
Probabilistic Selection of Inducing Points in Sparse Gaussian Processes

Anders Kirk Uhrenholt¹

Valentin Charvet¹

Bjørn Sand Jensen¹

¹School of Computing Science,
University of Glasgow,
United Kingdom.

Abstract

Sparse Gaussian processes and various extensions thereof are enabled through inducing points, that simultaneously bottleneck the predictive capacity and act as the main contributor towards model complexity. However, the number of inducing points is generally not associated with uncertainty which prevents us from applying the apparatus of Bayesian reasoning for identifying an appropriate trade-off. In this work we place a point process prior on the inducing points and approximate the associated posterior through stochastic variational inference. By letting the prior encourage a moderate number of inducing points, we enable the model to learn which and how many points to utilise. We experimentally show that fewer inducing points are preferred by the model as the points become less informative, and further demonstrate how the method can be employed in deep Gaussian processes and latent variable modelling.

1 INTRODUCTION

Gaussian processes (GP) constitute an attractive modelling tool when the amount of data is limited and/or uncertainty quantification is critical, e.g. global optimisation (Moćkus [1975], Shahriari et al. [2015]), medical modelling (Lorenzi et al. [2015]), and reinforcement learning (Deisenroth and Rasmussen [2011]). While the original formulation presents shortcomings in scalability and expressiveness, a vast amount of enhancements have been presented over the years that enable large-scale modelling (Hensman et al. [2013]), deep architectures (Damianou and Lawrence [2013]), inter-domain covariance mappings (Lázaro-Gredilla and Figueiras-Vidal [2009]), and various combinations thereof (Blomqvist et al. [2019]).

Fundamental to these advancements is the methodology of

sparse Gaussian processes (Snelson and Ghahramani [2006], Quiñonero-Candela and Rasmussen [2005], Titsias [2009]). Here, we introduce the assumption that the true posterior over functions may be adequately approximated by conditioning on a relatively small set of *inducing points*, acting as a representative proxy for the observed data (see Figure 1a). The task of model optimisation now reduces to quantifying and minimising the discrepancy between the approximate posterior and that of the full GP. Since the inducing points communicate all information between observations and predictions, it is crucial that they are initialised and optimised to provide good coverage of the input domain (Burt et al. [2019]). At the same time, they are the main contributors to the complexity of the model, and so both the number and locations of inducing points must be selected with care.

While well-developed and long underway, the inducing point methodology has always viewed the number of points as a design choice. We may add or remove inducing points adaptively and optimise their locations in input space, but the set of points ultimately appear as a deterministic quantity in the model specification. This has the immediate consequence that the question of which and how many inducing points to utilise is not part of the inference. As such, we cannot “learn” how many points to include in order to balance capacity and complexity, and there is no principled way of identifying inducing points that may not be contributing much towards explaining the data.

In this work we present a solution to the above limitations by including the *selection* of inducing points in the inference. Concretely, we expand the hierarchical model by first sampling the inducing points from a point process prior, which in a Bayesian fashion encodes the beliefs we have about our point set – namely that the model should be economical in its use of points. We then approximate the true posterior with a variational point process (Figure 1b) that must find a trade-off between adhering to the expectations of the prior and explaining the data. As a result the model learns to sample only those points that it deems sufficiently informative, yielding a data-driven approach for optimally

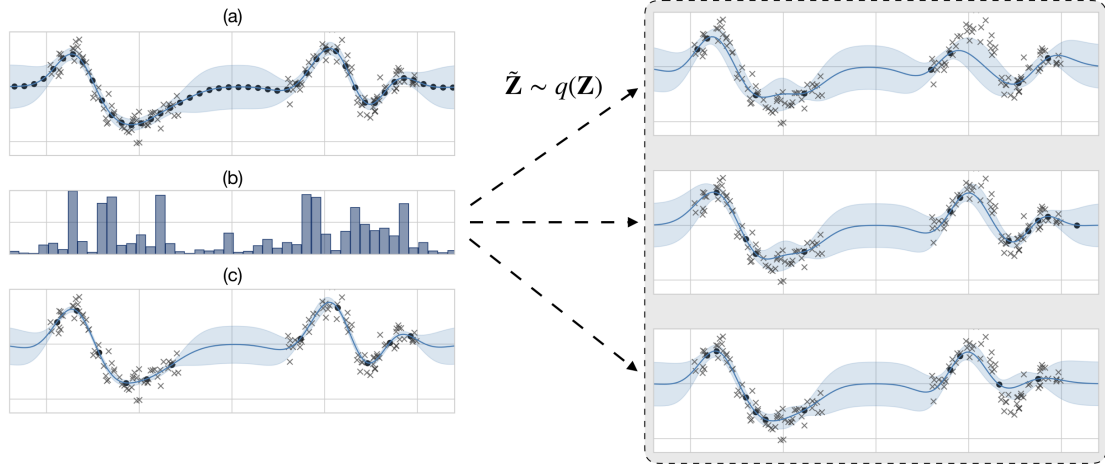


Figure 1: Illustration of our method. A sparse GP infers an approximate posterior (blue shading) over observed data (grey crosses) by conditioning on a set of inducing points (black circles). (a) Posterior when using an initial candidate set of 50 inducing points. (b) Marginal probability of inclusion assigned to each inducing point by the variational point process midway through training. Note that uninformative points are less likely to be included. Right panel shows three samples from the point process. (c) Result after training where a small set of highly informative inducing points remain.

exploiting the available resources (Figure 1c).

We present a prior that naturally encourages selectivity while allowing for an analytically tractable expression for the Kullback-Leibler divergence when certain statistical measures are known about the variational posterior. The evidence lower bound can then be maximised through score function estimation (Williams [1992], Fu [2006]), making it straightforward to implement in any existing Gaussian process framework.¹

We demonstrate on various datasets that when reducing the informativeness of inducing points, our model does indeed select fewer points while achieving comparable or better performance than standard approaches. We then consider the problem of allocating inducing points to the layers of a deep Gaussian process, and show that we can solve this otherwise combinatorial problem by jointly inferring the posterior set of points for each layer. Finally, we apply the method to the Gaussian process latent variable model where we learn a representation for a high-dimensional single cell dataset whilst adapting the suitable number of inducing points throughout optimisation, thus avoiding costly cross-validation.

2 BACKGROUND

2.1 GAUSSIAN PROCESSES

We consider the supervised learning setting of having N datapoints, $(\mathbf{y}, \mathbf{X}) = \{y_i, \mathbf{x}_i\}_{i=1}^N$, that are assumed condi-

tionally independent given a latent function f :

$$p(\mathbf{y} | f, \mathbf{X}) = \prod_{i=1}^N p(y_i | f(\mathbf{x}_i)).$$

By placing a Gaussian process prior (Rasmussen and Williams [2006]) on the latent function, $f \sim \mathcal{GP}(m(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}'))$, we assert that any finite set of evaluations, $\mathbf{f} = \{f(\mathbf{x}_i)\}_{i=1}^N$, follows a multivariate Gaussian with statistics given by mean function, m , and covariance function, κ :

$$p(\mathbf{f} | \mathbf{X}) = \mathcal{N}(\mathbf{f} | \mu, \Sigma),$$

$$\mu_i = m(\mathbf{x}_i; \theta), \quad \Sigma_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j; \theta).$$

Here θ comprises the hyper-parameters of m and κ that characterise the Gaussian process. When the data distribution, $p(\mathbf{y} | \mathbf{f})$, is an isotropic normal, the marginal likelihood is available in closed form, providing a means for model selection and a (differentiable) objective for estimating θ and the observation noise of $p(\mathbf{y} | \mathbf{f})$. Predicting for unseen data, $(\mathbf{y}_*, \mathbf{X}_*)$, is carried out by conditioning on the training data and marginalising out the latent function, yielding a tractable expression for $p(\mathbf{y}_* | \mathbf{X}_*, \mathbf{y}, \mathbf{X})$.

2.2 SPARSE VARIATIONAL GAUSSIAN PROCESSES

As a result of marginalising out the latent function, the datapoints are no longer independent leading to a model that scales poorly in the number of observations. This has motivated the development of sparse Gaussian processes (Snelson and Ghahramani [2006], Quiñonero-Candela and Rasmussen [2005]), in which we only condition on a set

¹Implementation is available at https://github.com/akuhren/selective_gp.

of M inducing points, $(\mathbf{u}, \mathbf{Z}) = \{u_j, \mathbf{z}_j\}_{j=1}^M$, $u_j = f(\mathbf{z}_j)$. These points are not restricted to being part of the observed data and can thus be optimised so as to best represent \mathbf{X} and \mathbf{y} . We will adopt the notation of $\mathbf{f}_{\neq \mathbf{u}}$ being all evaluations except for \mathbf{u} , and $\mathbf{f} \triangleq [\mathbf{f}_{\neq \mathbf{u}}, \mathbf{u}]$. Modelling proceeds by making the simplifying assumption that the function posterior can be approximated through these inducing points, which for $M \ll N$ leads to a vast reduction in complexity. While various approaches encode this assumption in different manners, we will focus on the variational sparse formulation of Titsias [2009]. By introducing a variational distribution, Q , to approximate the true posterior, $p(\mathbf{f} | \mathbf{y}, \mathbf{X})$, the evidence lower bound (ELBO) is derived through Jensen’s inequality:

$$\log p(\mathbf{y} | \mathbf{X}) \geq \int Q \log \frac{p(\mathbf{y}, \mathbf{f}_{\neq \mathbf{u}} | \mathbf{X})}{Q} d\mathbf{f}_{\neq \mathbf{u}} \triangleq \mathcal{L}.$$

Next, we augment the probability space with inducing points (\mathbf{u}, \mathbf{Z}) .² Note that these points are already implicitly present in the generative model of $p(\mathbf{y} | \mathbf{X})$ through the infinite-dimensional vector, f , but since $\mathbf{f}_{\neq \mathbf{u}}$ is a sufficient statistic for \mathbf{y} , the inducing outputs are marginalised out with no effect to the distribution over observed data. However, by including the points in the variational distribution s.t. $Q = p(\mathbf{f}_{\neq \mathbf{u}} | \mathbf{X}, \mathbf{u}, \mathbf{Z})q(\mathbf{u} | \mathbf{Z})$ we obtain the much simplified expression:

$$\begin{aligned} \mathcal{L} &= \int Q \log \frac{p(\mathbf{y}, \mathbf{f} | \mathbf{X}, \mathbf{Z})}{p(\mathbf{f}_{\neq \mathbf{u}} | \mathbf{X}, \mathbf{u}, \mathbf{Z})q(\mathbf{u} | \mathbf{Z})} d\mathbf{f} \\ &= \int Q \log \frac{p(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}})p(\mathbf{f}_{\neq \mathbf{u}} | \mathbf{u}, \mathbf{X}, \mathbf{Z})p(\mathbf{u} | \mathbf{Z})}{p(\mathbf{f}_{\neq \mathbf{u}} | \mathbf{X}, \mathbf{u}, \mathbf{Z})q(\mathbf{u} | \mathbf{Z})} d\mathbf{f} \\ &= \mathbb{E}_Q [\log p(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}})] - \text{KL}[q(\mathbf{u} | \mathbf{Z}) || p(\mathbf{u} | \mathbf{Z})], \quad (1) \end{aligned}$$

where $\text{KL}[\cdot]$ is the Kullback-Leibler divergence. Our approximate posterior, $Q \approx p(\mathbf{f} | \mathbf{X}, \mathbf{y})$, thus relies on the initially redundant inducing points to make the function evaluations pertaining to \mathbf{y} conditionally independent, leading to a vast reduction in complexity.

In the special case of $p(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}}) = \mathcal{N}(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}}, \mathbf{I}\sigma_y^2)$ the collapsed bound is analytically tractable and yields the variational distribution $q(\mathbf{u}) = \mathcal{N}(\mathbf{u} | \mathbf{m}, \mathbf{S})$ with:

$$\begin{aligned} \mathbf{S}^{-1} &= \sigma_y^{-2} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} \mathbf{K}_{NM} \mathbf{K}_{MM}^{-1} + \mathbf{K}_{MM}^{-1}, \\ \mathbf{m} &= \sigma_y^{-2} \mathbf{S} \mathbf{K}_{MM}^{-1} \mathbf{K}_{MN} \mathbf{y}, \end{aligned}$$

where $\mathbf{K}_{MN} = \mathbf{K}_{NM}^\top = \kappa(\mathbf{Z}, \mathbf{X})$ and $\mathbf{K}_{MM} = \kappa(\mathbf{Z}, \mathbf{Z})$. However, this bound does not allow us to subsample (\mathbf{y}, \mathbf{X}) which in turn prohibits stochastic variational inference. To circumvent this limitation, Hensman et al. [2013] let (\mathbf{m}, \mathbf{S}) be free parameters. This further enables settings where $p(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}})$ is not Gaussian (Hensman et al. [2015a]), in

which case the expectation over conditionally independent data points can be estimated by Monte Carlo sampling:

$$\begin{aligned} \mathcal{L} &= \sum_{i=1}^M \mathbb{E}_{q(f_i; \mathbf{m}, \mathbf{S})} [\log p(y_i | f_i)] \\ &\quad - \text{KL}[q(\mathbf{u} | \mathbf{Z}; \mathbf{m}, \mathbf{S}) || p(\mathbf{u} | \mathbf{Z})], \\ q(f_i; \mathbf{m}, \mathbf{S}) &\sim \mathcal{N}(\beta_i \mathbf{m}, \kappa(\mathbf{x}_i, \mathbf{x}_i) - \beta_i (\mathbf{K}_{MM} - \mathbf{S}) \beta_i^T), \\ \beta_i &= \kappa(\mathbf{x}_i, \mathbf{Z}) \mathbf{K}_{MM}^{-1}. \end{aligned}$$

In this work we focus on both the collapsed and uncollapsed bounds and explore how the probabilistic model is affected by associating \mathbf{Z} with a point process. The sparse variational GP (SVGP) model is straightforwardly extended to more exotic variants such as deep Gaussian processes (Damianou and Lawrence [2013]) and latent variable models (Lawrence [2004], Titsias and Lawrence [2010]), which we also explore in our experiments. However, we defer the derivation of these to Appendix A.

3 PROBABILISTIC SELECTION OF INDUCING POINTS

In the SVGP model the inducing points communicate all information between observed data and new predictions while contributing $\mathcal{O}(M^2)$ time complexity in doing so. This implies that the main trade-off between model capacity and complexity hinges on the cardinality of \mathbf{Z} . In many scenarios this does not pose much of a dilemma. If, for instance, we are fitting a SVGP and know the amount of available training data, we would opt for using as many inducing points as possible while keeping within computational constraints. In other situations, however, the choice is less obvious. We might have a hierarchical model with various compositions of SVGP’s (e.g. Hamelijnck et al. [2019]), each of which will model response surfaces of varying complexity and thus require differing amounts of inducing points. Or we may be presented with batches of data in an online fashion (Bui et al. [2017]) and should ideally determine adaptively how many points to utilise for a given batch dependent on the characteristics of the data.

In such situations it seems unsatisfactory to have the cardinality of \mathbf{Z} be a deterministic choice. Rather, we ought to adopt a Bayesian approach and explicitly model any uncertainty we might have about the number of inducing points to utilise. We propose going about this by extending the generative model to include the sampling of inducing points from a point process prior, $p(\mathbf{Z})$, which assigns probability to sets belonging to its domain, \mathcal{Z} , based on their cardinality.

This is theoretically justified since, as mentioned in the previous section, the inducing points are implicitly present in the generative model before the variational approximation is made. We are thus free to put a prior on this quantity (as has been done in previous work, e.g. Hensman et al.

²Matthews et al. [2016] provide a rigorous analysis of the validity of the augmentation strategy and show it to be valid under very general conditions.

[2015b]) even though this prior does not affect the marginal distribution over \mathbf{y} . However, once the variational approximation is made the inducing points cease to be redundant and, consequently, $p(\mathbf{Z})$ ceases to be redundant as well.

Note that $p(\mathbf{Z})$ does *not* specify a distribution over continuous values; rather it is a discrete distribution over subsets drawn from \mathcal{Z} . The prior will in traditional Bayesian fashion encode our expectations about any finite set of points not belonging to the observed data, allowing for a natural way to encourage a discriminatory selection of inducing points.

To accommodate this hierarchical expansion we also update our proposal distribution, \mathcal{Q} , to include a variational point process, $q(\mathbf{Z})$:

$$\mathcal{Q} = q(\mathbf{f}, \mathbf{Z}) = p(\mathbf{f}_{\neq \mathbf{u}} | \mathbf{X}, \mathbf{u}, \mathbf{Z})q(\mathbf{u} | \mathbf{Z})q(\mathbf{Z}).$$

Here, $q(\mathbf{u} | \mathbf{Z})$ is just the marginal distribution conditioned on a set of inputs sampled from \mathcal{Z} . Under the collapsed bound this distribution is, as always, available in closed form. In the uncollapsed case, where $\mathcal{N}(\mathbf{u}^* | \mathbf{m}^*, \mathbf{S}^*)$ is the variational distribution pertaining to all points in \mathcal{Z} and $(\mathbf{m}^*, \mathbf{S}^*)$ are variational parameters, we appeal to the marginalisation properties of the normal distribution for obtaining $q(\mathbf{u} | \mathbf{Z})$. That is, letting $\mathcal{I}_{\mathbf{Z}}$ be the index set associated with a particular $\mathbf{Z} \in \mathcal{Z}$, we have $q(\mathbf{u} | \mathbf{Z}) = \mathcal{N}(\mathbf{u} | \mathbf{m}_{\mathcal{I}_{\mathbf{Z}}}^*, \mathbf{S}_{\mathcal{I}_{\mathbf{Z}}}^*)$, with mean and covariance being the subvector and submatrix indexed by $\mathcal{I}_{\mathbf{Z}}$.

Including the prior and variational distribution over \mathbf{Z} yields the updated ELBO:

$$\begin{aligned} \log p(\mathbf{y} | \mathbf{X}) &\geq \mathbb{E}_{\mathcal{Q}} \left[\log \frac{p(\mathbf{y} | \mathbf{f}_{\neq \mathbf{u}})p(\mathbf{u} | \mathbf{Z})p(\mathbf{Z})}{q(\mathbf{u} | \mathbf{Z})q(\mathbf{Z})} \right] \\ &= \mathbb{E}_{q(\mathbf{Z})} [\mathcal{L}(\mathbf{Z})] - \text{KL}[q(\mathbf{Z}) || p(\mathbf{Z})] \quad (2) \\ &\triangleq \tilde{\mathcal{L}}, \end{aligned}$$

where $\mathcal{L}(\mathbf{Z})$ is the original ELBO from (1) evaluated for a given subset $\mathbf{Z} \in \mathcal{Z}$. This new objective encodes the trade-off between capacity and complexity since the first term increases and the second term decreases in the number of inducing points drawn from $q(\mathbf{Z})$.

3.1 OPTIMISING THE UPDATED ELBO

The first term in (2) is a discrete sum over a large (potentially infinite) domain of subsets and is therefore not straightforward to optimise. One strategy is to apply the reparameterisation trick (Maddison et al. [2017]) which allows for gradients to be propagated through samples from a continuous relaxation of the discrete distribution. For the current work we did derive an approach that utilises this method; however, as it restricts the form of $q(\mathbf{Z})$ we defer it to Appendix B.

Instead we rely on the more general framework of score function estimation (Williams [1992], Fu [2006]), which enables us to obtain unbiased, noisy gradients through Monte

Carlo sampling. Letting λ be the variational parameters of $q_{\lambda}(\mathbf{Z})$ and noting that $\nabla_{\lambda} q_{\lambda}(\mathbf{Z}) = q_{\lambda}(\mathbf{Z}) \nabla_{\lambda} \log q_{\lambda}(\mathbf{Z})$, we have:

$$\begin{aligned} \nabla_{\lambda} \mathbb{E} [\mathcal{L}(\mathbf{Z})] &= \sum_{\mathbf{Z} \in \mathcal{Z}} \mathcal{L}(\mathbf{Z}) q_{\lambda}(\mathbf{Z}) \nabla_{\lambda} \log q_{\lambda}(\mathbf{Z}) \\ &= \mathbb{E}_{q_{\lambda}(\mathbf{Z})} [\mathcal{L}(\mathbf{Z}) \nabla_{\lambda} \log q_{\lambda}(\mathbf{Z})], \end{aligned}$$

which can be approximated by

$$\begin{aligned} \nabla_{\lambda} \mathbb{E} [\mathcal{L}(\mathbf{Z})] &\approx \frac{1}{S} \sum_{s=1}^S \mathcal{L}(\tilde{\mathbf{Z}}^{(s)}) \nabla_{\lambda} \log q_{\lambda}(\tilde{\mathbf{Z}}^{(s)}), \quad (3) \\ \tilde{\mathbf{Z}}^{(s)} &\sim q_{\lambda}(\mathbf{Z}). \end{aligned}$$

In the basic form this method may suffer from high variance in the gradients, but various strategies have been developed to alleviate this problem (Glasserman [2013]), many of which include subtracting a baseline from $\mathcal{L}(\tilde{\mathbf{Z}}^{(s)})$. We found that using a decaying average of the samples from $\mathcal{L}(\tilde{\mathbf{Z}}^{(s)})$ as baseline proved sufficient for obtaining stable optimisation. Note that since $\mathcal{L}(\mathbf{Z})$ is computed as in (1), this method can be straightforwardly implemented in any existing GP framework with automatic differentiation. The usual parameters of $\mathcal{L}(\mathbf{Z})$ can furthermore be included in the gradient approximation of (3) allowing for jointly learning the SVGP and its supporting inducing set.

Score function estimation increases the time complexity with a factor of S compared to standard SVGP optimisation. However, when smaller sets of inducing points are sampled from $q_{\lambda}(\mathbf{Z})$, the evaluation of $\mathcal{L}(\tilde{\mathbf{Z}})$ becomes cheaper so the effective complexity depends on the degree to which the prior, $p(\mathbf{Z})$, promotes sparsity.

In the following section we present point processes that make the KL divergence analytically tractable. This crucially implies that in order to maximise (2) we only need to be able to sample from $q_{\lambda}(\mathbf{Z})$ and evaluate the gradients of the associated log probability mass.

3.2 SPECIFYING THE POINT PROCESSES

The main motivation for this work is to allow the model to infer which inducing points to include based on their informativeness w.r.t. the observed data. We therefore find it natural to encourage fewer points by having the prior assign probability according to the squared cardinality of \mathbf{Z} :

$$p_{\alpha}(\mathbf{Z}) = C \cdot e^{-\alpha |\mathbf{Z}|^2}.$$

Here we have introduced α as a (fixed) hyper-parameter that indicates the strength of the prior while C is a normalising constant. We use the squared cardinality because, for the models of our focus, the time complexity is $\mathcal{O}(M^2)$ when making new predictions. However, this is ultimately a modelling choice. Note that because we never sample from the

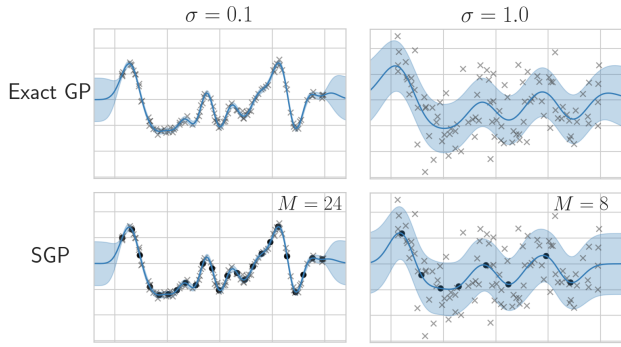


Figure 2: Interaction between observation noise, σ , and the informativeness of inducing points. As the noise increases, more of the true function is lost in the noise floor, and the exact posterior over functions (top row) becomes more uncertain. As a result, fewer inducing points are required to provide an accurate approximation of the posterior (bottom row).

prior when evaluating the objective in (2), we only have to define the probability mass function which affords us a lot of freedom in designing $p(\mathbf{Z})$. Another idea that was recently explored in Rossi et al. [2021] is to use the prior to encourage dispersion amongst the inducing inputs. This could potentially be combined with our approach in future work.

We let $q_\lambda(\mathbf{Z})$ be a discrete Poisson point process (PPP) that is constrained to a pre-defined set of candidate points, \mathbf{Z}^* (Streit [2010]). This process associates each point in \mathbf{Z}^* with an independent probability of inclusion:

$$q_\lambda(\mathbf{Z}) = \prod_{\mathbf{z}_k \in \mathbf{Z}} \lambda_k \prod_{\mathbf{z}_k \notin \mathbf{Z}} (1 - \lambda_k),$$

with $\lambda = \{\lambda_k\}_{k=1}^K$ and K being the total number of candidate points, as illustrated in Figure 1b. Under these definitions of $p_\alpha(\mathbf{Z})$ and $q_\lambda(\mathbf{Z})$, the KL divergence becomes closed-form computable. To see this, first note that $|\mathbf{Z}|$ under $q_\lambda(\mathbf{Z})$ follows a Poisson binomial distribution, and so

$$\mathcal{E} \triangleq \mathbb{E}[|\mathbf{Z}|] = \sum_{k=1}^K \lambda_k, \quad \mathcal{V} \triangleq \text{Var}[|\mathbf{Z}|] = \sum_{k=1}^K \lambda_k(1 - \lambda_k).$$

Decomposing into entropy and cross-entropy terms, we have:

$$\begin{aligned} \text{KL}[q_\lambda(\mathbf{Z}) \parallel p_\alpha(\mathbf{Z})] &= CE(\lambda, \alpha) - H(\lambda), \\ H(\lambda) &= - \sum_{k=1}^K (\lambda_k \log \lambda_k + (1 - \lambda_k) \log (1 - \lambda_k)), \\ CE(\lambda, \alpha) &= - \log C + \alpha \mathbb{E}_{q_\lambda(\mathbf{Z})}[|\mathbf{Z}|^2] \\ &= - \log C + \alpha(\mathcal{V} + \mathcal{E}^2). \end{aligned} \quad (4)$$

The normalisation constant is $C = 1 / \sum_{k=1}^K \binom{K}{k} e^{-\alpha \cdot k^2}$ but can be ignored as it contains no free parameters.

The new hyperparameter, α , is user-defined and reflects the sparsity level that we expect to be sufficient for obtaining a suitable approximation. In contrast to specifying the exact number of points as in standard methods, this means that the model can identify an informative subset out of the entire set of candidates, and additionally adapt the number of points according to the amount and characteristics of the observed data.

Remark: An interesting generalisation of our framework is to let $q_\lambda(\mathbf{Z})$ be a point process that is i) continuous, i.e. not restricted to a finite set of candidate points, and/or ii) capable of expressing correlation between individual points. Both of these qualities are afforded by determinantal point processes (DPP) (Kulesza and Taskar [2012]). However, this would render the entropy term in (4) intractable and require either that we drop it, thus loosening the bound of $\tilde{\mathcal{L}}$, or absorb it into the stochastic approximation of (3). To maintain focus on the paradigm of deriving posteriors over inducing point selection, we restrict this work to the simpler yet effective and robust PPP setting and leave the exploration of more flexible point processes as future work.

4 RELATED WORK

Before the introduction of inducing points, sparse GP approximations were obtained by conditioning on only a subset of the observed data (Williams and Seeger [2001], Csató and Opper [2002], Lawrence et al. [2002]). Here we find commonalities with our work, since the task is to identify the most informative subset of (observed) datapoints. Generally, these methods employ discrete optimisation and often rely on greedy heuristics. With the introduction of inducing points the focus naturally moved to M as the new bottleneck. Various strategies have been proposed for reducing complexity, e.g. decoupling the variational mean and covariance to enable distinct sets of inducing points (Salimbeni et al. [2018], Havasi et al. [2018]) or using the spectral representation of kernels to express the posterior through Fourier features (Lázaro-Gredilla et al. [2010], Hensman et al. [2018]). In work parallel with ours, Rossi et al. [2021] also expand the hierarchical model to include the inducing inputs and even consider the DPP as a prior. Their motivation differs fundamentally from ours, however, in that they keep the number of points fixed and use the prior to improve model fitting and subsequent inference. Identifying a good selection of inducing points has also been examined in theoretical work by Burt et al. [2019], who provide an asymptotic bound on the KL divergence as a function of the number of data- and inducing points, input distribution, and kernel smoothness. The latter results lend credence to common heuristics regarding how many number of inducing points to utilise for a given modelling task, and we will use them as a point of reference in our experiment section. However, they

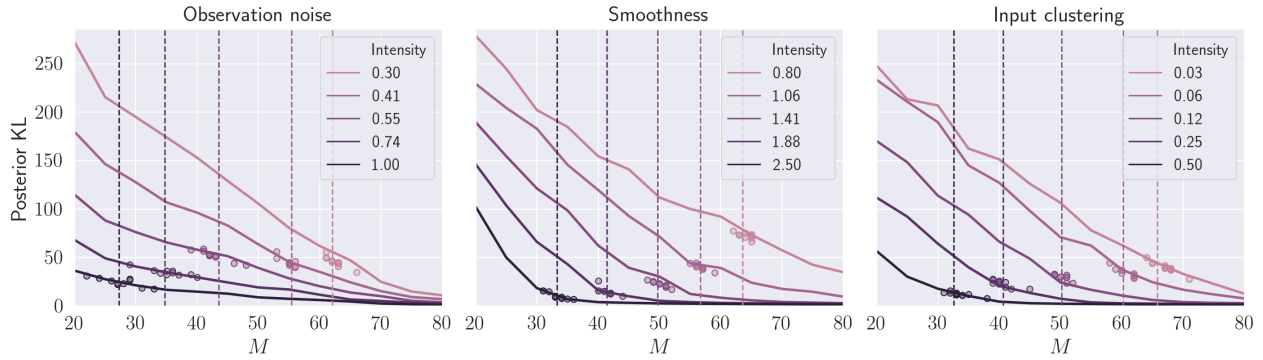


Figure 3: KL divergence between the approximate and exact posterior as a function of number of inducing points under different data characteristics (lower is better). Solid lines are the baseline SVGP with a fixed number of inducing points. As a given intensity increases, fewer points are required to obtain the same KL divergence. The vertical, dashed lines show the expected number of inducing points inferred by the point process under different characteristics, and the circles are 10 samples of subsets drawn from each. Our method adaptively reduces the number of points as informativeness decreases, and in certain cases it slightly improves upon the baseline.

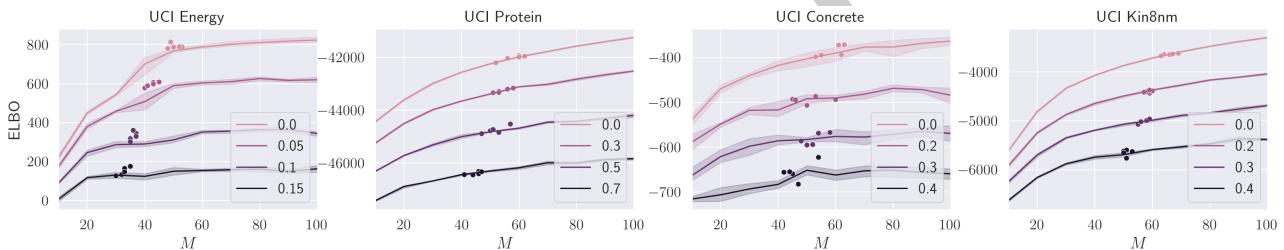


Figure 4: ELBO as a function of number of inducing points, when adding Gaussian noise of varying intensity (higher is better). The setup is similar to that of Figure 3.

are less helpful when the influential factors are unknown prior to observing the data.

The above body of work offer various strategies for exploiting available resources when doing sparse approximations. However, a common thread is that the number of inducing points can generally be categorised as a design choice. Our contribution lies in associating this choice with uncertainty and thus reframing the problem of balancing capacity and complexity as being part of the Bayesian inference.

5 EXPERIMENTS

In this section we provide empirical evidence for the efficacy of our approach in various contexts of SVGP modelling. We first show for synthetic and real-world datasets that as the informativeness of inducing points decreases due to data characteristics, the model prunes away more points. Next, we turn to the practical problem of dynamically allocating inducing points amongst the layers of a deep Gaussian process. Lastly, we demonstrate how the method can be applied in latent variable modelling (GP-LVM) to jointly learn a low-dimensional representation along with a supporting set of inducing points.

5.1 INFORMATIVENESS OF INDUCING POINTS

Previous work has identified various characteristics that determine how effective the inducing points are in communicating information about the observed data. In the works of Burt et al. [2019] and Hensman and Lawrence [2014] it is shown that fewer points are required to achieve the same quality of approximation when i. the observation noise increases, ii. the kernel becomes more smooth, or iii. the input data is more clustered. This effect is shown for the case of observation noise in Figure 2 and illustrates how data characteristics, that are often not known a priori, influence the capacity/complexity trade-off.

In this experiment we vary the intensity of these three characteristics – observation noise, smoothness, and input clustering – and demonstrate that our method does indeed adapt to those conditions by selecting fewer inducing points when their informativeness decreases. As a baseline for our method we fitted SVGP’s with increasing number of inducing points to noisy observations, for which we had adjusted the intensity of each characteristic separately. See Figure 3 where the solid lines are the posterior KL divergence as a function of number of inducing points for 5 different intensity levels. For higher intensities, the divergence reaches a

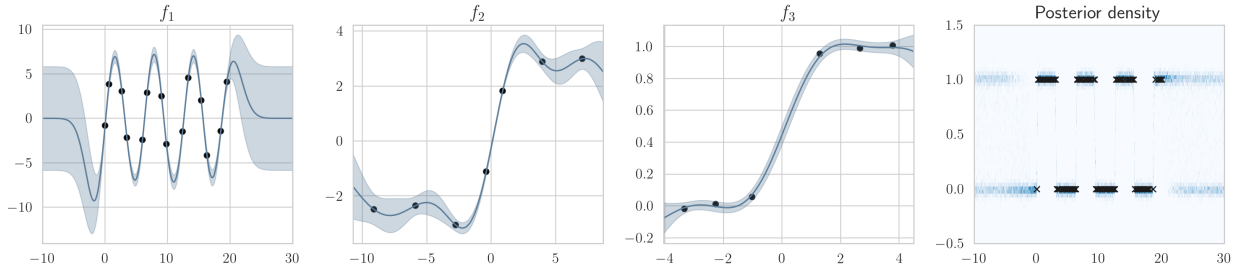


Figure 5: Illustration of different function fidelities in a 3 layer DGP. The three left plots show the learnt SVGP’s for each layer with the inducing being the black circles. The target observations, drawn from a square wave, are plotted to the right on top of the DGP’s posterior density.

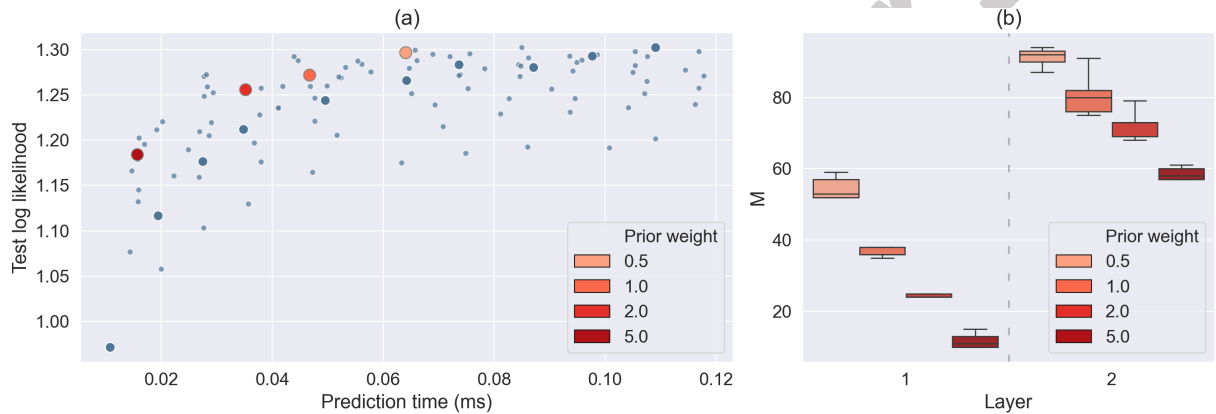


Figure 6: Dynamically allocating inducing points in a 2 layer DGP. (a) Each dot corresponds to the accuracy and performance of a specific configuration of inducing points. Grey dots are the outcomes of a comprehensive and tedious grid-search, with the 10 large dots representing those configurations that have equal number of inducing points both layers. Red dots are the configurations found by our point process for four different settings of the prior weight, α . Note that the latter all fall on the frontier of optimal trade-offs between accuracy and performance. (b) The distribution of points across layers over 5 folds. We see a clear preference towards placing more inducing points in the last layer.

plateau more quickly, showing a saturation of the information communicated by the inducing points.

We then applied our method for each intensity level, using an initial set of 80 inducing points. The vertical lines in Figure 3 show the expected number of points after fitting the point process for a given intensity; for each we have drawn 10 samples from the point process. The results demonstrate that the number of inferred points decrease monotonically as the intensities increase, showing that the point process does indeed adapt to the characteristics of the data; in fact, for some intensities, the adaptive approach achieves a slightly higher reduction of KL divergence.

The same experiment was repeated for 4 real-world datasets by adding varying degrees of Gaussian noise to the training outputs. Results are displayed in Figure 3 where we use the ELBO after 5000 epochs as metric. Again we see that as more noise is added, fewer points are deemed necessary by the points process for emulating the exact GP. The baseline for each noise level shows the trade-off between capacity and complexity that would normally have to be identified

through expensive search over M , but that our method finds adaptively as part of the inference. We give a detailed description of the experimental setups in Appendix C.

5.2 DEEP GAUSSIAN PROCESS

When specifying a deep Gaussian process (DGP) in the sparse framework, each layer is associated with its own set of inducing points. Without evidence to the contrary, the most natural choice might be to use the same number of points for each layer. However, we put forth the conjecture that the layers could potentially model varying levels of fidelity and thus require different numbers of inducing points in order to reach an optimal trade-off between accuracy and complexity. As a motivating example, consider Figure 5 where we apply a three-layer DGP to approximate a square-wave. Such discontinuous data are notoriously difficult for shallow GP’s using a smooth, stationary kernel. The DGP finds a good fit by making a smooth approximation in the first layer and accentuating the jumps by step functions in the subsequent layers. Here we have applied our method to

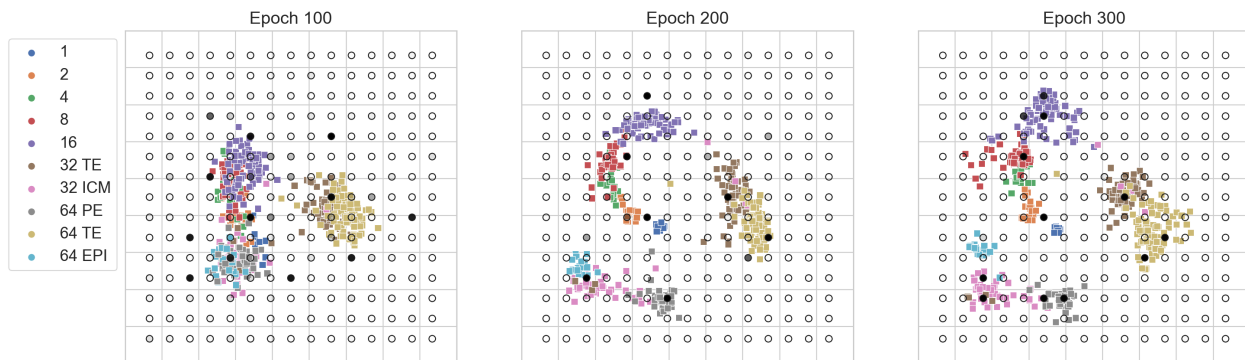


Figure 7: Demonstration of our method applied to a GP-LVM on a single-cell qPCR dataset where observations are coloured according to cell stage. The inducing points are fixed on a 15×15 grid with filling indicating the marginal probability of inclusion (filled means probability of one; open means probability of zero). As the latent inputs change throughout optimisation, different inducing points are activated.

jointly learn a set of inducing points across all layers. From the figure it can be seen that the approximated functions become less complex as we move through the layers, which in turn causes fewer points to be retained by the point process.

We demonstrate that the same approach can be applied for a benchmark dataset, UCI Kin8nm (Dua and Graff [2017]), to incorporate the search for an optimum between capacity and complexity directly in the training of a two-layer DGP. The model follows the doubly-stochastic formulation from Salimbeni and Deisenroth [2017]. In accordance with the findings in Duvenaud et al. [2014], the observed input was added to the second layer which was empirically verified to increase stability. As baseline we first carried out a grid-search over $[10, 20, \dots, 90, 100]$ inducing points in each layer, yielding 100 different configurations. For each configuration we recorded the test log likelihood and prediction time over 5 folds. These results are plotted as grey dots in Figure 6a, and show a significant variance in prediction time for models with similar accuracy. In our adaptive strategy we first pre-fitted the model with 150 inducing points in each layer without the point processes, estimated a posterior set of inducing points, and then re-fitted the pruned model. See Appendix C for a detailed description. The results are plotted in Figure 6a for four different settings of the prior, each yielding different near-optimal trade-offs between accuracy and complexity. Figure 6b shows the hierarchical distribution of inducing points, illustrating a clear preference towards allocating points in the second layer. Finding such optima in the standard setting would effectively be a combinatorial problem where one would have to re-fit the model for each new configuration.

5.3 LATENT VARIABLE MODELLING

Lastly, we apply our method in a Bayesian Gaussian Process latent variable model (GP-LVM) where the latent inputs are learnt along with the function. We use two latent dimensions to infer a representation of 48 single cell gene expressions

for 437 samples.³ The aim is to avoid the need for the extensive cross-validation required to identify a reasonable set of inducing points as the case in for example Ahmed et al. [2018]. The latent variables were initialised with PCA and the locations where fixed to a 15×15 grid for illustrative purposes. All parameters (except for inducing inputs) were trained jointly along with the point process for 300 epochs. Figure 7 illustrates how nearby inducing points are activated as the latent coordinates move around in the input space. In effect our method allows for the model to adapt to the change in function complexity throughout model training and negates the need for finding an appropriate number of inducing points through cross-validation.

6 CONCLUSION

We have presented a Bayesian approach for choosing the number of inducing points in sparse Gaussian processes, thus introducing a probabilistic paradigm for balancing capacity and complexity. By extending the standard probabilistic model with a point process prior we encourage discriminative selection of inducing points, effectively allowing the model to choose only those points that it finds most informative. We apply variational inference to learn an approximate posterior jointly along with the usual model parameters, and present choices of prior and variational point processes that make the resulting evidence lower bound straightforward to optimise. The efficacy of this approach is verified in a controlled experiment on synthetic and real-world data, and its practical applicability is demonstrated for deep Gaussian process regression and latent variable modelling.

Acknowledgements: BSJ and VC acknowledge support from EPSRC grant EP/R018634/1: Closed-Loop Data Science for Complex, Computationally- and Data-Intensive Analytics.

³qPCR dataset: <https://github.com/sods/ods>

REFERENCES

- S. Ahmed, M. Rattray, and A. Boukouvalas. Grandprix: Scaling up the Bayesian GPLVM for single-cell data. *Bioinformatics*, 35(1):47–54, 07 2018.
- K. Blomqvist, S. Kaski, and M. Heinonen. Deep convolutional Gaussian processes. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 582–597. Springer, 2019.
- T. D. Bui, C. Nguyen, and R. E. Turner. Streaming sparse Gaussian process approximations. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3299–3307. 2017.
- D. Burt, C. E. Rasmussen, and M. Van Der Wilk. Rates of convergence for sparse variational Gaussian process regression. In *Proceedings of the International Conference on Machine Learning*, pages 862–871, 2019.
- L. Csató and M. Opper. Sparse online Gaussian processes. *Neural Computation*, 14(3):641–668, 2002.
- A. Damianou and N. Lawrence. Deep Gaussian Processes. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 207–215, 2013.
- M. Deisenroth and C. E. Rasmussen. PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning*, pages 465–472, 2011.
- D. Dua and C. Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- D. Duvenaud, O. Rippel, R. Adams, and Z. Ghahramani. Avoiding pathologies in very deep networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 202–210, 2014.
- M. C. Fu. Gradient estimation. *Handbooks in operations research and management science*, 13:575–616, 2006.
- P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.
- O. Hamelijnck, T. Damoulas, K. Wang, and M. Girolami. Multi-resolution multi-task Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 14025–14035. 2019.
- M. Havasi, J. M. Hernández-Lobato, and J. J. Murillo-Fuentes. Deep Gaussian processes with decoupled inducing inputs. *arXiv:1801.02939*, 2018.
- J. Hensman and N. D. Lawrence. Nested variational compression in deep Gaussian processes. *arXiv preprint arXiv:1412.1370*, 2014.
- J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 282–290, 2013.
- J. Hensman, A. Matthews, and Z. Ghahramani. Scalable variational Gaussian process classification. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, 2015a.
- J. Hensman, A. G. Matthews, M. Filippone, and Z. Ghahramani. MCMC for variationally sparse Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 1648–1656, 2015b.
- J. Hensman, N. Durrande, and A. Solin. Variational Fourier features for Gaussian processes. *Journal of Machine Learning Research*, 18, 2018.
- A. Kulesza and B. Taskar. *Determinantal Point Processes for Machine Learning*. Now Publishers Inc., Hanover, MA, USA, 2012. ISBN 1601986289.
- N. Lawrence, M. Seeger, and R. Herbrich. Fast sparse Gaussian process methods: The informative vector machine. In *Advances in Neural Information Processing Systems*, 2002.
- N. D. Lawrence. Gaussian process latent variable models for visualisation of high dimensional data. In *Advances in Neural Information Processing Systems*, pages 329–336, 2004.
- M. Lázaro-Gredilla and A. Figueiras-Vidal. Inter-domain Gaussian processes for sparse inference using inducing features. In *Advances in Neural Information Processing Systems*, pages 1087–1095, 2009.
- M. Lázaro-Gredilla, J. Quiñonero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal. Sparse spectrum Gaussian process regression. *The Journal of Machine Learning Research*, 11:1865–1881, 2010.
- M. Lorenzi, G. Ziegler, D. C. Alexander, and S. Ourselin. Efficient Gaussian process-based modelling and prediction of image time series. In *Proceedings of the International Conference on Information Processing in Medical Imaging*, pages 626–637. Springer, 2015.
- C. J. Maddison, A. Mnih, and Y. W. Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017.

- A. G. d. G. Matthews, J. Hensman, R. Turner, and Z. Ghahramani. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. *Journal of Machine Learning Research*, 51:231–239, 2016.
- J. Moćkus. On Bayesian methods for seeking the extremum. In *IFIP technical conference*, pages 400–404. Springer, 1975.
- J. Quiñonero-Candela and C. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, USA, Jan. 2006.
- S. Rossi, M. Heinonen, E. Bonilla, Z. Shen, and M. Filippone. Sparse Gaussian processes revisited: Bayesian approaches to inducing-variable approximations. In *International Conference on Artificial Intelligence and Statistics*, pages 1837–1845. PMLR, 2021.
- H. Salimbeni and M. P. Deisenroth. Doubly stochastic variational inference for deep Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 4589–4600, 2017.
- H. Salimbeni, C.-A. Cheng, B. Boots, and M. Deisenroth. Orthogonally decoupled variational Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 8711–8720, 2018.
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pages 1257–1264, 2006.
- R. L. Streit. *Poisson point processes: Imaging, tracking, and sensing*. Springer Science & Business Media, 2010.
- M. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, pages 567–574, 2009.
- M. Titsias and N. D. Lawrence. Bayesian Gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 844–851, 2010.
- C. K. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems*, pages 682–688, 2001.
- R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992.