
Deep Sigma Point Processes

Martin Jankowiak*
The Broad Institute

Geoff Pleiss
Cornell University

Jacob R. Gardner
University of Pennsylvania

Abstract

We introduce Deep Sigma Point Processes, a class of parametric models inspired by the compositional structure of Deep Gaussian Processes (DGPs). Deep Sigma Point Processes (DSPPs) retain many of the attractive features of (variational) DGPs, including mini-batch training and predictive uncertainty that is controlled by kernel basis functions. Importantly, since DSPPs admit a simple maximum likelihood inference procedure, the resulting predictive distributions are not degraded by any posterior approximations. In an extensive empirical comparison on univariate and multivariate regression tasks we find that the resulting predictive distributions are significantly better calibrated than those obtained with other probabilistic methods for scalable regression, including variational DGPs—often by as much as a nat per datapoint.

1 INTRODUCTION

As machine learning becomes utilized for increasingly high risk applications such as medical treatment suggestion and autonomous driving, calibrated uncertainty estimation becomes critical. As a result, substantial effort has been devoted to the development of flexible probabilistic models. Gaussian Processes (GPs) have emerged as an important class of models in cases where predictive uncertainty estimates are essential (Rasmussen, 2003). Unfortunately, the simplest GP models often fail to match the expressive power of their neural network counterparts.

This limited flexibility has motivated the introduction of Deep Gaussian Processes (DGPs), which compose multiple layers of latent functions to build up more flexible

function priors (Damianou and Lawrence, 2013). While this class of models has shown considerable promise, inference remains a serious challenge, with empirical evidence suggesting—not surprisingly—that posterior approximations (e.g. factorizing across layers) can degrade the performance of the predictive distribution (Havasi et al., 2018). Indeed, targeting posterior approximations as an optimization objective may fail both to achieve good predictive performance and to faithfully approximate the exact posterior. Clearly, this motivates investigating alternative approaches.

In this work we take a different approach to constructing flexible probabilistic models. In particular we formulate a class of *fully parametric* models that retain many of the attractive features of variational deep GPs, while posing a much easier (maximum likelihood) inference problem. This pragmatic approach is motivated by the recognition that—especially in settings where predictive performance is paramount—it is essential to consider the interplay between modeling and inference. In particular, as motivated above, a compelling class of models may be of limited predictive utility if approximations in the inference procedure severely degrade the posterior predictive distribution. Conversely, it can be a significant advantage if a class of models admits a simple inference procedure.

Similarly to variational deep GPs, our regression models utilize predictive distributions that are mixtures of Normal distributions whose mean and variance functions make use of hierarchical composition and kernel interpolation. In contrast to deep GPs, however, our parametric perspective allows us to directly target the predictive distribution in the training objective. As we show empirically in Sec. 5 the model we introduce—the *Deep Sigma Point Process* (DSPP)—exhibits excellent predictive performance and dramatically outperforms a number of strong baselines, including Deep Kernel Learning (Calandra et al., 2016; Wilson et al., 2016) and variational Deep Gaussian Processes (Salimbeni and Deisenroth, 2017).

*Correspondence to: jankowiak@gmail.com. This work was completed while MJ and JG were at Uber AI.

2 BACKGROUND

This section is organized as follows. In Sec. 2.1-2.2 we review the basics of Gaussian Processes and inducing point methods. In Sec. 2.3 we review Deep Gaussian Processes. In Sec. 2.4 we review PPGPR (Jankowiak et al., 2020), as it serves as motivation for DSPPs in Sec. 3. We also use this section to establish our notation.

2.1 GAUSSIAN PROCESS REGRESSION

In probabilistic modeling Gaussian Processes offer flexible non-parametric function priors that are useful in various regression and classification tasks (Rasmussen, 2003). For a given input space \mathbb{R}^d GPs are entirely specified by a kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and a mean function $\mu : \mathbb{R}^d \rightarrow \mathbb{R}$. Different choices of μ and k permit the modeler to encode prior information about the generative process. In the prototypical case of univariate regression¹ the joint density takes the form

$$p(\mathbf{y}, \mathbf{f}|\mathbf{X}) = p(\mathbf{y}|\mathbf{f}, \sigma_{\text{obs}}^2)p(\mathbf{f}|\mathbf{X}) \quad (1)$$

where \mathbf{y} are the real-valued targets, \mathbf{f} are the latent function values, $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ are the N inputs with $\mathbf{x}_i \in \mathbb{R}^d$, $p(\mathbf{f}|\mathbf{X})$ is a multivariate Normal distribution with covariance $\mathbf{K}_{NN} = k(\mathbf{X}, \mathbf{X})$, and σ_{obs}^2 is the variance of the Normal likelihood $p(\mathbf{y}|\cdot)$. The marginal likelihood takes the form

$$p(\mathbf{y}|\mathbf{X}) = \int d\mathbf{f} p(\mathbf{y}|\mathbf{f}, \sigma_{\text{obs}}^2)p(\mathbf{f}|\mathbf{X}) \quad (2)$$

Eqn. 2 can be computed analytically, but doing so is computationally prohibitive for large datasets, necessitating approximate methods when N is large.

2.2 SPARSE GAUSSIAN PROCESSES

Recent years have seen significant progress in scaling Gaussian Process inference to large datasets. This advance has been enabled by the development of inducing point methods (Snelson and Ghahramani, 2006; Titsias, 2009; Hensman et al., 2013), which we now review. One begins by introducing inducing variables \mathbf{u} that depend on variational parameters $\{\mathbf{z}_m\}_{m=1}^M$, with each $\mathbf{z}_m \in \mathbb{R}^d$ and where $M = \dim(\mathbf{u}) \ll N$. One then augments the GP prior with the auxiliary variables \mathbf{u}

$$p(\mathbf{f}|\mathbf{X}) \rightarrow p(\mathbf{f}|\mathbf{u}, \mathbf{X}, \mathbf{Z})p(\mathbf{u}|\mathbf{Z})$$

and then appeals to Jensen’s inequality to lower bound the log joint density over the inducing variables and the

¹Note that here and throughout this work we focus on the regression case.

targets:

$$\begin{aligned} \log p(\mathbf{y}, \mathbf{u}|\mathbf{X}, \mathbf{Z}) &= \log \int d\mathbf{f} p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{u})p(\mathbf{u}) \quad (3) \\ &\geq \mathbb{E}_{p(\mathbf{f}|\mathbf{u})} [\log p(\mathbf{y}|\mathbf{f}) + \log p(\mathbf{u})] \\ &= \sum_{i=1}^N \log \mathcal{N}(y_i | \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{u}, \sigma_{\text{obs}}^2) \\ &\quad - \frac{1}{2\sigma_{\text{obs}}^2} \text{Tr} \tilde{\mathbf{K}}_{NN} + \log p(\mathbf{u}) \end{aligned}$$

where $\tilde{\mathbf{K}}_{NN}$ is given by

$$\tilde{\mathbf{K}}_{NN} = \mathbf{K}_{NN} - \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} \quad (4)$$

and

$$\begin{aligned} \mathbf{K}_{MM} &= k(\mathbf{Z}, \mathbf{Z}) \quad \mathbf{k}_i = k(\mathbf{x}_i, \mathbf{Z}) \\ \mathbf{K}_{NM} &= \mathbf{K}_{MN}^T = k(\mathbf{X}, \mathbf{Z}) \end{aligned} \quad (5)$$

Eqn. 3 can be used to construct a variety of algorithms for scalable GP inference; here we limit our discussion to SVGP (Hensman et al., 2013).

2.2.1 VARIATIONAL INFERENCE: SVGP

If we apply standard techniques from variational inference to the lower bound in Eqn. 3, we obtain SVGP, a popular algorithm for scalable GP inference. In more detail, SVGP proceeds by introducing a multivariate Normal variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ and computing the ELBO, which is the expectation of Eqn. 3 w.r.t. $q(\mathbf{u})$ together with an entropy term $H[q(\mathbf{u})]$:

$$\begin{aligned} \mathcal{L}_{\text{svgp}} &= \mathbb{E}_{q(\mathbf{u})} [\log p(\mathbf{y}, \mathbf{u}|\mathbf{X}, \mathbf{Z})] + H[q(\mathbf{u})] \\ &= \sum_{i=1}^N \left\{ \log \mathcal{N}(y_i | \mu_f(\mathbf{x}_i), \sigma_{\text{obs}}^2) - \frac{1}{2} \frac{\sigma_f(\mathbf{x}_i)^2}{\sigma_{\text{obs}}^2} \right\} \quad (6) \\ &\quad - \text{KL}(q(\mathbf{u})|p(\mathbf{u})) \end{aligned}$$

Here KL denotes the Kullback-Leibler divergence, $\mu_f(\mathbf{x}_i)$ is the predictive mean function given by

$$\mu_f(\mathbf{x}_i) = \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{m} \quad (7)$$

and $\sigma_f(\mathbf{x}_i)^2 \equiv \text{Var}[f_i|\mathbf{x}_i]$ denotes latent function variance

$$\sigma_f(\mathbf{x}_i)^2 = \tilde{\mathbf{K}}_{ii} + \mathbf{k}_i^T \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{k}_i \quad (8)$$

Note that the complete variational distribution used in SVGP is given by

$$q(\mathbf{f}, \mathbf{u}) = p(\mathbf{f}|\mathbf{u}, \mathbf{X})q(\mathbf{u}) \quad (9)$$

with a marginal distribution given by

$$q(\mathbf{f}) = \int d\mathbf{u} q(\mathbf{f}, \mathbf{u}) = \mathcal{N}(\mathbf{f}|\mu_f(\mathbf{X}), \Sigma_f(\mathbf{X})) \quad (10)$$

where $\Sigma_f(\mathbf{X})$ is the $N \times N$ covariance matrix

$$\Sigma_f(\mathbf{X}) = \tilde{\mathbf{K}}_{NN} + \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} \quad (11)$$

The objective $\mathcal{L}_{\text{svgp}}$, which depends on $\mathbf{m}, \mathbf{S}, \mathbf{Z}, \sigma_{\text{obs}}$ and the various kernel hyperparameters, can then be maximized with gradient methods. Since the expected log likelihood in Eqn. 6 factorizes as a sum over data points (y_i, \mathbf{x}_i) it is amenable to stochastic gradient methods, and thus SVGP can be applied to very large datasets.

2.3 DEEP GAUSSIAN PROCESSES

Deep Gaussian Processes (Damianou and Lawrence, 2013) are a natural generalization of GPs in which a sequence of GP layers form a hierarchical model in which the outputs of one GP layer become the inputs of the subsequent layer, resulting in a flexible, compositional function prior. In the simplest case of a 2-layer DGP with a univariate continuous output y and with a GP layer of width W fed into the topmost GP, the joint likelihood for a dataset (\mathbf{y}, \mathbf{X}) is given by²

$$p(\mathbf{y}, \mathbf{f}, \mathbf{G} | \mathbf{X}) = p(\mathbf{y} | \mathbf{f}, \sigma_{\text{obs}}^2) p(\mathbf{f} | \mathbf{G}) p(\mathbf{G} | \mathbf{X}) \quad (12)$$

Here \mathbf{G} is a matrix of latent function values of size $N \times W$, i.e. g_{iw} denotes the latent function value of the w^{th} GP in the first layer evaluated at the i^{th} input \mathbf{x}_i . For $i = 1, \dots, N$ the vector $\mathbf{g}_i \equiv g_{i,:}$ of dimension W then serves as the input to the second GP denoted by \mathbf{f} . See Fig. 1 for an illustration. Throughout this work we assume that the prior over \mathbf{G} factorizes as $p(\mathbf{G} | \mathbf{X}) = \prod_{w=1}^W p(\mathbf{g}_w | \mathbf{X})$, with each GP governed by its own kernel. Inference in this model is intractable, necessitating approximate methods. A popular approach is variational inference, which we now briefly review.

2.3.1 DOUBLY STOCHASTIC VARIATIONAL INFERENCE

The stochastic variational inference approach described in Sec. 2.2.1 can be generalized to the DGP setting (Salimbeni and Deisenroth, 2017). Proceeding in analogy to SVGP, we introduce inducing variables and inducing point locations for each GP and form a factorized variational distribution $Q(\mathbf{f}, \mathbf{u}_f, \dots, \mathbf{g}_W, \mathbf{u}_{g_W})$ with each factor of the form in Eqn. 9. The variational distribution Q can then be used to form the ELBO

$$\mathcal{L}_{\text{dsvi}} = \mathbb{E}_Q [\log p(\mathbf{y} | \mathbf{f}, \sigma_{\text{obs}}^2)] - \sum \text{KL} \quad (13)$$

where $\sum \text{KL}$ denotes a sum over the various KL divergences for the inducing variables $\{\mathbf{u}_f, \dots, \mathbf{u}_{g_W}\}$. Crucially, Q is easy to sample from since it suffices to sample

²We limit our discussion to this case to simplify notation; generalizations to multiple outputs and multiple layers are straightforward.

from the various marginals $\{q(f_i), \dots, q(g_{iw})\}$ and so the expected log likelihood term in Eqn. 13 factorizes³ into a sum over data points (y_i, \mathbf{x}_i) , enabling mini-batch training. As in SVGP, the topmost layer of latent function values \mathbf{f} can be integrated out analytically. All remaining latent variables (namely \mathbf{G}) must be sampled, necessitating the use of the reparameterization trick (Price, 1958; Salimans et al., 2013) and resulting in ‘doubly’ stochastic gradients. For more details we refer the reader to (Salimbeni and Deisenroth, 2017).

2.3.2 DGP PREDICTIVE DISTRIBUTIONS

We review the structure of the posterior predictive distribution that results from the variational inference procedure described in the previous section, as this will serve as motivation for the introduction of Deep Sigma Point Processes in Sec. 3. We first note that in the (single-layer) SVGP case the predictive distribution at input \mathbf{x}_* is given by the Normal distribution

$$p(y_* | \mathbf{x}_*) = \mathcal{N}(y_* | \mu_f(\mathbf{x}_*), \sigma_f(\mathbf{x}_*)^2 + \sigma_{\text{obs}}^2) \quad (14)$$

where $\mu_f(\mathbf{x}_*)$ is the predictive mean function in Eqn. 7 and $\sigma_f(\mathbf{x}_*)^2$ is the latent function variance in Eqn. 8. The predictive distribution for the DGP is an immediate generalization of Eqn. 14. In particular in the DGP case the predictive distribution is given by a *continuous mixture* of Normal distributions of the form in Eqn. 14. In more detail, for the 2-layer DGP in Eqn. 12, the predictive distribution is given by

$$\mathbb{E}_{\prod_{w=1}^W q(g_{*w} | \mathbf{x}_*)} [\mathcal{N}(y_* | \mu_f(\mathbf{g}_*), \sigma_f(\mathbf{g}_*)^2 + \sigma_{\text{obs}}^2)] \quad (15)$$

Since this expectation is intractable, in practice it must be approximated with Monte Carlo samples, i.e. as a finite mixture of Normal distributions.

2.4 PPGPR

In this section we review a recently introduced approach to scalable GP regression (PPGPR; Jankowiak et al. (2020)), as it will serve as motivation for Deep Sigma Point Processes in Sec. 3. In PPGPR one takes the family of predictive distribution in Eqn. 14—parameterized by $\mathbf{m}, \mathbf{S}, \mathbf{Z}$, and the kernel hyperparameters—as the model class, and fits the model using gradient-based maximum likelihood estimation. Jankowiak et al. (2020) argue that this class of models achieves good performance because the training objective directly targets the distribution used to make predictions at test time. In more detail the PPGPR

³Since e.g. $\Sigma_f(\mathbf{X})_{ii} = \sigma_f(\mathbf{x}_i)^2$ only depends on \mathbf{x}_i .

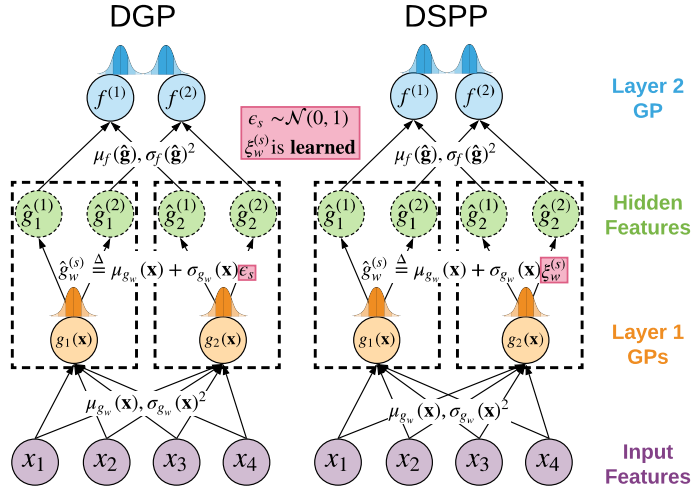


Figure 1: We depict the computational flow of a 2-layer DGP (left; see Sec. 2.3) and a 2-layer DSPP (right; see Sec. 3). Input features \mathbf{x} are fed through a set of W (here $W = 2$) Gaussian Processes by computing $\mu_{g_w}(\mathbf{x})$ and $\sigma_{g_w}(\mathbf{x})$ (Eqns. 7 and 8). Hidden features for the next GP layer (f) are computed from each $\mu_{g_w}(\mathbf{x})$ and $\sigma_{g_w}(\mathbf{x})$ by sampling from $\mathcal{N}(\mu_{g_w}(\mathbf{x}), \sigma_{g_w}(\mathbf{x})^2)$ in the DGP case, or by applying one of the quadrature rules discussed in Sec. 3.2 (here QR3), with the $\xi_w^{(s)}$ treated as learnable parameters. The final predictive distribution is a mixture of S Gaussians, where each Gaussian depends on one of the S sampled feature sets in the DGP case, or on one of the S deterministic quadrature-dependent feature sets in the DSPP case. The DGP is trained by gradient descent on $\mathcal{L}_{\text{dsvi}}$ (Eqn. 13), and the DSPP is trained via $\mathcal{L}_{\text{dspp}}$ (Eqn. 21). See Sec. 3.3 for more discussion on the relationship between DGPs and DSPPs.

objective is given by

$$\begin{aligned} \mathcal{L}_{\text{ppgpr}} &= \sum_{i=1}^N \log p(y_i | \mathbf{x}_i) - \beta_{\text{reg}} \text{KL}(q(\mathbf{u}) | p(\mathbf{u})) \\ &= \sum_{i=1}^N \log \mathcal{N}(y_i | \mu_f(\mathbf{x}_i), \sigma_f(\mathbf{x}_i)^2 + \sigma_{\text{obs}}^2) \\ &\quad - \beta_{\text{reg}} \text{KL}(q(\mathbf{u}) | p(\mathbf{u})) \end{aligned} \quad (16)$$

where $\beta_{\text{reg}} > 0$ is a hyperparameter and $\text{KL}(q(\mathbf{u}) | p(\mathbf{u}))$ serves as a regularizer. Note that the objective in Eqn. 16 looks deceptively similar to the SVGP objective in Eqn. 6; the crucial difference is where $\sigma_f(\mathbf{x}_i)^2$ appears. This difference is a result of the fact that the PPGPR objective directly targets the predictive distribution, while the SVGP objective targets minimizing a KL divergence between the variational distribution and the exact posterior. In SVGP the posterior predictive is then formed in a second step and does not directly enter the training procedure. PPGPR will serve as one of our baselines in Sec. 5.

3 DEEP SIGMA POINT PROCESSES

We now describe the class of models that is the focus of this work. Our approach is motivated by the good empirical performance exhibited by PPGPR, a class of parametric GP models explored in (Jankowiak et al., 2020) and reviewed in Sec. 2.4. Crucially, this approach to scalable GP regression relies on an objective function that is formulated in terms of the predictive distribution. We would like to apply this approach to the DGP setting but, unfortunately, the form of the predictive distribution

for DGPs (see Eqn. 15) presents an immediate obstacle: Eqn. 15 is a continuous mixture of Normal distributions and cannot be computed in closed form. In other words, the analog of the PPGPR objective in Eqn. 16 would involve the logarithm of the expectation in Eqn. 15. While this quantity could be approximated with a Monte Carlo estimator, because of the outer logarithm the result would be a biased estimator. Instead of trying to address this obstacle directly and construct an unbiased (gradient) estimator, we instead adopt a simpler solution: we replace the continuous mixture with a (parametric) *finite* mixture.⁴

To define a finite parametric family of mixture distributions, we essentially apply a sigma point approximation or quadrature-like integration rule to Eqn. 15. To make this concrete, suppose the width of the first GP layer in Eqn. 15 is $W = 2$. Then for an input \mathbf{x}_i we can write

$$p_{\text{dspp}}(y_i | \mathbf{x}_i) = \int d\mathbf{g}_i \mathcal{N}(y_i | \mu_f(\mathbf{g}_i), \sigma_f(\mathbf{g}_i)^2 + \sigma_{\text{obs}}^2) \prod_{w=1}^2 q(g_{iw} | \mathbf{x}_i) \quad (17)$$

The simplest ansatz for converting Eqn. 17 into a finite mixture uses Gauss-Hermite quadrature, i.e. we approximate $q(g_{i1} | \mathbf{x}_i)$ with a S -component mixture of Dirac delta distributions controlled by weights $\omega_1^{(s)}$ and quadra-

⁴In Sec. C we report empirical results comparing to the ‘direct’ (biased) approach in which the predictive distribution is given as a continuous mixture. We find that this model variant is outperformed by the DSPP, presumably because of the additional model flexibility that results from the learned quadrature rules we adopt in Sec. 3.2.

ture points $\xi_1^{(s)}$

$$q(g_{i1}|\mathbf{x}_i) = \sum_{s_1=1}^S \omega_1^{(s_1)} \delta\left(g_{i1} - \left(\mu_{g_1}(\mathbf{x}_i) + \xi_1^{(s_1)} \sigma_{g_1}(\mathbf{x}_i)\right)\right) \quad (18)$$

with an analogous ansatz for $q(g_{i2}|\mathbf{x}_i)$. For example for $S = 3$ we would have

$$\begin{aligned} \xi^{(1)} &= -\sqrt{3} & \xi^{(2)} &= 0 & \xi^{(3)} &= \sqrt{3} \\ \omega^{(1)} &= \frac{1}{6} & \omega^{(2)} &= \frac{2}{3} & \omega^{(3)} &= \frac{1}{6} \end{aligned} \quad (19)$$

for both g_{i1} and g_{i2} . Making these replacements in Eqn. 17 then yields a mixture with S^2 components:

$$p_{\text{dspp}}(y_i|\mathbf{x}_i) = \sum_{s_1}^S \sum_{s_2}^S \omega_1^{(s_1)} \omega_2^{(s_2)} \times \quad (20)$$

$$\mathcal{N}(y_i|\mu_f(\mu_{g_1}(\mathbf{x}_i) + \xi_1^{(s_1)} \sigma_{g_1}(\mathbf{x}_i), \mu_{g_2}(\mathbf{x}_i) + \xi_1^{(s_2)} \sigma_{g_2}(\mathbf{x}_i)), \sigma_f(\mu_{g_1}(\mathbf{x}_i) + \xi_1^{(s_1)} \sigma_{g_1}(\mathbf{x}_i), \mu_{g_2}(\mathbf{x}_i) + \xi_1^{(s_2)} \sigma_{g_2}(\mathbf{x}_i)))$$

Thus for a 2-layer model where the first GP layer has width W this particular quadrature rule leads to a predictive distribution that is a mixture of S^W Normal distributions, each of which is parameterized by composition of mean and variance functions of the form in Eqn. 7 and Eqn. 8. This exponential growth in the number of mixture components is potentially problematic; we defer a more detailed discussion of alternative—in particular more compact—quadrature rules to Sec. 3.2.

3.1 TRAINING OBJECTIVE

Now that we have defined the class of parametric regression models we are interested in, we can define our training objective. As in Jankowiak et al. (2020), we define an objective function that corresponds to regularized maximum likelihood estimation

$$\mathcal{L}_{\text{dspp}} = \sum_{i=1}^N \log p_{\text{dspp}}(y_i|\mathbf{x}_i) - \beta_{\text{reg}} \sum \text{KL} \quad (21)$$

where $\beta_{\text{reg}} > 0$ is an optional regularization constant and $\sum \text{KL}$ is a sum over KL divergences of inducing variables (one for each GP) just as in the DGP objective, Eqn. 13. Just as in ‘Doubly Stochastic Variational Inference’ for DGPs (Sec. 2.3.1), this objective—which depends on σ_{obs} as well as parameters \mathbf{m} , \mathbf{S} , \mathbf{Z} and various kernel hyperparameters for each GP—can be optimized using stochastic gradient methods. In contrast to DGPs, this optimization is only ‘singly stochastic,’ i.e. we only subsample data points and not latent function values.

3.2 QUADRATURE RULES

The Gauss-Hermite quadrature rule used to motivate DSPPs in Eqn. 18 is intuitive, but is of limited practical use, since it leads to an exponential blow-up in the number of mixture components used to define the DSPP. It is therefore essential to consider alternative quadrature rules that are more compact. We emphasize at the outset that our goal is *not* to accurately estimate the intractable expectation in Eqn. 17. Rather, our goal is to construct a flexible family of parametric distributions that are governed by well-behaved mean and variance functions that benefit from compositional structure. Consequently, there is no need to restrict ourselves to quadrature rules derived from gaussian integrals—indeed empirically we find that the quadrature rule in Eqn. 18 is too rigid. We now describe three more flexible alternatives.

QR1 In the first quadrature rule we choose quadrature points that follow the same factorized structure that is evident in the double sum in Eqn. 20, i.e. we still use a grid of S^W quadrature points for a 2-layer DSPP where the first GP layer has width W . That is we make the substitution

$$\begin{aligned} \prod_{w=1}^W q(g_{iw}|\mathbf{x}_i) &\rightarrow \quad (22) \\ \sum_{\mathbf{s}} \omega^{(\mathbf{s})} \prod_{w=1}^W \delta\left(g_{iw} - \left(\mu_{g_w}(\mathbf{x}_i) + \xi_w^{(s_w)} \sigma_{g_w}(\mathbf{x}_i)\right)\right) \end{aligned}$$

where $\mathbf{s} = (s_1, \dots, s_W) \in \{1, \dots, S\}^W$ is a multi-index. In contrast to the Gauss-Hermite quadrature rule, however, we choose the quadrature points $\{\xi_w^{(s_w)}\}$ to be learnable parameters (for a total of $S \times W$ real parameters). In addition we replace the Gauss-Hermite weights for each mixture—which are given as a product over the W GPs in the first layer, i.e. $\prod_w \omega_w^{(s_w)}$ for each multi-index \mathbf{s} —by S^W -many learnable parameters that sum to unity $\{\omega^{(\mathbf{s})}\}$.

QR2 The second rule is identical to QR1 except the quadrature points are forced to be symmetric, i.e. $\xi_w^{(s_w)} = -\xi_w^{(S+1-s_w)}$ for $s_w = 1, \dots, S$ and $w = 1, \dots, W$.

QR3 In the third quadrature rule we abandon the factorized structure of QR1 and QR2, thus liberating us from the exponential growth of mixture components. To do this we effectively ‘line-up’ the quadrature points across the different GPs g_1, \dots, g_W , making the substitution

$$\begin{aligned} \prod_{w=1}^W q(g_{iw}|\mathbf{x}_i) &\rightarrow \quad (23) \\ \sum_{s=1}^S \omega^{(s)} \prod_{w=1}^W \delta\left(g_{iw} - \left(\mu_{g_w}(\mathbf{x}_i) + \xi_w^{(s)} \sigma_{g_w}(\mathbf{x}_i)\right)\right) \end{aligned}$$

where there is a set of S learnable quadrature weights $\{\omega^{(s)}\}_{s=1}^S$ and S learnable quadrature points $\{\xi_w^{(s)}\}_{s=1}^S$ defined by $S \times W$ real parameters. Here S is a parameter that we control; in particular it has no relationship to W and can be as small as $S = 1$. We explore the empirical performance of these quadrature rules in Sec. 5.1.

3.3 DISCUSSION

We use this section to clarify the relationship between variational DGPs and DSPPs. DSPPs differ from DGPs in two important respects (also see Fig. 1):

1. **Objective function:** The DGP is trained with an ELBO (Eqn. 13) and the DSPP is trained via a regularized maximum likelihood objective (Eqn. 21) that directly targets the DSPP predictive distribution.
2. **Treatment of latent function values:** In the DGP latent function values not at the top of the hierarchy (e.g. \mathbf{G} in Eqn. 12) are *sampled* while in the DSPP they are *parameterized* via a learnable quadrature rule as in Eqn. 23.

Indeed this latter point is made explicit by our quadrature rules—see Eqn. 22 and Eqn. 23—which should be compared to the reparameterization trick used during DGP training, which implicitly makes the substitution

$$q(g_{iw}|\mathbf{x}_i) \rightarrow \mathbb{E}_{\mathcal{N}(\epsilon|0,1)} \left[\delta(g_{iw} - (\mu_{g_w}(\mathbf{x}_i) + \epsilon\sigma_{g_w}(\mathbf{x}_i))) \right]$$

and approximates the expectation using Monte Carlo samples $\{\epsilon_s\}$ with $\epsilon_s \sim \mathcal{N}(\cdot|0, 1)$.

Note that in other respects variational DGPs and DSPPs are very similar. For example, apart from the parameters defining the learned quadrature rule in a DSPP, they make use of the same parameters. Similarly, their predictive distributions have similar forms, with the difference that the DSPP utilizes a finite mixture of Normal distributions, while the DGP predictive distribution is a continuous mixture of Normal distributions.

4 RELATED WORK

We discuss some work related to DSSPs, noting that we review some of the relevant literature in Sec. 2. The use of pseudo-inputs and inducing point methods to scale-up Gaussian Process inference has spawned a large literature, especially in the context of variational inference (Snelson and Ghahramani, 2006; Titsias, 2009; Hensman et al., 2013; Cheng and Boots, 2017). While variational inference remains the most popular inference algorithm for scalable GPs, researchers have also explored different

Table 1: Average ranking of different quadrature rules (further to the left is better). CRPS is the Continuous Ranked Probability Score, a popular calibration metric for regression (Gneiting and Raftery, 2007). Rankings are aggregated across the smallest 8 UCI datasets and train/test/validation splits. See Sec. 5.1 for details.

	DSPP-QR1	DSPP-QR2	DSPP-QR3
NLL	2.01	2.44	1.37
RMSE	1.79	2.27	1.62
CRPS	1.73	2.51	1.51

Table 2: Average ranking of methods in Sec. 5.2 (lower is better). Rankings are aggregated across all 120 pairs of dataset and train/test/validation split.

	OD-SVGP	PPGPR	DGP	γ -DGP	DSPP
NLL	4.38	2.33	3.55	3.60	1.15
RMSE	3.28	2.85	2.98	3.32	2.58
CRPS	4.42	2.47	3.66	2.93	1.52

variants of Expectation Propagation (Hernández-Lobato and Hernández-Lobato, 2016) as well as Stochastic gradient Hamiltonian Monte Carlo (Havasi et al., 2018).

Deep Gaussian Processes were introduced in (Damianou and Lawrence, 2013), with recent approaches to variational inference for DGPs described by Salimbeni and Deisenroth (2017). Cutajar et al. (2017) introduce a hybrid model formulated with random feature expansions that combines features of DGPs and neural networks. This class of models bears some resemblance to DSPPs; this is especially true for the VAR-FIXED variant, in which spectral frequencies are treated deterministically. Importantly, since Cutajar et al. (2017) rely on variational inference, their training objective does not directly target the predictive distribution. As we show empirically in Sec. 5, the mismatch between the training objective and test time predictive distributions for variational DGPs severely degrades predictive performance; we expect this is also true of the approach in (Cutajar et al., 2017).

DSPPs can also be motivated by Direct Loss Minimization, which emerges from a view of approximate inference as regularized loss minimization (Sheth and Khordon, 2017). This connection is somewhat loose, however, since our fully parametric models dispose of approximate posterior (or quasi-posterior) distributions entirely.

5 EXPERIMENTS

In this section we explore the empirical performance of the Deep Sigma Point Process introduced in Sec. 3. Throughout DSPPs use diagonal (i.e. mean field) covari-




Figure 2: We depict negative log likelihoods (NLLs) for the 12 univariate regression datasets in Sec. 5.2 (further to the left is better). Results are averaged over ten random train/test/validation splits. Here and throughout uncertainty bars depict standard errors.




Figure 3: We depict root mean squared errors (RMSEs) for the 12 univariate regression datasets in Sec. 5.2 (further to the left is better). Results are averaged over ten random train/test/validation splits.




Figure 4: We depict the Continuous Ranked Probabilistic Score (CRPS; Gneiting and Raftery (2007)) for the 12 univariate regression datasets in Sec. 5.2 (further to the left is better). CRPS is a popular calibration metric for regression. Results are averaged over ten random train/test/validation splits.

ance matrices S . Except for Sec. 5.4 where we consider 3-layer models, all DGPs and DSPPs considered in our experiments have two layers. For details about kernels, mean functions, layer widths, numbers of inducing points, etc., refer to Sec. A in the appendix.

5.1 QUADRATURE RULE ABLATION STUDY

We begin by exploring the impact of the three different quadrature rules defined in Sec. 3.2. To do so we train 2-layer DSPPs on the 8 smallest univariate regression datasets described in the next section. In particular we compare QR1 and QR2 with⁵ $S = 3$ to QR3 with $S = 10$. The results are summarized in Table 1, with more detailed results to be found in Sec. C in the appendix. The upshot is that the three quadrature rules give largely comparable performance, with some preference for the most flexible quadrature rule, QR3. Since this quadrature rule avoids exponentially many quadrature points—and the compu-

⁵Since we consider $W \in \{3, 4\}$ this corresponds to $S^W \in \{27, 81\}$ mixture components.

tational cost can be carefully controlled by choosing S appropriately—we use QR3 in the remainder of our experiments. Indeed this choice is essential for training 3-layer DSPPs in Sec. 5.4, which would otherwise be too computationally expensive.

5.2 UNIVARIATE REGRESSION

We reproduce a univariate regression experiment described in (Jankowiak et al., 2020). In particular we consider consider twelve datasets from the UCI repository (Dua and Graff, 2017), with the number of data points in the range $10^4 \lesssim N \lesssim 10^6$ and the number of input dimensions in the range $3 \leq \text{Dim}(\mathbf{x}) \leq 380$.

We consider four strong GP baselines—two single-layer models: (OD-SVGP) the orthogonal basis decoupling method described in Cheng and Boots (2017) and Salimbeni et al. (2018); (PPGPR) the Parametric Predictive GP regression method described in Sec. 2.4; as well as two 2-layer models: (DGP) a variational DGP as described in Sec. 2.3; and (γ -DGP) the robust DGP described in

Table 3: Average ranking of methods in Sec. 5.3 (lower is better). Rankings are aggregated across all 25 pairs of dataset and train/test/validation split.

	SVGP	PPGPR	DGP	γ -DGP	DSPP
NLL	4.16	2.20	4.24	3.40	1.00
MRMSE	2.52	4.72	2.04	3.60	2.12

Knoblauch (2019); Knoblauch et al. (2019).⁶ Results for OD-SVGP and PPGPR are reproduced from (Jankowiak et al., 2020).

Our results are summarized in Figs. 2-4 and Table 2. We find that in aggregate the DSPP outperforms all the baselines in terms of log likelihood, RMSE, and CRPS (also see Table 6 in Sec. C in the appendix). In particular, averaging across all twelve datasets, the DSPP outperforms the DGP by ~ 0.75 nats and the PPGPR by ~ 0.47 nats w.r.t. log likelihood. Strikingly, the second strongest baseline is the (single-layer) PPGPR described in Sec. 2.4. The fact that the PPGPR is able to outperform a 2-layer DGP highlights the advantages of a training procedure that directly targets the predictive distribution. Since the DSPP is in effect a much more flexible version of the PPGPR, it yields even better predictive performance, especially with respect to log likelihood. Indeed while the DGP achieves good RMSE performance on most datasets, posterior approximations degrade the calibration of the test time predictive distribution (as measured by log likelihood and CRPS).

5.3 MULTIVARIATE REGRESSION

We conduct a multivariate regression experiment using five robotics datasets, two of which were collected from real-world robots and three of which were generated using the MuJoCo physics simulator (Todorov et al., 2012). In all five datasets the input and output dimensions correspond to various joint positions/velocities/etc. of the robot, with the number of data points in the range $10^4 \lesssim N \lesssim 10^5$, the number of input dimensions in the range $10 \leq \text{Dim}(\mathbf{x}) \leq 23$, and the number of output dimensions in the range $7 \leq \text{Dim}(\mathbf{y}) \leq 10$. These datasets have been used in a number of papers, including (Vijayakumar and Schaal, 2000; Cheng and Boots, 2017).

All of our models follow the structure of the ‘linear model of coregionalization’ (Alvarez et al., 2012), specifically along the lines of ‘Semiparametric latent factor models’

⁶This robust DGP can be viewed as a variant of the DGP described in Sec. 2.3 in which the inference procedure is modified by replacing the expected log likelihood loss with a gamma divergence in which the likelihood is raised to a power: $\log p(\mathbf{y}|\mathbf{f}) \rightarrow p(\mathbf{y}|\mathbf{f})^{\gamma-1}$.

(Seeger et al., 2005); for more details see Sec. A.3 in the supplementary materials.

We consider four strong GP baselines. In particular we consider two single-layer models: (SVGP) as described in Sec. 2.2.1 and (Hensman et al., 2013); (PPGPR) the Parametric Predictive GP regression method described in Sec. 2.4; as well as two 2-layer models: (DGP) a variational DGP as described in Sec. 2.3; and (γ -DGP) the robust DGP described in Knoblauch (2019); Knoblauch et al. (2019).

Our results are summarized in Fig. 5 and Table 3; see Sec. C in the appendix for additional results. We find that the DSPP outperforms all the baselines w.r.t. NLL, and achieves comparable MRMSE performance to the DGP. Averaged across all five datasets, the DSPP outperforms the DGP by ~ 0.82 nats and PPGPR by ~ 0.17 nats w.r.t. log likelihood. Note that while the PPGPR achieves good performance on log likelihoods, its MRMSE performance is substantially worse than the DSPP.

5.4 MULTILAYER MODELS

In the above experiments we have limited ourselves to 2-layer DSPPs. Here we investigate the predictive performance of 3-layer models, in particular comparing to 3-layer DGPs. Our results for four univariate regression datasets are summarized in Fig. 6. For both DGPs and DSPPs we find that, depending on the dataset, adding a third layer can improve NLLs, but that these gains are somewhat marginal compared to the gains from moving from single-layer to two-layer models. DSPPs exhibit the best log likelihoods, while RMSE results are somewhat more mixed, with the DSPP and PPGPR obtaining the lowest RMSE depending on the dataset.

5.5 COMPARISON WITH DEEP KERNEL LEARNING

From the previous sets of results we see that DSPP models tend to outperform their single-layer counterparts, both in terms of NLL and RMSE. A natural question is whether or not these performance improvements can be obtained with other classes of hierarchical models. More specifically, could we achieve similar results if the hierarchical features are extracted using traditional neural networks rather than GPs? To answer this question, we consider deep kernel learning (DKL) variants of SVGP, PPRPR, DGP, and DSPP in which the input features to each model are extracted by a neural network (Calandra et al., 2016; Wilson et al., 2016). The neural network parameters are optimized end-to-end alongside the GP/DGP/DSPP parameters. We conduct experiments on 4 medium-sized UCI regression datasets. For all DKL models we use a




Figure 5: We depict NLLs (left) and mean root mean squared errors (MRMSEs, right)—i.e. RMSEs averaged across all output dimensions—for the 5 multivariate regression datasets in Sec. 5.3 (lower is better). Results are averaged over five random train/test/validation splits. Note that NLLs are normalized by the number of output dimensions.




Figure 6: We depict NLLs (left) and RMSEs (right) for the multi-layer experiment in Sec. 5.4 (lower is better). Results for 1L and 2L (respectively, 3L) models are averaged over ten (respectively, five) random train/test/validation splits.




Figure 7: We compare NLLs (left) and RMSEs (right) of neural-network modulated (deep kernel learning) variants of GP/DGP/DSPP models (lower is better). Results are presented for 4 of the UCI regression datasets in Sec. 5.2, averaged over ten random train/test/validation splits. See Sec. 5.5 for details.

5-layer neural network architecture proposed by Wilson et al. (2016) to extract features (see Sec. A.5 for details).

In Fig. 7 we compare the neural network modulated models (denoted by “NN+”) with their standard counterparts. We find that adding a neural network to DSPP models has limited impact on performance, improving on some datasets and not on others. For SVGP/PPGPR/DGP models, neural networks improve RMSE but have mixed effects on NLL. On three of the four datasets, none of the NN + SVGP/PPGPR/DGP models matches the NLL of the standard DSPP. This is particularly notable for the NN+PPGPR model, as it only differs from the standard DSPP model in terms of its “feature extractor” layer (neural networks versus GP/quadrature). These results suggest that hidden GP layers in DSPPs extract features that

are complementary to those extracted by neural networks, while enjoying favorable regularization properties.

6 DISCUSSION

We motivated Deep Sigma Point Processes as a finite family of parametric distributions whose structure mirrors the DGP predictive distribution in Eqn. 15. It would be interesting to consider model variants that retain a continuous mixture distribution as in Eqn. 15 while leveraging the flexibility inherent in the learned quadrature rules described in Sec. 3.2. Another open question is whether DSPPs can be fruitfully applied to other likelihoods, for example those that arise in classification. We leave the exploration of these directions to future work.

References

- Mauricio A Alvarez, Lorenzo Rosasco, Neil D Lawrence, et al. Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3): 195–266, 2012.
- Roberto Calandra, Jan Peters, Carl Edward Rasmussen, and Marc Peter Deisenroth. Manifold gaussian processes for regression. In *International Joint Conference on Neural Networks*, 2016.
- Ching-An Cheng and Byron Boots. Variational inference for gaussian process models with linear complexity. In *Advances in Neural Information Processing Systems*, 2017.
- Kurt Cutajar, Edwin V Bonilla, Pietro Michiardi, and Maurizio Filippone. Random feature expansions for deep gaussian processes. In *International Conference on Machine Learning*, 2017.
- Andreas Damianou and Neil Lawrence. Deep gaussian processes. In *Artificial Intelligence and Statistics*, 2013.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Marton Havasi, José Miguel Hernández-Lobato, and Juan José Murillo-Fuentes. Inference in deep gaussian processes using stochastic gradient hamiltonian monte carlo. In *Advances in Neural Information Processing Systems*, 2018.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, 2013.
- Daniel Hernández-Lobato and José Miguel Hernández-Lobato. Scalable gaussian process classification via expectation propagation. In *Artificial Intelligence and Statistics*, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, 2015.
- Martin Jankowiak, Geoff Pleiss, and Jacob R Gardner. Parametric gaussian process regressors. In *International Conference on Machine Learning*, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learned Representations*, 2015.
- Jeremias Knoblauch. Robust deep gaussian processes. *arXiv preprint arXiv:1904.02303*, 2019.
- Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. Generalized variational inference. *arXiv preprint arXiv:1904.02063*, 2019.
- Robert Price. A useful theorem for nonlinear devices having gaussian inputs. *IRE Transactions on Information Theory*, 4(2):69–72, 1958.
- Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.
- Tim Salimans, David A Knowles, et al. Fixed-form variational posterior approximation through stochastic linear regression. *Bayesian Analysis*, 8(4):837–882, 2013.
- Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep gaussian processes. In *Advances in Neural Information Processing Systems*, 2017.
- Hugh Salimbeni, Ching-An Cheng, Byron Boots, and Marc Deisenroth. Orthogonally decoupled variational gaussian processes. In *Advances in neural information processing systems*, 2018.
- Matthias Seeger, Yee-Whye Teh, and Michael Jordan. Semiparametric latent factor models. Technical report, 2005.
- Rishit Sheth and Roni Khardon. Excess risk bounds for the bayes risk using variational inference in latent gaussian models. In *Advances in Neural Information Processing Systems*, 2017.
- Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, 2006.
- Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, 2009.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *International Conference on Intelligent Robots and Systems*, 2012.
- Sethu Vijayakumar and Stefan Schaal. Locally weighted projection regression: An $o(n)$ algorithm for incremental real time learning in high dimensional space. In *International Conference on Machine Learning*, 2000.
- Andrew G Wilson, Zhiting Hu, Russ R Salakhutdinov, and Eric P Xing. Stochastic variational deep kernel learning. In *Advances in Neural Information Processing Systems*, 2016.