Deep Neural Networks as Point Estimates for Deep Gaussian Processes

Anonymous Author(s)
Affiliation
Address
email

Abstract

Neural networks and Gaussian processes are complementary in their strengths and weaknesses. Having a better understanding of their relationship comes with the promise to make each method benefit from the strengths of the other. In this work, we establish an equivalence between the forward passes of neural networks and (deep) sparse Gaussian process models. The theory we develop is based on interpreting activation functions as interdomain inducing features through a rigorous analysis of the interplay between activation functions and kernels. This results in models that can either be seen as neural networks with improved uncertainty prediction or deep Gaussian processes with increased prediction accuracy. These claims are supported by experimental results on toy datasets, UCI benchmarks, MNIST, Fashion-MNIST and CIFAR-10.

1 Introduction

Neural networks (NNs) [1] and Gaussian processes (GPs) [2] are well-established frameworks for solving regression or classification problems, with complementary strengths and weaknesses. NNs work well when given very large datasets, and are computationally scalable enough to handle them. GPs, on the other hand, are challenging to scale to large datasets, but provide robust solutions with uncertainty estimates in low-data regimes where NNs struggle. Ideally, we would have a single model that provides the best of both approaches: the ability to handle low and high dimensional inputs, and to make robust uncertainty-aware predictions from the small to big data regimes.

Damianou and Lawrence [3] introduced the Deep Gaussian process (DGP) as a promising attempt to obtain such a model. DGPs replicate the structure of deep NNs by stacking multiple GPs as layers, with the goal of gaining the benefits of depth while retaining high quality uncertainty. Delivering this potential in practice requires an efficient and accurate approximate Bayesian training procedure, which is highly challenging to develop. Significant progress has been made in recent years, which has led to methods outperform both GPs and NNs in various medium-dimensional tasks [4, 5]. In addition, some methods [6, 5] train DGPs in ways that closely resemble backpropagation in NNs, which has also greatly improved efficiency compared to early methods [3]. However, despite recent progress, DGPs are still cumbersome to train compared to NNs. The similarity between the training procedures sheds light on a possible reason: current DGP models are forced to choose activation functions that are known to behave poorly in NNs (e.g., radial basis functions).

In this work we aim to further unify DGP and NN training, so practices that are known to work for NNs can be applied in DGPs. We do this by developing a DGP that has an identical forward pass to a typical ReLU NN. This link provides practical advantages for both models. The training of DGPs can be improved by taking best practices from NNs, to the point where a DGP can even be initialised from a NN trained in its usual way. Conversely, NN solutions can be endowed with better uncertainty estimates by continued training with the DGP training objective.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

2 Related Work

Many different relationships between GPs and NNs have been established over the years. These relationships mainly arise from Bayesian approaches to neural networks. Finding the posterior distribution on neural network weights is challenging, as closed-form expressions do not exist. As a consequence, developing accurate approximations has been a key goal since the earliest works on Bayesian NNs [7]. While investigating single-layer NN posteriors, Neal [8] noticed that randomly initialised NNs converged to a *Gaussian process* (GP) in the limit of infinite width. Like NNs, GPs represent functions, although they do not do so through weights. Instead, GPs specify function values at observed points, and a kernel which describes how function values at different locations influence each other.

This relationship was of significant practical interest, as the mathematical properties of GPs could **1)** represent highly flexible NNs with *infinite* weights, and **2)** perform Bayesian inference without approximations. This combination was highly successful for providing accurate predictions with reliable uncertainty estimates [9, 10]. To obtain models with various properties, GP analogues were found for infinitely-wide networks with various activation functions [11, 12] including the ReLU [13]. More recently, Meronen et al. [14] investigated the relationship in the opposite direction, by deriving activation functions such that the infinite width limit converges to a given GP prior.

Since the growth of modern deep learning, relationships have also been established between infinitely wide *deep* networks and GPs [15, 16, 17]. Given these relationships, one may wonder whether GPs can supersede NNs, particularly given the convenience of Bayesian inference in them. Empirically though, finite NNs outperform their GP analogues [16, 18, 19] on high-dimensional tasks such as images. MacKay [12] explained this by noting that NNs lose their ability to learn features in the infinite limit, since every GP can be represented as a single-layer NN, with *fixed* features [20, 21]. This observation justifies the effort of performing approximate Baysian inference in *finite* deep networks, so both feature learning and uncertainty can be obtained. Renewed effort in Bayesian training procedures has focussed on computationally scalable techniques that take advantage of modern large datasets, and have provided usable uncertainty estimates [e.g., 22, 23, 24, 25].

However, questions remain about the accuracy of these approximations. For accurate Bayesian inference, marginal likelihoods are expected to be usable for hyperparameter selection [26, 27]. For most current approximations, there is either no positive or explicitly negative [23] evidence for this holding, although recent approaches do seem to provide usable estimates [28, 29].

DGPs [3] provide an alternative approach to deep NNs, which use GP layers instead of weight-based layers, in order to take advantage of the improved uncertainty estimates afforded by having an infinite number of weights. The DGP representation is particularly promising because both early and recent work [3, 30, 31] shows that marginal likelihood estimates are usable for hyperparmeter selection, which indicates accurate inference. However, currently scalability and optimisation issues hinder widespread use.

3 Background

In this section, we review Deep Gaussian processes (DGPs), with a particular focus on the structure of the approximate posterior and its connection to deep NNs. We then discuss how inducing variables control the DGP activation function, which is a key ingredient for our method.

3.1 Deep Gaussian Processes and Sparse Variational Approximate Inference

GPs are defined as an infinite collection of random variables $\{f(\boldsymbol{x})\}_{x\in\mathbb{R}^d}$ such that for all $n \in \mathbb{N}$, the distribution of any finite subset $\{f(\boldsymbol{x}_1), \dots, f(\boldsymbol{x}_n)\}$ is multivariate normal. The GP is fully specified by its mean function $\mu(\boldsymbol{x}) = \mathbb{E}[f(\boldsymbol{x})]$ and its kernel $k(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{C}(f(\boldsymbol{x}), f(\boldsymbol{x}'))$, shorthanded as $f(\cdot) \sim \mathcal{GP}(\mu(\cdot), k(\cdot,\cdot))$. GPs are often used as priors over functions in supervised learning, as for Gaussian likelihoods the posterior $f(\cdot) \,|\, \mathcal{D}$ given a dataset $\mathcal{D} = \{\boldsymbol{x}_i \in \mathbb{R}^d, y_i \in \mathbb{R}\}_{i=1}^N$ is tractable.

Surprisingly diverse properties of the GP prior can be specified by the kernel [2, 32, 33, 34, 35], although many problems require the flexible feature learning that deep NNs provide. DGPs [3] propose to solve this in a fully Bayesian way by composing several layers of simple GP priors,

$$\mathcal{F}(\cdot) = f_L(\dots f_2(f_1(\cdot))) \qquad \text{where} \qquad f_\ell(\cdot) \sim \mathcal{GP}\big(0, k_\ell(\cdot,\cdot)\big). \tag{1}$$

Inference in DGPs is challenging as the composition is no longer a GP, leading to an intractable posterior $\mathcal{F}(\cdot) \,|\, \mathcal{D}$. We follow the variational approach by Salimbeni and Deisenroth [5] due to its similarity to backpropagation in NNs. They use an approximate posterior consisting of an independent sparse GP for each layer. Each sparse GP is constructed by conditioning the prior on M inducing variables [36, 37, 38], commonly function values $\boldsymbol{u}_\ell = \{u_\ell^m = f_\ell(\boldsymbol{w}_\ell^m)\}_{m=1}^M$, and then specifying their marginal distribution as $q(\boldsymbol{u}_\ell) = \mathcal{N}(\boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$. This results in the approximate posterior process:

$$q(f_\ell(\cdot)) = \mathcal{GP}\Big(\boldsymbol{b}_\ell^\top \boldsymbol{c}_{\boldsymbol{u}_\ell}(\cdot); \quad k_\ell(\cdot,\cdot') + \boldsymbol{c}_{\boldsymbol{u}_\ell}^\top(\cdot)\mathbf{C}_{\boldsymbol{u}_\ell\boldsymbol{u}_\ell}^{-1}\left(\boldsymbol{\Sigma}_\ell - \mathbf{C}_{\boldsymbol{u}_\ell\boldsymbol{u}_\ell}\right)\mathbf{C}_{\boldsymbol{u}_\ell\boldsymbol{u}_\ell}^{-1}\boldsymbol{c}_{\boldsymbol{u}_\ell}(\cdot')\Big), \tag{2}$$

where $\boldsymbol{b}_\ell = \mathbf{C}_{\boldsymbol{u}_\ell\boldsymbol{u}_\ell}^{-1}\boldsymbol{\mu}_\ell$ (or $\mathbf{B}_\ell$ in the multi-output case) and $\boldsymbol{c}_{\boldsymbol{u}_\ell}(\cdot) = \mathbb{C}(f_\ell(\cdot), \boldsymbol{u}_\ell)$ and $\mathbf{C}_{\boldsymbol{u}_\ell\boldsymbol{u}_\ell} = \mathbb{C}(\boldsymbol{u}_\ell, \boldsymbol{u}_\ell)$. It is worth noting that, while the mean has a finite number of parameters, the variance does take the infinite degrees of freedom from the prior into account. The variational parameters $\boldsymbol{\mu}_\ell \in \mathbb{R}^M$ and $\boldsymbol{\Sigma}_\ell \in \mathbb{R}^{M\times M}$ are selected by reducing the Kullback-Leibler (KL) divergence from the variational distribution to the true posterior. This is equivalent to maximising the Evidence Lower BOund (ELBO). Taking the evaluations of the GP as $\boldsymbol{h}_\ell = f_\ell(\boldsymbol{h}_{\ell-1})$, the ELBO becomes [5]:

$$\text{ELBO} = \sum\nolimits_i \mathbb{E}_{q(h_{i,L})}[\log p(y_i \,|\, h_{i,L})] - \sum\nolimits_\ell \text{KL}[q(\boldsymbol{u}_\ell) \,\|\, p(\boldsymbol{u}_\ell)] \leq \log p(\boldsymbol{y}). \tag{3}$$

3.2 Connection between Deep Gaussian processes and Deep Neural Networks

Hensman and Lawrence [6] observed that propagating an input through the variational means becomes:

$$\mathbb{E}[f_L(\ldots\mathbb{E}[f_2(\mathbb{E}[f_1(\boldsymbol{x})])])] = \mathbf{B}_L^\top \boldsymbol{c}_{\boldsymbol{u}_L}(\ldots\mathbf{B}_2^\top \boldsymbol{c}_{\boldsymbol{u}_2}(\mathbf{B}_1^\top \boldsymbol{c}_{\boldsymbol{u}_1}(\boldsymbol{x}))). \tag{4}$$

This resembles the forward pass through fully-connected layers with non-linearity $\sigma(\cdot)$:

$$\mathcal{F}_{\text{DNN}}(\boldsymbol{x}) = \mathbf{V}_L^\top \sigma(\mathbf{W}_L \ldots \mathbf{V}_2^\top \sigma(\mathbf{W}_2 \mathbf{V}_1^\top \sigma(\mathbf{W}_1 \boldsymbol{x}))), \tag{5}$$

with $\mathbf{W}_\ell$ and $\mathbf{V}_\ell$ the pre-activation and output weights, respectively. Both models are visualised in Fig. E.1. If we can formulate an approximation that makes the covariance $\boldsymbol{c}_{\boldsymbol{u}}(\cdot)$ the same as a typical neural net activation function, we will have a formal mathematical equivalence to the forward pass of a DNN. This is one of the main contributions of this work.

3.3 Interdomain Inducing Features

The basis functions used in the approximate posterior mean (Eq. (2)) are determined by the covariance between the inducing variables and other function evaluations $[\boldsymbol{c}_{\boldsymbol{u}}(\cdot)]_m = \mathbb{C}(f(\cdot), u_m)$. Commonly, the inducing variables are taken to be function values $u_m = f(w_m)$, which leads to the kernel becoming the basis function $[\boldsymbol{c}_{\boldsymbol{u}}(\cdot)]_m = k(\boldsymbol{w}_m, \cdot)$. *Interdomain* inducing variables [39] select different properties of the GP (e.g. integral transforms $u_m = \int f(\boldsymbol{x}) g_m(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}$), which modifies this covariance (see [40, 41] for an overview), and therefore gives control over the basis functions. Most current interdomain methods are designed to improve computational properties [42, 43, 44]. Our aim is to control $\boldsymbol{c}_{\boldsymbol{u}}(\cdot)$ to be a typical NN actvation function like a ReLU or Softplus. This was also investigated by Sun et al. [45], although they found less common saturating activation functions.

3.4 The Arc Cosine Kernel and its associated RKHS

The first order Arc Cosine kernel mimics the computation of infinitely wide fully connected layers with ReLU activations. Cho and Saul [13] showed that for $\sigma(t) = \max(0, t)$, the covariance between function values of $f(\boldsymbol{x}) = \sigma(\boldsymbol{w}^\top \boldsymbol{x})$ for $\boldsymbol{w} \sim \mathcal{N}\big(0, d^{-1/2}\mathbf{I}_d\big)$ and $\boldsymbol{w} \in \mathbb{R}^d$ is given by

$$k(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}_{\boldsymbol{w}}\big[\sigma(\boldsymbol{w}^\top \boldsymbol{x})\,\sigma(\boldsymbol{w}^\top \boldsymbol{x}')\big] = \underbrace{||\boldsymbol{x}||||\boldsymbol{x}'||}_{\text{radial}}\ \underbrace{\frac{1}{\pi}\big(\sqrt{1-t^2} + t\,(\pi - \arccos t)\big)}_{\text{angular (shape function) } s(t)}, \tag{6}$$

where $t = \frac{\boldsymbol{x}^\top \boldsymbol{x}'}{||\boldsymbol{x}||||\boldsymbol{x}'||}$. The factorisation of the kernel in a radial and angular factor leads to an RKHS consisting of functions of the form $f(\boldsymbol{x}) = ||\boldsymbol{x}||\, g(\frac{\boldsymbol{x}}{||\boldsymbol{x}||})$, where $g(\cdot)$ is defined on the unit hypersphere $\mathbb{S}^{d-1} = \{\boldsymbol{x} \in \mathbb{R}^d : ||\boldsymbol{x}||_2 = 1\}$ but fully determines the function on $\mathbb{R}^d$.

The shape function can be interpreted as a kernel itself, since it is the restriction of $k(\cdot,\cdot)$ to the unit hypersphere. Furthermore its expression only depends on the dot-product between the inputs so it is a zonal kernel (also known as a dot-product kernel [46]). This means that the eigenfunctions of the angular part of $k(\cdot,\cdot)$ are the spherical harmonics $\phi_{n,j}$ (we index them with a level n and an index within each level $j \in \{1,\dots,N_n^d\}$) [47, 44]. Their associated eigenvalues only depend on n:

$$\lambda_n = \frac{\omega_d}{C_n^{(\alpha)}(1)} \int_{-1}^{1} s(t)\, C_n^{(\alpha)}(t)\, (1-t^2)^{\frac{d-3}{2}} \mathrm{d}t, \tag{7}$$

where $C_n^{(\alpha)}(\cdot)$ is the Gegenbauer polynomial[1] of degree n, $\alpha = \frac{d-2}{2}$, ω_d is a constant that depends on the surface area of the hypersphere. Analytical expressions of λ_n are provided in Appendix C. The above implies that k admits the Mercer representation:

$$k(\boldsymbol{x},\boldsymbol{x}') = ||\boldsymbol{x}||\,||\boldsymbol{x}'|| \sum_{n=0}^{\infty}\sum_{j=1}^{N_n^d} \lambda_n \phi_{n,j}\left(\frac{\boldsymbol{x}}{||\boldsymbol{x}||}\right)\phi_{n,j}\left(\frac{\boldsymbol{x}'}{||\boldsymbol{x}'||}\right), \tag{8}$$

and that the inner product between any two functions $g,\ h \in \mathcal{H}_k$ is given by:

$$\langle g, h\rangle_{\mathcal{H}_k} = \sum_{n,j} \frac{g_{n,j} h_{n,j}}{\lambda_n} \tag{9}$$

where $g_{n,j}$ and $h_{n,j}$ are the Fourier coefficients of g and h, i.e. $g(\boldsymbol{x}) = \sum_{n,j} g_{n,j} ||\boldsymbol{x}|| \phi_{n,j}(\boldsymbol{x})$.

4 Method: Gaussian Process Layers with Activated Inducing Variables

The concepts introduced in the background section can now be used to summarise our approach: We consider DGP models with GP layers that have an Arc Cosine kernel and Variational Fourier Feature style inducing variables $u_m = \langle f(\cdot), g_m(\cdot)\rangle_{\mathcal{H}_k}$ [42]. We then choose inducing functions $g_m(\cdot)$ that have the same shape as neural network activation functions (Section 4.1). This yields basis functions for the SVGP model that correspond to activation functions (Section 4.2), and thus to a model whose mean function can be interpreted as a classic feed forward neural network. Section 4.3 covers the mathematical intricacies associated with the construction described above.

4.1 Activated Inducing Functions and their Spherical Harmonic Decomposition

The RKHS of the Arc Cosine kernel consists solely of functions that are equal to zero at the origin, i.e. $\forall f \in \mathcal{H}_k : f(\mathbf{0}) = 0$. To circumvent this problem we artificially increase the input space dimension by concatenating a constant to each input vector. In other words, the data space is embedded in a larger space with an offset such that it does not contain the origin anymore. This is analogous to the bias unit in multi-layer perceptron (MLP) layers in neural networks. For convenience we will denote by $(d-1)$ the dimension of the original data space (i.e. the number of input variables), and by d the dimension of the extended space on which the Arc Cosine kernel is defined.

The inducing functions $g_m(\cdot)$ play an important role because they determine the shape of the SVGP's basis functions. Ideally they should be defined such that their restriction to the $(d-1)$-dimensional original data space matches classic activation functions, such as the ReLU or the Softplus, exactly. However, this results in an angular component for $g_m(\cdot)$ that is not necessarily zonal. Since this property will be important later on, we favour the following alternative definition that enforces zonality:

$$g_m : \mathbb{R}^d \to \mathbb{R}, \qquad \boldsymbol{x} \mapsto ||\boldsymbol{x}||\,||\boldsymbol{w}_m||\,\sigma\left(\frac{\boldsymbol{w}_m^\top \boldsymbol{x}}{||\boldsymbol{w}_m||\,||\boldsymbol{x}||}\right), \tag{10}$$

with $\boldsymbol{w}_m \in \mathbb{R}^d$ a parameter, and $\sigma : [-1,1] \to \mathbb{R}$ the function that determines the value of $g_m(\cdot)$ on the unit hypersphere. In Fig. 1, we show that choosing $\sigma(\cdot)$ to be a ReLU ($\sigma(t) = \max(0,t)$) or a Softplus ($\sigma(t) = \log(1+\exp(3\,t))$) activation function leads to inducing functions that closely resemble the classic ReLU and Softplus on the data space. In the specific case of the ReLU it can actually be shown that the match is exact. In both cases, The parameter $\boldsymbol{w}_m \in \mathbb{R}^d$ determines the

[1] See Appendix B for a primer on Gegenbauer polynomials and spherical harmonics.

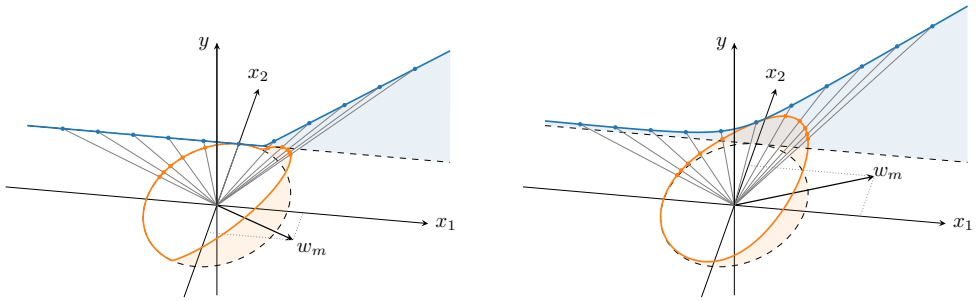

(a) ReLU: $\sigma(t) = \max(0, t)$ (b) Softplus: $\sigma(t) = \log(1 + \exp(3\,t))$

Figure 1: Activated inducing function $g_m(\boldsymbol{x}) = ||\boldsymbol{x}||\,||\boldsymbol{w}_m||\,\sigma\big(\boldsymbol{w}_m^\top \boldsymbol{x} \,/\, ||\boldsymbol{w}_m||\,||\boldsymbol{x}||\big)$ where $\sigma(\cdot)$ correspond to the ReLU **(a)** and Softplus **(b)**. Although the input domain is $\mathbb{R}^2$ we only plot the value of the function on the unit circle $\mathbb{S}^1$ (orange), and on the subspace that has an offset of 1 in the x_2 direction (blue). The linear radial component of $g_m(\cdot)$ creates a one-to-one mapping between the blue curve and the upper half of the orange one.

orientation and slope of the activation function—they play the same role as the pre-activation weights in a NN.

The zonality that we enforced in Eq. (10) is particularly convenient when it comes to representing $g_m(\cdot)$ in the basis of the eigenfunctions of $\mathcal{H}_k$, which is required for computing inner products. It indeed allows us to make use of the Funk-Hecke theorem (see Appendix B) and to eventually obtain

$$g_m(\boldsymbol{x}) = ||\boldsymbol{x}||\,||\boldsymbol{w}_m|| \sum_{n=0}^{\infty} \sum_{j=1}^{N_n^d} \sigma_n \phi_{n,j}\left(\frac{\boldsymbol{w}_m}{||\boldsymbol{w}_m||}\right) \phi_{n,j}\left(\frac{\boldsymbol{x}}{||\boldsymbol{x}||}\right),$$
$$\text{where} \quad \sigma_n = \frac{\omega_d}{C_n^{(\alpha)}(1)} \int_{-1}^{1} \sigma(t)\, C_n^{(\alpha)}(t)\, (1-t^2)^{\frac{d-3}{2}}\, \mathrm{d}t. \tag{11}$$

Analytical expressions for σ_n when $\sigma(t) = \max(0, t)$ are given in Appendix C.

4.2 Activated Interdomain Inducing Variables

We define our *activated* interdomain inducing variables as

$$u_m = \big\langle f(\cdot), g_m(\cdot) \big\rangle_{\mathcal{H}_k}. \tag{12}$$

Since the GP samples do not belong to the RKHS there are mathematical subtleties associated with such a definition, which are detailed in Section 4.3. Assuming for now that they are indeed well defined, using these interdomain inducing variables as part of the SVGP framework requires access to two quantities: (i) their pairwise covariance, and (ii) the covariance between the GP and the inducing variables. The pairwise covariance, which is needed to populate $\mathbf{C}_{\boldsymbol{uu}}$, is given by

$$\mathbb{C}(u_m, u_{m'}) = \langle g_m(\cdot), g_{m'}(\cdot) \rangle_{\mathcal{H}_k} = \sum_{\substack{n=0 \\ \lambda_n \neq 0}}^{\infty} \frac{\sigma_n^2}{\lambda_n} \frac{n+\alpha}{\alpha} C_n^{(\alpha)}\left(\frac{\boldsymbol{w}_m^\top \boldsymbol{w}_{m'}}{||\boldsymbol{w}_m||\,||\boldsymbol{w}_{m'}||}\right). \tag{13}$$

The above is obtained using the RKHS inner product from Eq. (9), the Fourier coefficients from Eq. (11) and the addition theorem for spherical harmonics from Appendix B. Secondly, the covariance between the GP and u_m, which determines $[\boldsymbol{c}_{\boldsymbol{u}}(\cdot)]_m$, is given by:

$$\mathbb{C}(u_m, f(\boldsymbol{x})) = \langle k(\boldsymbol{x}, \cdot), g_m(\cdot) \rangle_{\mathcal{H}_k} = g_m(\boldsymbol{x}) \tag{14}$$

as a result of the reproducing property of the RKHS. It becomes clear that this procedure gives rise to basis functions that are equal to our inducing functions. By construction, these inducing functions match neural network activation functions in the data plane, as shown in Fig. 1. Using these inducing variables thus leads to an approximate posterior GP (Eq. (2)) which has a mean that is equivalent to a fully-connected layer with a non-linear activation function (e.g. ReLU, Softplus, Swish).

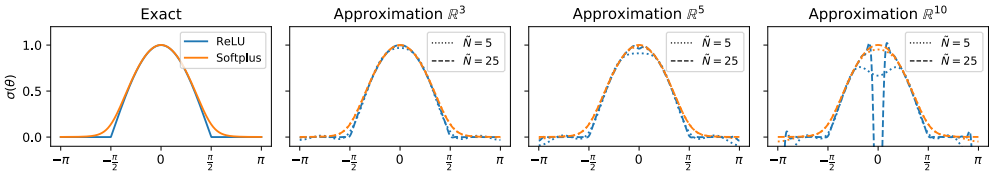

Figure 2: The exact ReLU and Softplus activation function and its approximation for different truncation levels and dimensions. These functions correspond to the orange function in Fig. 1 plotted on a line rather than on the circle.

4.3 Analysis of the interplay between kernels and inducing functions

In this section we describe the mathematical pitfalls that can be encountered [e.g., 45] when defining new inducing variables of the form of Eq. (12), and how we address them. We discuss two problems: 1) the GP and the inducing function are not part of the RKHS, 2) the inducing functions are not expressive enough to explain the prior. Both problems manifest themselves in an approximation that is overly smooth and over-estimates the predictive variance.

The Mercer representation of the kernel given in Eq. (8) implies that we have direct access to the Karhunen–Loève representation of the GP:

$$f(\boldsymbol{x}) = \sum_{n=0}^{\infty} \sum_{j=1}^{N_n^d} \xi_{n,j} \sqrt{\lambda_n} ||\boldsymbol{x}|| \phi_{n,j}\left(\frac{\boldsymbol{x}}{||\boldsymbol{x}||}\right), \quad \text{where the } \xi_{n,j} \text{ are i.i.d. } \mathcal{N}(0,1). \tag{15}$$

Using this expression to compute the RKHS norm of a GP sample $f(\cdot)$ yields $||f||^2 = \sum_{n,j} \xi_{n,j}^2$, which is a diverging series. This is a clear indication that the GP samples do not belong to the RKHS [48], and that expressions such as $\langle f(\cdot), g(\cdot)\rangle_{\mathcal{H}_k}$ should be manipulated with care. According to the definition given in Eq. (9), the RKHS inner product is a series that converges if computed for any two elements $g(\cdot), h(\cdot) \in \mathcal{H}_k$. The inner product operator can however be extended to the case where one input lives in a space larger than $\mathcal{H}_k$, provided that restrictions are introduced on the second input to guarantee the convergence of the series. In other words, even if the decay of the Fourier coefficients of $f(\cdot)$ is too slow to make it an element of $\mathcal{H}_k$, if the Fourier coefficients of $g(\cdot)$ decay quickly enough for the series $\sum_{n,j} \xi_{n,j} g_{n,j}/\sqrt{\lambda_n}$ to converge then $\langle f(\cdot), g(\cdot)\rangle_{\mathcal{H}_k}$ is well defined.

The above reasoning indicates that, for a given kernel $k(\cdot,\cdot)$, some activation functions $g_m(\cdot)$ will result in inducing variables $u_m = \langle f(\cdot), g_m(\cdot)\rangle_{\mathcal{H}_k}$ that are well defined whereas other activation functions do not. For example, if we consider the case of the Arc Cosine kernel and the ReLU inducing function, the decay rate of σ_n is proportional to the square root of λ_n [49, 46]. This implies that the inner product series diverges and that this kernel and inducing variable cannot be used together. Alternatively, using smoother activation functions for $\sigma(t)$, such as the Softplus, results in a faster decay of the coefficients σ_n and can guarantee that inducing variables are well defined.

An alternative to ensure the series convergence for any combination of kernel and activation function is to use a truncated approximation of the activation function $\tilde{g}_m(\cdot)$ where all the Fourier coefficients above a given level $\tilde{N}$ are set to zero, which basically turns the inner product series into a finite sum. Figure 2 shows how the true and truncated activation functions for the ReLU and Softplus compare. These correspond to the orange functions in Fig. 1, but are now plotted on a line. In the low to medium dimensional regime, we see that even for small truncation levels we approximate the ReLU and Softplus well. In higher dimensions this becomes more challenging for the ReLU.

Unexpressive inducing variables through truncation (Fig. 3) The main concern with this truncation approach, however, comes from elsewhere: the larger $\tilde{N}$ is, the closer $\tilde{g}_m(\cdot)$ is to $g_m(\cdot)$, but the larger $||\tilde{g}_m(\cdot)||_{\mathcal{H}_k}$ becomes (to the point where it may be arbitrarily large). Similarly to ridge regression where the norm acts as a regulariser, using inducing functions with a large norm in SVGP models comes with a penalty which enforces more smoothness in the approximate posterior and limits its expressiveness. Figure 3 shows how the norm of our ReLU inducing functions grow in the RKHS. So by making $\tilde{N}$ larger such that we approximate the ReLU better, we incur a greater penalty

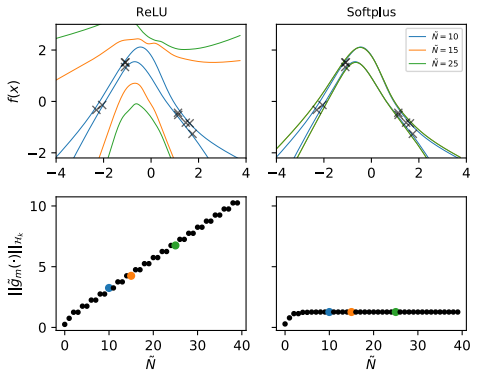

Figure 3: **Top:** Predictive variance an SVGP fit on a synthetic dataset using our ReLU (left) and Softplus (right) inducing variables for $\tilde{N} = 10, 15$ and 25. **Bottom:** the norm of the inducing function $\tilde{g}_m(\cdot)$ as a function of $\tilde{N}$.

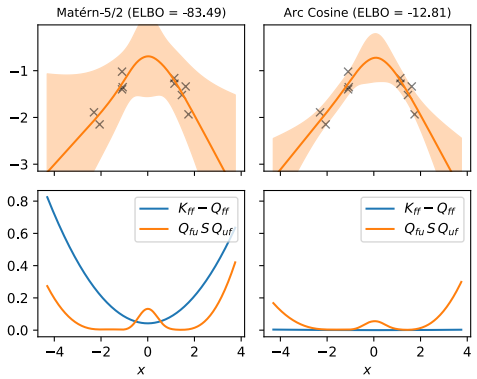

Figure 4: **Top:** Predictive mean and variance for an SVGP model using our Softplus inducing variables for the Matérn-5/2 (left) and Arc Cosine (right) kernel. **Bottom:** The two terms that constitute the predictive variance.

in the ELBO for using them. This leads to unexpressive inducing variables, which can be seen by the growing predictive uncertainty. The Softplus, which is part of the RKHS, does not suffer from this.

Unexpressive inducing variables through spectra mismatch (Fig. 4) Any stationary kernel whose inputs $\boldsymbol{x}, \boldsymbol{x}' \in \mathbb{S}^{d-1}$ are restricted to the unit hypersphere is a zonal kernel (i.e., the kernel only depends on the dot-product). This means that we are not limited to only the Arc Cosine because we can replace the shape function in Eq. (6) by any stationary kernel, and our approach would still hold. For example, we could use the Matérn-5/2 shape function with $s_{\text{mat-5/2}}(t) = \left(1 + \sqrt{5}t + 5t^2/3\right)\exp\left(-\sqrt{5}t\right)$. However, in Fig. 4 we compare the fit of an SVGP model using a Matérn-5/2 kernel (left) to a model using an Arc Cosine kernel (right). While both models use our Softplus inducing variables, we clearly observe that the Matérn kernel gives rise to a worse posterior model (lower ELBO and an over-estimation of the variance). In what follows we explain why this is the case.

In the bottom panel in Fig. 4 we see that for the Matérn-5/2, the Nyström approximation $\mathbf{Q}_{\boldsymbol{ff}} = \mathbf{C}_{\boldsymbol{fu}}\mathbf{C}_{\boldsymbol{uu}}^{-1}\mathbf{C}_{\boldsymbol{fu}}^{\top}$ is unable to explain the prior $\mathbf{K}_{\boldsymbol{ff}}$, imposed by the kernel. This leads to the overestimation of the predictive variance in the top plot. The reason for this problem is the mismatch between the eigenvalues λ_n (Eq. (7)) of the Matérn and the Fourier coefficients σ_n (Eq. (11)) of the Softplus inducing function. As shown in Fig. C.2, the Matérn kernel has a full spectrum (i.e., $\lambda_n \neq 0,\ \forall n \in \mathbb{N}$), whereas the coefficients for the Softplus σ_n are zero at levels $n = 3, 5, 7, \cdots$. We are thus trying to approximate our prior kernel, containing all levels, by a set of functions that is missing many. This problem does not occur for the combination of Softplus (or ReLU) inducing functions and the Arc Cosine kernel (right-hand side) because their spectral decomposition match. In other words, the Arc Cosine kernel has zero coefficients for the same levels as our activated inducing functions.

5 Experiments

We have shown that that we can build SVGP models with basis functions that behave like neural net activations. This equivalence has the practical advantage that we can train the means of the GP layers in our DGP as if they are a neural network model — making use of the great advances made by the deep learning community. Once the mean of the DGP is trained, we can further optimise the remaining model hyper- and variational parameters w.r.t. the ELBO, which is a more principled objective [50]. This approach allows us to exploit the benefit of both, the efficient training of the DNN in combination with the principled uncertainty estimate of the DGP. For all SVGP models we use the Arc Cosine kernel and inducing variables obtained with the Softplus activation (Section 4.1) with $\tilde{N} = 20$, for the reasons explained in Section 4.3.

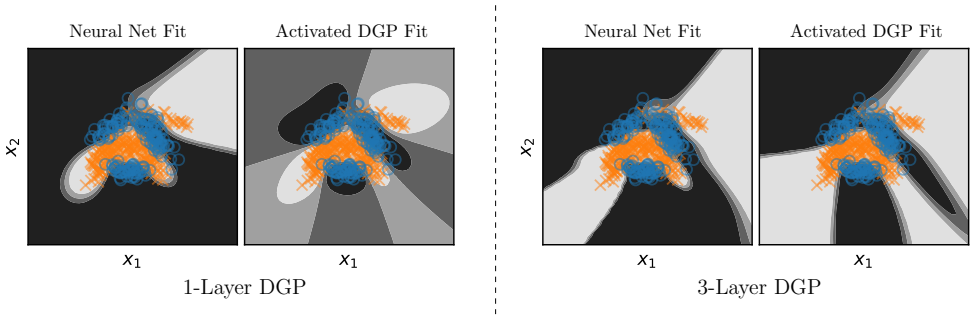

Figure 5: We compare the fit of a single-layer and three-layer DNN optimised using binary cross-entropy and it's equivalent DGP trained using the ELBO. The DNNs are very confident, even far away from the data. We used the DNNs to initialise the equivalent DGP posterior mean before optimising the ELBO, which in both situations leads to a model exhibiting more calibrated uncertainty.

The aim of the experiments is to highlight that (i) our method leads to valid inducing variables, (ii) our initialisation improves DGPs in terms of accuracy, and (iii) we are able to improve on Dropout-based [24] neural networks in terms of calibrated uncertainty. We acknowledge that the NNs we benchmark against are the models for which we can build an equivalent DGP. While this leads to a fair comparison, it excludes recent improvements such as Transformer and Batch Norm layers.

From Neural Network to Activated DGP Figure 5 shows the difference in predictive probability $p(y \mid \boldsymbol{x})$ for a DNN and our activated DGP, in the single-layer and three-layer case. We configure the models with a Softplus activation function and set the number of both inducing variables of the GP and hidden units of the DNN to 100. In this experiment, the first step is to optimise the DNN w.r.t. the binary cross-entropy objective, upon convergence we initialise the DGP with this solution and resume training of the DGP w.r.t. the ELBO. Especially in the single-layer case, we notice how the sharp edges from the NN are relaxed by the GP fit, and we see how the GP expresses uncertainty away from the data by letting $p(y \mid \boldsymbol{x}) \approx 0.5$. This is thanks to the ELBO, which balances both data fit and model complexity and simultaneously trains the uncertainty.

Regression on UCI benchmarks We compare a suit of models on a range of regression problems. The important aspect is that we keep the model configuration and training procedure fixed across datasets. We use three-layered models with 256 inducing variables or hidden units. In each layer, the number of output heads is equal to the input dimensionality of the data. The Activated DGP (ADGP) and neural network approaches (Vanilla and Dropout) use Softplus activation functions. The Dropout model [24] uses a rate of 0.1 during train and test, and the Vanilla model is a deterministic neural network that uses the training MSE as the empirical variance during prediction.

The DGP and ADGP both use the Arc Cosine kernel. The main difference is that the DGP has standard inducing points $u_m = f(\boldsymbol{z}_m)$, whereas ADGP makes use of our activated inducing variables $u_m = \langle f, g_m \rangle_{\mathcal{H}_k}$. The ADGP is trained in two steps: we first train the mean of the approximate posterior w.r.t. the MSE, and then optimise all parameters w.r.t. the ELBO.

In Fig. 6 we see that in general our ADGP model is more accurate than its neural network initialisation (Vanilla model) in terms of RMSE. This is a result of the second stage of training in which we use the ELBO rather than the MSE, which is especially beneficial to prevent overfitting on the smaller datasets. When it comes to NLPD, ADGP shows improvements over its NN initialisation for 5 datasets out of 7 and consistently outperforms classic DGPs.

Large scale image classification In this experiment we measure the performance of our models under dataset shifts [51]. For MNIST and Fashion-MNIST the out-of-distribution (OOD) test sets consist of rotated digits — from 0°(i.e. the original test set) to 180°. For CIFAR-10 we apply four different types of corruption to the test images with increasing intensity levels from 0 to 5. For MNIST and FASHION-MNIST the models consist of two convolutional and max-pooling layers, followed by two dense layers with 128 units and 10 output heads. The dense layers are either fully-connected neural network layers using a Softplus activation function (Vanilla and Dropout), or our Activated

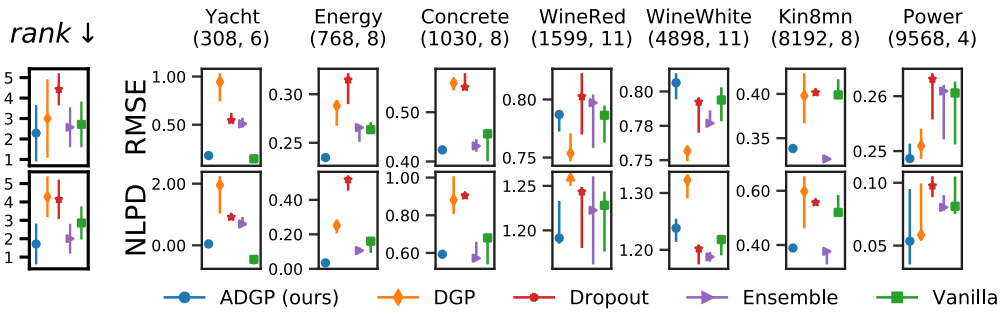

Figure 6: UCI: Root Mean Squared Error (RMSE) and Negative Log Predictive Density (NLPD) with 25% and 75% quantile error bars based on 5 splits. Dataset size and dimension given in parentheses.

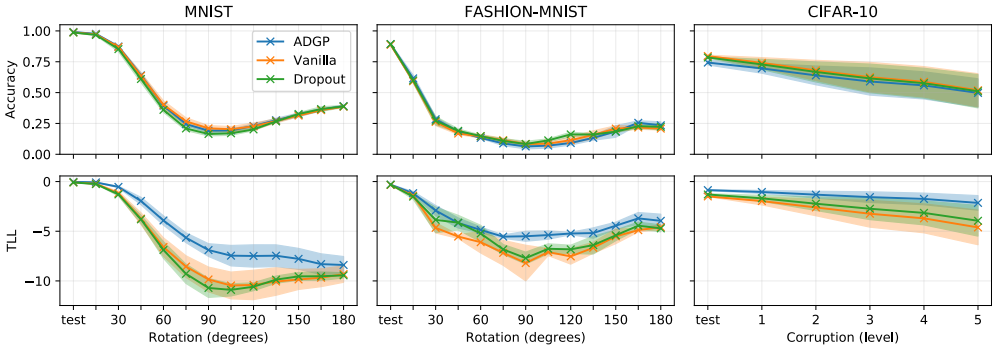

Figure 7: Results on the rotated MNIST, FASHION-MNIST and corrupted CIFAR-10, showing the mean and std. dev. of the accuracy **(top)**, and test log-likelihood (TLL) **(bottom)**.

GP layers using the Arc Cosine kernel and Softplus inducing variables (ADGP). For the CIFAR-10 models, we use the residual convolutional layers from a ResNet [52] to extract useful features before passing them to our dense GP or NN layers. Details of the model architectures are given in Appendix E. As previously, the ADGP model is initialised to the solution of the Vanilla model, and training is then continued using the ELBO. In Fig. 7 we observe that the models perform very similar in terms of prediction accuracy, but that ADGP better account for uncertainty as evidenced by the Test Log Likelihood metric.

6 Conclusion

In this work, we establish a connection between fully-connected neural networks and the posterior of deep sparse Gaussian processes. We use a specific flavour of interdomain inducing variables based on the RKHS inner product to obtain basis functions for the SVGP that match activation functions from the neural network literature. By composing such SVGPs together, we obtain a DGP for which a forward pass through the mean of each layer is equivalent to a forward pass in a DNN. We also address important mathematical subtleties to ensure the validity of the approach and to gain insights on how to choose the activation function. As demonstrated in the experiments, being able to interpret the same mathematical expression either as a deep neural network or as the mean of an deep Gaussian process can benefit both approaches. On the one hand it allows us to improve the prediction accuracy and the uncertainty representation of the neural network by regularising it with the ELBO obtained from a GP prior. On the other hand, it enables us to better optimise deep Gaussian process model by initialising them with a pre-trained neural network. We believe that by providing an equivalence between the two models, not in the limit but in the operational regime opens many future research directions allowing for beneficial knowledge transfer between the two domains.

References

[1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. The MIT Press, 2016.

[2] Carl E. Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[3] Andreas Damianou and Neil D. Lawrence. "Deep Gaussian Processes". In: *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2013.

[4] Thang Bui, Daniel Hernandez-Lobato, Jose Hernandez-Lobato, Yingzhen Li, and Richard E. Turner. "Deep Gaussian Processes for Regression using Approximate Expectation Propagation". In: *Proceedings of The 33rd International Conference on Machine Learning (ICML)*. 2016.

[5] Hugh Salimbeni and Marc P. Deisenroth. "Doubly Stochastic Variational Inference for Deep Gaussian Processes". In: *Advances in Neural Information Processing Systems 30 (NIPS)*. 2017.

[6] James Hensman and Neil D. Lawrence. "Nested Variational Compression in Deep Gaussian Processes". In: *arXiv preprint arXiv:1412.1370* (2014).

[7] David J. C. MacKay. "A Practical Bayesian Framework for Backpropagation Network". In: *Neural Computation* (1992).

[8] Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer, 1995.

[9] Christopher KI Williams and Carl Edward Rasmussen. "Gaussian processes for regression". In: (1996).

[10] Carl Edward Rasmussen. "Evaluation of Gaussian processes and other methods for non-linear regression". PhD thesis. University of Toronto Toronto, Canada, 1997.

[11] Christopher K. I. Williams. "Computing with Infinite Networks". In: *Advances in Neural Information Processing Systems 11 (NIPS 1997)* (1998).

[12] David J. C. MacKay. "Introduction to Gaussian Processes". In: (1998).

[13] Youngmin Cho and Lawrence K. Saul. "Kernel Methods for Deep Learning". In: *Advances in Neural Information Processing Systems 22 (NIPS)*. 2009.

[14] Lassi Meronen, Christabella Irwanto, and Arno Solin. "Stationary Activations for Uncertainty Calibration in Deep Learning". In: *Advances in Neural Information Processing Systems 33 (NeurIPS)*. 2020.

[15] Alexander G. de G. Matthews, Mark Rowland, Jiri Hron, Richard E. Turner, and Zoubin Ghahramani. "Gaussian process behaviour in wide deep neural networks". In: *Proceedings of the 6th International Conference on Learning Representations (ICLR)* (2018).

[16] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. "Deep neural networks as Gaussian processes". In: *6th International Conference on Learning Representation (ICLR)*. 2018.

[17] Greg Yang. "Wide feedforward or recurrent neural networks of any architecture are Gaussian processes". In: *Advances in Neural Information Processing Systems 32 (NeurIPS)*. 2019.

[18] Adrià Garriga-Alonso, Carl E. Rasmussen, and Laurence Aitchison. "Deep Convolutional Networks as shallow Gaussian Processes". In: *Proceedings of the 7th International Conference on Learning Representations (ICLR)* (2018).

[19] Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Jiri Hron, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. "Bayesian deep convolutional networks with many channels are Gaussian processes". In: *Proceedings of the 7th International Conference on Learning Representations (ICLR)* (2019).

[20] James Mercer. "Functions of positive and negative type, and their connection the theory of integral equations". In: *Philosophical transactions of the royal society of London. Series A, containing papers of a mathematical or physical character* (1909).

[21] Carl E. Rasmussen and Christopher K. I. Williams. In: *Gaussian Processes for Machine Learning*. 2006. Chap. Regression (§2.1 & §2.2).

[22] Durk Kingma, Tim Salimans, and Max Welling. "Variational Dropout and the Local Reparameterization Trick". In: *Advances in Neural Information Processing Systems 28 (NIPS)*. 2015.

[23] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. "Weight Uncertainty in Neural Network". In: *Proceedings of The 32nd International Conference on Machine Learning (ICML)*. 2015.

[24] Yarin Gal and Zoubin Ghahramani. "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning Zoubin Ghahramani". In: *Proceedings of The 33rd International Conference on Machine Learning (ICML)*. 2016.

[25] Christos Louizos and Max Welling. "Structured and efficient variational deep learning with matrix gaussian posteriors". In: *Proceedings of The 33rd International Conference on Machine Learning*. 2016.

[26] David J. C. MacKay. "Bayesian Model Comparison and Backprop Nets". In: *Advances in Neural Information Processing Systems 4 (NIPS 1991)*. 1992.

[27] David J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.

[28] Sebastian W. Ober and Laurence Aitchison. *Global inducing point variational posteriors for Bayesian neural networks and deep Gaussian processes*. 2020.

[29] Alexander Immer, Matthias Bauer, Vincent Fortuin, Gunnar Rätsch, and Mohammad Emtiyaz Khan. *Scalable Marginal Likelihood Estimation for Model Selection in Deep Learning*. 2021.

[30] Andreas Damianou. "Deep Gaussian Processes and Variational Propagation of Uncertainty". PhD thesis. University of Sheffield, 2015.

[31] Vincent Dutordoir, Mark van der Wilk, Artem Artemev, and James Hensman. "Bayesian Image Classification with Deep Convolutional Gaussian Processes". In: *Proceedings of the 23th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2020.

[32] David Duvenaud, James Lloyd, Roger Grosse, Joshua Tenenbaum, and Ghahramani Zoubin. "Structure discovery in nonparametric regression through compositional kernel search". In: *International Conference on Machine Learning*. PMLR. 2013.

[33] Andrew Wilson and Ryan Adams. "Gaussian process kernels for pattern discovery and extrapolation". In: *International conference on machine learning*. PMLR. 2013.

[34] David Ginsbourger, Nicolas Durrande, and Olivier Roustant. "Kernels and designs for modelling invariant functions: from group invariance to additivity". In: *mODa 10–Advances in Model-Oriented Design and Analysis*. 2013.

[35] Mark van der Wilk, Matthias Bauer, S. T. John, and James Hensman. "Learning Invariances using the Marginal Likelihood". In: *Advances in Neural Information Processing Systems 31 (NeurIPS)*. 2018.

[36] Michalis K. Titsias. "Variational Learning of Inducing Variables in Sparse Gaussian Processes". In: *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2009.

[37] James Hensman, Nicolo Fusi, and Neil D. Lawrence. "Gaussian Processes for Big Data". In: *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI)*. 2013.

[38] Alexander G. de G. Matthews, James Hensman, Turner E. Richard, and Zoubin Ghahramani. "On Sparse Variational Methods and the Kullback-Leibler Divergence between Stochastic Processes". In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2016.

[39] Miguel Lázaro-Gredilla and Aníbal Figueiras-Vidal. "Inter-domain Gaussian Processes for Sparse Inference using Inducing Features". In: *Advances in Neural Information Processing Systems 22 (NIPS)*. 2009.

[40] Mark van der Wilk, Vincent Dutordoir, S. T. John, Artem Artemev, Vincent Adam, and James Hensman. "A Framework for Interdomain and Multioutput Gaussian Processes". In: *arXiv preprint arXiv:2003.01115* (2020).

[41] Felix Leibfried, Vincent Dutordoir, S. T. John, and Nicolas Durrande. "A Tutorial on Sparse Gaussian Processes and Variational Inference". In: *arXiv preprint arXiv:2012.13962* (2020).

[42] James Hensman, Nicolas Durrande, and Arno Solin. "Variational Fourier Features for Gaussian Processes". In: *Journal of Machine Learning Research* (2018).

[43] David R Burt, Carl Edward Rasmussen, and Mark van der Wilk. "Variational orthogonal features". In: *arXiv preprint arXiv:2006.13170* (2020).

[44] Vincent Dutordoir, Nicolas Durrande, and James Hensman. "Sparse Gaussian Processes with Spherical Harmonic Features". In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. 2020.

[45] Shengyang Sun, Jiaxin Shi, and Roger B. Grosse. "Neural Networks as Inter-Domain Inducing Points". In: *3rd Symposium on Advances in Approximate Bayesian Inference*. 2021.

[46] Alberto Bietti and Francis Bach. "Deep Equals Shallow for ReLU Networks in Kernel Regimes". In: *arXiv preprint arXiv:2009.14397* (2020).

[47] Holger Wendland. *Scattered Data Approximation*. Cambridge University Press, 2005.

[48] Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K Sriperumbudur. "Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences". In: *arXiv preprint arXiv:1807.02582* (2018).

[49] Francis Bach. "Breaking the Curse of Dimensionality with Convex Neural Networks". In: *Journal of Machine Learning Research* (2017).

[50] Edwin Fong and Chris Holmes. "On the marginal likelihood and cross-validation". In: *arXiv preprint arXiv:1905.08737* (2019).

[51] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. "Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift". In: *Advances in Neural Information Processing Systems 32 (NeurIPS)*. 2019.

[52] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. "Deep residual learning for image recognition". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016.

[53] Feng Dai and Yuan Xu. *Approximation Theory and Harmonic Analysis on Spheres and Balls*. Springer, 2013.

[54] Costas Efthimiou and Christopher Frye. *Spherical Harmonics in p Dimensions*. World Scientific Publishing, 2014.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] The main claim is given in the introduction and accurately reflects the paper's contributions.
   (b) Did you describe the limitations of your work? [Yes] Yes, in the introductory paragraph of the experiments and in Section 4.3.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] This paper is situated in the context of supervised learning, where we are given examples of inputs and outputs, and are tasked to learn the mapping between. As such, we encourage anyone to familiarise themselves with the literature on common ethical pitfalls encountered in supervised learning, such as dataset bias. Often, the learned mapping is a first component in a larger system that is used to inform on a decision (e.g., correcting the steering of a car, deciding on medical treatment, etc.). Therefore, we believe that research into models that accurately account for the uncertainty in the prediction is extremely valuable, since it allows the risks associated with those decisions to be controlled. By providing models that offer a better quantification prediction uncertainty we are aiming at empowering people with greater understanding and safer decision making.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We only ran standardised benchmarks on publicly available dataset. The code is part of open source toolboxes which we point to upon publication to avoid breaking anonymity.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We use well-known benchmarks (e.g., UCI, MNIST, Fashion-MNIST and CIFAR-10) with standardised train and test splits. The limited number of hyperparameters are specified in Section 5 and Appendix E.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] All experiments are ran over multiple seeds and report error bars.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] Our experiments are not computationally very demanding. Every setting can be ran in less than 30 minutes on a mid-range GPU.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [N/A]

(b) Did you mention the license of the assets? [N/A]

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]