \sum_{j \geq 1} w_j^2 = 1.$$

Then, it is straight-forward to see the maximum value of R_{11} is the largest ground truth eigenvalue μ_1 . The condition to make the maximization hold is that (w_1, w_2, \dots) is a one-hot vector with the first element as 1, namely, $\hat{\psi}_1 = \psi_1$. Thereby, we prove that solving the first problem uncovers the first eigenvalue as well as the associated eigenfunction of the kernel κ .

We then consider solving the second problem given $\hat{\psi}_1 = \psi_1$. Compared to the first problem, there is one more constraint:

$$\begin{aligned} R_{12} &= 0 \\ \Rightarrow \iint \hat{\psi}_1(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') \hat{\psi}_2(\mathbf{x}') q(\mathbf{x}') q(\mathbf{x}) d\mathbf{x}' d\mathbf{x} &= 0 \\ \Rightarrow \int \hat{\psi}_2(\mathbf{x}') q(\mathbf{x}') \int \hat{\psi}_1(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') q(\mathbf{x}) d\mathbf{x} d\mathbf{x}' &= 0 \\ \Rightarrow \int \hat{\psi}_2(\mathbf{x}') q(\mathbf{x}') \int \psi_1(\mathbf{x}) \kappa(\mathbf{x}, \mathbf{x}') q(\mathbf{x}) d\mathbf{x} d\mathbf{x}' &= 0 \\ \Rightarrow \int \hat{\psi}_2(\mathbf{x}') q(\mathbf{x}') \mu_1 \psi_1(\mathbf{x}') d\mathbf{x}' &= 0 \\ \Rightarrow \langle \psi_1, \hat{\psi}_2 \rangle &= 0. \end{aligned}$$

Namely, $\hat{\psi}_2$ is constrained in the orthogonal complement of the subspace spanned by ψ_1 . Given such a minor difference between the second problem and the first problem, we can apply an analysis similar to that for the first problem to solve the second problem. Note that μ_2 is the largest eigenvalue in the orthogonal complement of the subspace spanned by ψ_1 , $(R_{22}, \hat{\psi}_2)$ would hence converge to the 2-th largest eigenvalue and the associated eigenfunction of κ .

Applying this procedure incrementally to the additional problems then finishes the proof. \square

A.2. Justification of NeuralEF in Practice

Though, ideally, $\hat{\psi}_j$ would converge to the ground truth eigenfunction ψ_j by Theorem 1, in practice, when performing optimization according to Equation (14) we cannot guarantee that at the convergence $\hat{\psi}_j$ is exactly equivalent to ψ_j . May the small deviation in the solutions to the preceding problems significantly bias the solutions to the following problems? We show that it is not the case below.

To simplify the analysis, we consider only the first two problems in Equation (14), namely, learning the first two approximate eigenfunctions $\hat{\psi}_1$ and $\hat{\psi}_2$. As done in the proof of Theorem 1, we represent $\hat{\psi}_1$ and $\hat{\psi}_2$ in the axis system specified by $\{\psi_j\}_{j \geq 1}$:

$$\hat{\psi}_1 = \sum_{i \geq 1} v_i \psi_i \quad \hat{\psi}_2 = \sum_{i \geq 1} w_i \psi_i.$$

At first, we assume that solving $\max_{\hat{\psi}_1} R_{11}$ s.t.: $C_1 = 1$, whose optima is ψ_1 , leads to that $\|\hat{\psi}_1 - \psi_1\|^2 < 2\epsilon_1 < 2$, where $\|\cdot\|$ refers to the norm induced by the inner product in $L^2(\mathcal{X}, q)$. Then,

$$\|\hat{\psi}_1 - \psi_1\|^2 < 2\epsilon_1 \Rightarrow 2 - 2\langle \hat{\psi}_1, \psi_1 \rangle < 2\epsilon_1 \Rightarrow 2 - 2v_1 < 2\epsilon_1 \Rightarrow v_1 > 1 - \epsilon_1.$$

As $\hat{\psi}_1$ is normalized, we have:

$$v_1^2 + v_2^2 \leq 1 \Rightarrow v_2^2 \leq 1 - v_1^2 \leq 1 - (1 - \epsilon_1)^2 = 2\epsilon_1 - \epsilon_1^2 \Rightarrow |v_2| < \sqrt{2\epsilon_1 - \epsilon_1^2}.$$

We then assume that solving $\max_{\hat{\psi}_2} -\frac{R_{12}^2}{R_{11}}$ results in that $|R_{12}| < \epsilon_2 < 1$. Given that

$$\begin{aligned}
 R_{12} &= \iint \hat{\psi}_1(\mathbf{x})\kappa(\mathbf{x}, \mathbf{x}')\hat{\psi}_2(\mathbf{x}')q(\mathbf{x}')q(\mathbf{x})d\mathbf{x}'d\mathbf{x} \\
 &= \iint \left(\sum_{i \geq 1} v_i \psi_i(\mathbf{x}) \right) \left(\sum_{j \geq 1} \mu_j \psi_j(\mathbf{x}) \psi_j(\mathbf{x}') \right) \hat{\psi}_2(\mathbf{x}')q(\mathbf{x}')q(\mathbf{x})d\mathbf{x}'d\mathbf{x} \\
 &= \sum_{i \geq 1} v_i \sum_{j \geq 1} \mu_j \int \psi_i(\mathbf{x})\psi_j(\mathbf{x})q(\mathbf{x})d\mathbf{x} \int \psi_j(\mathbf{x}')\hat{\psi}_2(\mathbf{x}')q(\mathbf{x}')d\mathbf{x}' \\
 &= \sum_{i \geq 1} v_i \sum_{j \geq 1} \mu_j \mathbb{1}[i = j] \langle \psi_j, \hat{\psi}_2 \rangle \\
 &= \sum_{i \geq 1} v_i \mu_i \langle \psi_i, \hat{\psi}_2 \rangle \\
 &= \sum_{i \geq 1} v_i \mu_i w_i,
 \end{aligned}$$

we have $|\sum_{i \geq 1} v_i \mu_i w_i| < \epsilon_2$.

We consider the optima $\hat{\psi}_2^*$ of the problem $\max_{\hat{\psi}_2} R_{22} - \frac{R_{12}^2}{R_{11}}$ s.t.: $C_2 = 1$, which equals to

$$\max_{\{w_j\}_{j \geq 1}} \sum_{j \geq 1} \mu_j w_j^2 \text{ s.t.: } \left| \sum_{i \geq 1} v_i \mu_i w_i \right| < \epsilon_2, \sum_{j \geq 1} w_j^2 = 1.$$

It is easy to see at the maximum, $w_i^2 = 0, \forall i > 2$. Namely, $\hat{\psi}_2^* = w_1 \psi_1 + w_2 \psi_2$ with $w_1^2 + w_2^2 = 1$.

Without loss of generality, we assume $w_2 > 0$ and measure the distance between $\hat{\psi}_2^*$ and ψ_2 to estimate the induced bias:⁸

$$\|\hat{\psi}_2^* - \psi_2\|^2 = 2 - 2\langle \hat{\psi}_2^*, \psi_2 \rangle = 2 - 2w_2 = 2 - 2\sqrt{1 - w_1^2} \leq 2 - 2(1 - w_1^2) = 2w_1^2.$$

Recall that $|v_1 \mu_1 w_1 + v_2 \mu_2 w_2| < \epsilon_2, v_1 > 1 - \epsilon_1$, and $|v_2| < \sqrt{2\epsilon_1 - \epsilon_1^2}$, then

1) when $w_1 > 0$:

$$(1 - \epsilon_1)\mu_1 w_1 < v_1 \mu_1 w_1 < \epsilon_2 - v_2 \mu_2 w_2 \leq \epsilon_2 + \sqrt{2\epsilon_1 - \epsilon_1^2} \mu_2,$$

then $w_1 < \frac{1}{(1 - \epsilon_1)\mu_1} (\epsilon_2 + \sqrt{2\epsilon_1 - \epsilon_1^2} \mu_2)$;

2) when $w_1 < 0$:

$$(1 - \epsilon_1)\mu_1 w_1 > v_1 \mu_1 w_1 > -\epsilon_2 - v_2 \mu_2 w_2 \geq -\epsilon_2 - \sqrt{2\epsilon_1 - \epsilon_1^2} \mu_2,$$

then $w_1 > \frac{1}{(1 - \epsilon_1)\mu_1} (-\epsilon_2 - \sqrt{2\epsilon_1 - \epsilon_1^2} \mu_2)$.

Therefore, $|w_1|$ is small if $\|\hat{\psi}_1 - \psi_1\|^2$ and $|R_{12}|$ are small, and in turn $\|\hat{\psi}_2^* - \psi_2\|^2 = 2w_1^2$ is small.

Applying this procedure incrementally to the additional problems then finishes the whole justification.

A.3. Proof of Equation (22)

Proof. We denote $\partial_{\theta} g(\mathbf{x}, \theta_{\text{MAP}})$ as $\mathbf{J}_{\mathbf{x}}^*$ for compactness. We concatenate $\{\mathbf{J}_{\mathbf{x}_i}^* \in \mathbb{R}^{N_{\text{out}} \times \dim(\theta)}\}_{i=1}^N$ as a big matrix $\mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \in \mathbb{R}^{N N_{\text{out}} \times \dim(\theta)}$, and organize $\{\mathbf{A}_i \in \mathbb{R}^{N_{\text{out}} \times N_{\text{out}}}\}_{i=1}^N$ as a block-diagonal matrix $\mathbf{A}_{\mathbf{X}_{\text{tr}}} \in \mathbb{R}^{N N_{\text{out}} \times N N_{\text{out}}}$. Then, by Woodbury

⁸If $w_2 < 0$, we measure the distance between $\hat{\psi}_2^*$ and $-\psi_2$, and the consequence is the same.

matrix identity (Woodbury, 1950), Σ can be rephrased:

$$\begin{aligned}\Sigma &= \left[\sum_i \mathbf{J}_{\mathbf{x}_i}^* \mathbf{A}_i \mathbf{J}_{\mathbf{x}_i}^* + 1/\sigma_0^2 \mathbf{I}_{\dim(\theta)} \right]^{-1} \\ &= \left[\mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \mathbf{A}_{\mathbf{X}_{\text{tr}}} \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* + 1/\sigma_0^2 \mathbf{I}_{\dim(\theta)} \right]^{-1} \\ &= \sigma_0^2 \left(\mathbf{I}_{\dim(\theta)} - \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \mathbf{A}_{\mathbf{X}_{\text{tr}}} \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \right)^{-1}.\end{aligned}$$

Consequently, the covariance in Equation (21) becomes:

$$\begin{aligned}& \mathbf{J}_{\mathbf{x}}^* \Sigma \mathbf{J}_{\mathbf{x}'}^* \\ &= \sigma_0^2 \mathbf{J}_{\mathbf{x}}^* \left(\mathbf{I}_{\dim(\theta)} - \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \mathbf{A}_{\mathbf{X}_{\text{tr}}} \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \right)^{-1} \mathbf{J}_{\mathbf{x}'}^* \\ &= \sigma_0^2 \left(\mathbf{J}_{\mathbf{x}}^* \mathbf{J}_{\mathbf{x}'}^* - \mathbf{J}_{\mathbf{x}}^* \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \mathbf{A}_{\mathbf{X}_{\text{tr}}} \mathbf{J}_{\mathbf{X}_{\text{tr}}}^* \mathbf{J}_{\mathbf{x}'}^* \right) \\ &= \sigma_0^2 \left(\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') - \kappa_{\text{NTK}}(\mathbf{x}, \mathbf{X}_{\text{tr}}) \left[1/\sigma_0^2 \mathbf{A}_{\mathbf{X}_{\text{tr}}}^{-1} + \kappa_{\text{NTK}}(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}) \right]^{-1} \kappa_{\text{NTK}}(\mathbf{X}_{\text{tr}}, \mathbf{x}') \right) \\ & \quad \left(\kappa_{\text{NTK}}(\mathbf{x}, \mathbf{x}') := \partial_{\theta} g(\mathbf{x}, \boldsymbol{\theta}_{\text{MAP}}) \partial_{\theta} g(\mathbf{x}', \boldsymbol{\theta}_{\text{MAP}})^{\top} = \mathbf{J}_{\mathbf{x}}^* \mathbf{J}_{\mathbf{x}'}^* \right) \\ &\approx \sigma_0^2 \left(\tilde{\Psi}(\mathbf{x}) \tilde{\Psi}(\mathbf{x}')^{\top} - \tilde{\Psi}(\mathbf{x}) \tilde{\Psi}(\mathbf{X}_{\text{tr}})^{\top} \left[1/\sigma_0^2 \mathbf{A}_{\mathbf{X}_{\text{tr}}}^{-1} + \tilde{\Psi}(\mathbf{X}_{\text{tr}}) \tilde{\Psi}(\mathbf{X}_{\text{tr}})^{\top} \right]^{-1} \tilde{\Psi}(\mathbf{X}_{\text{tr}}) \tilde{\Psi}(\mathbf{x}')^{\top} \right) \quad (\text{Mercer's theorem}) \\ &= \tilde{\Psi}(\mathbf{x}) \sigma_0^2 \left(\mathbf{I}_k - \tilde{\Psi}(\mathbf{X}_{\text{tr}})^{\top} \left[1/\sigma_0^2 \mathbf{A}_{\mathbf{X}_{\text{tr}}}^{-1} + \tilde{\Psi}(\mathbf{X}_{\text{tr}}) \tilde{\Psi}(\mathbf{X}_{\text{tr}})^{\top} \right]^{-1} \tilde{\Psi}(\mathbf{X}_{\text{tr}}) \right) \tilde{\Psi}(\mathbf{x}')^{\top} \\ &= \tilde{\Psi}(\mathbf{x}) \left[\tilde{\Psi}(\mathbf{X}_{\text{tr}})^{\top} \mathbf{A}_{\mathbf{X}_{\text{tr}}} \tilde{\Psi}(\mathbf{X}_{\text{tr}}) + 1/\sigma_0^2 \mathbf{I}_k \right]^{-1} \tilde{\Psi}(\mathbf{x}')^{\top} \quad (\text{Woodbury matrix identity}) \\ &= \tilde{\Psi}(\mathbf{x}) \left[\sum_i \tilde{\Psi}(\mathbf{x}_i) \mathbf{A}_i \tilde{\Psi}(\mathbf{x}_i) + 1/\sigma_0^2 \mathbf{I}_k \right]^{-1} \tilde{\Psi}(\mathbf{x}')^{\top}\end{aligned}$$

where $\tilde{\Psi}(\mathbf{x}) := [\sqrt{\hat{\mu}_1} \hat{\psi}_1(\mathbf{x}), \dots, \sqrt{\hat{\mu}_k} \hat{\psi}_k(\mathbf{x})] \in \mathbb{R}^{N_{\text{out}} \times k}$ with $\hat{\psi}_i : \mathcal{X} \rightarrow \mathbb{R}^{N_{\text{out}}}$ as the approximate multi-output eigenfunction corresponding to the approximate i -th largest eigenvalues $\hat{\mu}_i$ of κ_{NTK} . Note that $\tilde{\Psi}(\mathbf{X}_{\text{tr}})$ is the concatenation of $\tilde{\Psi}(\mathbf{x}_1), \dots, \tilde{\Psi}(\mathbf{x}_N)$ and hence is of size $NN_{\text{out}} \times k$.

Thus, we obtain Equation (22). \square

B. Experiment Settings

Implementation of the Nyström method To find the top- k eigenpairs, we use the `scipy.linalg.eigh` API to eigendecompose the kernel matrix $\kappa(\mathbf{X}_{\text{tr}}, \mathbf{X}_{\text{tr}}) \in \mathbb{R}^{N \times N}$ with the argument `subset_by_index` as $[N - k, N - 1]$.

Experiments on classic kernels We use 3-layer MLP to instantiate the neural eigenfunctions. The hidden size of the MLP is set as 32. We use a mixture of `Sin` and `Cos` activations, namely, one half of the hidden neurons use `Sin` activation and the other half use `Cos` activation.

Experiments on MLP-GP kernels The architecture of concern is a 3-layer MLP. To set up the MLP-GP kernels, for every linear layer, the prior on the weights is set as $\mathcal{N}(0, 2/\omega)$ with ω as the layer width, and the prior on the bias is set as $\mathcal{N}(0, 1)$. To compute the kernel matrix on the training data, we instantiate a finitely wide MLP whose width is 16, and perform MC estimation by virtue of the strategies in Section 3.3. The number of MC samples S is set as 10000.

Experiments on CNN-GP kernels To set up the CNN-GP kernels, for every convolutional/linear layer, the prior on the weights is set as $\mathcal{N}(0, 2/\text{fan.in})$, and the prior on the bias is set as $\mathcal{N}(0, 0.01)$. To compute the kernel matrix on the training data, we instantiate a finitely wide CNN whose width is 16, and perform MC estimation by virtue of the strategies in Section 3.3. The number of MC samples S is set as 2000. We instantiate $\hat{\psi}$ as the CNNs with the same architecture as the

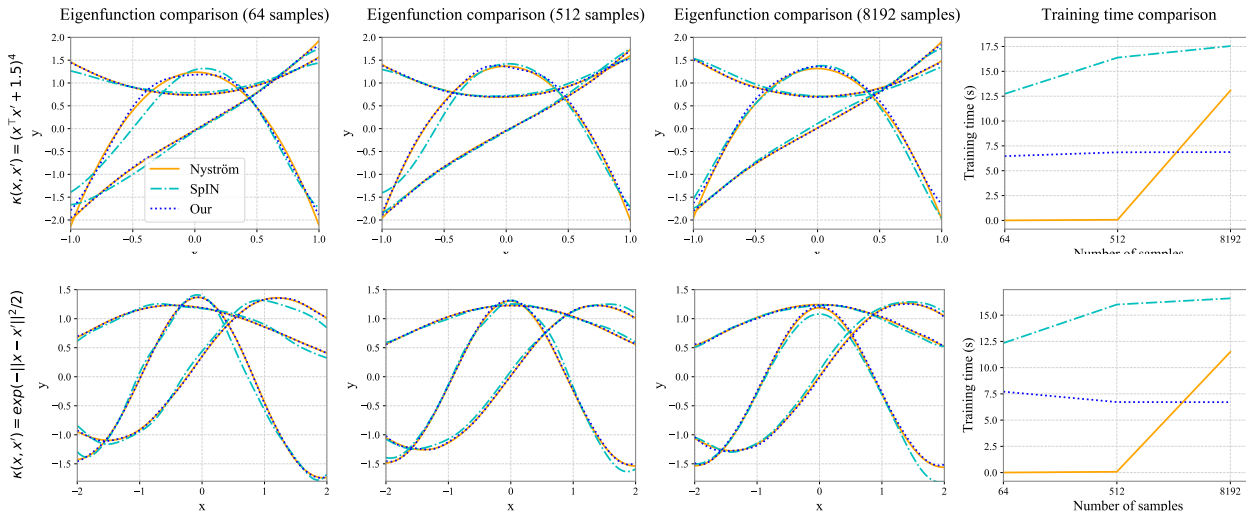


Figure 7: Estimate the eigenfunctions of the polynomial and RBF kernels with the Nyström method, SpIN, and NeuralEF (our). NeuralEF behaves as well as Nyström method but consumes *nearly constant* training time w.r.t. sample size.

concerned CNN-GP kernel but augmented with batch normalization (Ioffe & Szegedy, 2015), and set the the layers width as 32, 64, 128. We use 6000 randomly sampled MNIST training images (due to resource constraint) to perform the Nyström method. The used polynomial kernel and RBF kernel take the follows forms:

$$\begin{aligned} \kappa(x, x') &= (0.001x^T x' + 1)^{10} \\ \kappa(x, x') &= \exp(-\|x - x'\|^2/2/100), \end{aligned}$$

where the hyper-parameters are found by grid search.

Experiments on NTKs For the trained NN classifier, we fuse the BN layers into the convolutional layers to get a compact model that possesses 269, 034 parameters. To compute the kernel matrix on the training data, we perform MC estimation by virtue of the strategies in Section 3.3 and the number of MC samples S is set as 4000. We instantiate $\hat{\psi}$ as widened ResNet-20 with widening factor of 2.0.

Experiments on the linearised Laplace approximation with NeuralEF We train the CIFAR-10 classifiers with ResNet architectures for totally 150 epochs under MAP principle. The optimization settings are identical to the above ones. In particular, the weight decay is 10^{-4} , thus we can estimate the prior variance $\sigma_0^2 = \frac{1}{50000 \times 10^{-4}} = 0.2$ where 50000 is the number of training data N . After classifier training, we fuse the BN layers into the convolutional layers to get a compact model. To compute the kernel matrix on the training data, we perform MC estimation by virtue of the strategies in Section 3.3 and the number of MC samples S is set as 4000. We instantiate $\hat{\psi}$ as widened ResNet-20 with widening factor of 2.0. We learn the matrix-valued NTK kernel with the strategy provided in Section 3.2. When testing, we use 256 MC samples from the function-space posterior to empirically estimate the predictive distribution as the categorical likelihood is not conjugate of Gaussian. We scale the sampled noises by 20 given the observation that they are pretty tiny originally.

Experiments on the implicit kernel induced by SGD trajectory To collect the SGD iterates, we first train the CIFAR-10 classifiers for 150 epochs. We set the initial learning rate as 0.1, and scale the learning rate by 0.1 at 80-th and 120-th epochs. We use SGD optimizer with 0.9 momentum to train the the classifiers, with the batch size set as 128. From 150-th epoch to 155-th epoch, we linearly increase the learning rate from 0.001 to 0.05, and then keep it constant until the classifiers have been trained for 200 epochs. We totally collect $M = 50$ weight samples for SWA and SWAG (one at per epoch). We cannot use more weight samples for constructing the Gaussian covariance in SWAG due to memory constraint. However, with NeuralEF introduced for kernel approximation, we can use many weight samples for defining κ_{SGD} as eventually we save only the top eigenfunctions of the kernel instead of the kernel itself. In practice, we collect the function evaluations on the training set of around $M = 4000$ weights from the SGD trajectory, based on which we estimate the training kernel matrix. We instantiate $\hat{\psi}$ as widened ResNet-20 with widening factor of 2.0. When testing, we use 256 MC samples from the function-space posterior to empirically estimate the predictive distribution as the categorical likelihood is not conjugate of Gaussian.




Figure 8: Project MNIST test images to the 3-D space by the approximate top-3 eigenfunctions of the CNN-GP kernel found by NeuralEF. Each color represents an category.

C. More Experiments

C.1. More Results on Polynomial and RBF kernels

We provide more results of various kernel approximation methods for the aforementioned polynomial and RBF kernels in Figure 7.

C.2. Visualization of the Projections of MNIST Test Images

We plot the top-3 dimensions of the projections belonging to the MNIST test images produced by NeuralEF in Figure 8. The NeuralEF model is trained to deconstruct the CNN-GP kernel mentioned in Section 4.2.1. We see the projections form class-conditional clusters, implying that NeuralEF can learn the discriminative structures in the CNN-GP kernel.

C.3. Learn the Kernel Induced by SGD Trajectory

There has been a surge of interest in exploiting the SGD trajectory for Bayesian deep learning, attributed to the close connection between the stationary distribution of SGD iterates and the Bayesian posterior (Mandt et al., 2017). SWAG (Maddox et al., 2019) is a typical method in this line: it collects a bunch of NN weights from the SGD trajectory then constructs approximate weight-space posteriors. In fact, akin to the random feature approaches, the SGD iterates implicitly induce a kernel:

$$\kappa_{\text{SGD}}(\mathbf{x}, \mathbf{x}') = \frac{1}{M} \sum_{i=1}^M (g(\mathbf{x}, \boldsymbol{\theta}_i) - \bar{g}(\mathbf{x})) (g(\mathbf{x}', \boldsymbol{\theta}_i) - \bar{g}(\mathbf{x}'))^\top,$$

where g refers to an NN function, $\{\boldsymbol{\theta}_i\}_{i=1}^M$ are the weights SGD traverses, and $\bar{g}(\cdot) = \frac{1}{M}g(\cdot, \boldsymbol{\theta}_i)$ is the ensemble of SGD iterates. We can then define a Gaussian process $\mathcal{GP}(\bar{g}(\mathbf{x}), \kappa_{\text{SGD}}(\mathbf{x}, \mathbf{x}'))$ to approximate the true function-space posterior, which leads to the *posterior predictive*

$$p(\mathbf{x}_{\text{new}}) = \int \mathcal{GP}(f|\bar{g}(\mathbf{x}), \kappa_{\text{SGD}}(\mathbf{x}, \mathbf{x}'))p(\mathbf{x}_{\text{new}}|f)df.$$

Yet, the evaluation of both \bar{g} and κ_{SGD} on a datum \mathbf{x} entails M forward passes, resulting in the dilemma of deciding efficiency or exactness. Nonetheless, we can conjoin both by (i) taking the Stochastic Weight Averaging (SWA) (Izmailov et al., 2018), which predicts with the average weights, as a substitute for \bar{g} ⁹ and (ii) proceeding with the top eigenfunctions of κ_{SGD} instead of the kernel itself.

To verify, we experiment on the full CIFAR-10 dataset with ResNets. Following the common assumption in GP classification, we assume that the kernel correlations among output dimensions are negligible, so $\kappa^{\mathbf{X}, \mathbf{X}}$ degrades to N_{out} matrices of size $B \times B$. Viewing $\hat{\psi}_j^{\mathbf{X}}$ as N_{out} vectors of size B , we can then compute the losses Equation (17) dimension by dimension, sum them up, and invoke once backprop. With this strategy, we find the top-10 eigenpairs at per output dimension, and then use them for kernel recovery and posteriori prediction. We dub our model as SWA+NeuralEF. More experimental details

⁹Maddox et al. (2019) found that SWA is basically on par with \bar{g} in terms of performance.

NeuralEF: Deconstructing Kernels by Deep Neural Networks




Figure 9: Test accuracy \uparrow , NLL \downarrow , and ECE \downarrow comparisons among models on CIFAR-10.




Figure 10: NLL \downarrow and ECE \downarrow on CIFAR-10 corruptions for models trained with ResNet-20 architecture. The results across 19 types of skew are summarized in the boxes.

are presented in Appendix B. The baselines of concern include the vanilla models trained by SGD, SWA, and SWAG. The results on test accuracy, negative log-likelihood (NLL), and expected calibration error (ECE) (Guo et al., 2017) are given in Figure 9. We further assess the models with ResNet-20 architecture on CIFAR-10 corruptions (Hendrycks & Dietterich, 2019), a standard OOD testbed for CIFAR-10 models, with Figure 10 presenting the results.

We can see that SWA+NeuralEF is on par with or superior over SWAG across evaluation metrics, and outperforms SWA, especially in the aspect of ECE. We emphasize that another merit of SWA+NeuralEF is that, unlike SWAG, the storage cost of SWA+NeuralEF is *agnostic* to M .