
Doubly Sparse Variational Gaussian Processes

Vincent Adam Stefanos Eleftheriadis Nicolas Durrande Artem Artemev James Hensman
PROWLER.io, Cambridge, UK

Abstract

The use of Gaussian process models is typically limited to datasets with a few tens of thousands of observations due to their complexity and memory footprint. The two most commonly used methods to overcome this limitation are 1) the variational sparse approximation which relies on inducing points and 2) the state-space equivalent formulation of Gaussian processes which can be seen as exploiting some sparsity in the precision matrix. We propose to take the best of both worlds: we show that the inducing point framework is still valid for state space models and that it can bring further computational and memory savings. Furthermore, we provide the natural gradient formulation for the proposed variational parameterisation. Finally, this work makes it possible to use the state-space formulation inside deep Gaussian process models as illustrated in one of the experiments.

1 Introduction

Gaussian processes (GPs) provide a very powerful framework for statistical modelling in low data regimes (i.e., when the number of observations N is small). However, they typically scale as $\mathcal{O}(N^3)$ in computational complexity and $\mathcal{O}(N^2)$ in memory which makes them impractical for datasets containing more than a few thousand observations. This limitation has received a lot of attention, especially in the machine learning community (Rasmussen and Williams, 2006), and two frameworks have distinguished themselves so far.

The first approach focuses on state-space models (SSM) and exploits the underlying Markov property for computational efficiency. The Markov structure results in




Figure 1: Illustration of the proposed method for a regression task with a Matérn $5/2$ kernel. For each inducing location z_i , we introduce three inducing variables, i.e., $f(z_i)$, $f'(z_i)$ and $f''(z_i)$. Figure also shows samples from the model, where the inducing variables are represented by a quadratic Taylor expansion. For more details on the experiment settings see Section 4.1.

sparse precision matrices (Grigorievskiy et al., 2017; Durrande et al., 2019), or enables filtering algorithms (Kalman, 1960; Särkkä and Solin, 2019; Solin et al., 2018). These approaches lead to inference algorithms with $\mathcal{O}(N)$ complexity. Several model types can be tackled with this approach, such as Gaussian Markov random fields or linear stochastic differential equations. The second approach is the sparse variational Gaussian process (SVGP). It relies on the assumption that the dataset $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$ contains some redundant information, so that the GP posterior $f|\{f(x_i) = y_i\}_{i=1}^N$ can be approximated using a smaller set of inducing points (i.e. pseudo-inputs): $f|\{f(z_i) = u_i\}_{i=1}^M$ with $M \ll N$. Variational Inference (VI), which consists in minimising the Kullback–Leibler divergence between the approximate and the true posterior, is then used to find appropriate values for \mathbf{z} , for and for the model parameters (Titsias, 2009; Hensman et al., 2013).

In this paper, we focus on GPs with 1-dimensional input and propose a novel inference method that brings together the merits of both sparse GP approximation and SSM representation: for a state-space model with state dimension d , we introduce d -dimensional inducing variables and associate each element of this vector to a state component. For example, a GP f with a Matérn $5/2$ covariance is Markovian if you consider the state space

$(f(x), f'(x), f''(x))$. Given some inducing locations z_i , we approximate the posterior $f|\{f(x_i) = y_i\}_{i=1}^N$, by $f|\{f(z_j) = u_{j,1}, f'(z_j) = u_{j,2}, f''(z_j) = u_{j,3}\}_{j=1}^M$ (see Figure 1). Using the state space components as inducing features brings several advantages. First, the proposed method scales linearly both with the number of data and with the number of inducing points (whereas SVGP is quadratic in the later). Second, the number of variational parameters that need optimising is $\mathcal{O}(Md^2)$, which scales favourably compared to $\mathcal{O}(M^2)$ and $\mathcal{O}(Nd^2)$ as respectively required by SVGP and VI for SSM (Durrande et al., 2019). Finally, our approach allows for mini-batching as well as the use of SSM layers in deep-GP models.

In developing our approach, we quickly identified that optimisation of the objective function was cumbersome using standard gradient approaches. We developed a natural gradient approach for our method based on Salimbeni et al. (2018). To do this efficiently requires the formulation of the compact exponential family form of a Markov structured Gaussian distribution, with novel mathematical and computational operators.

2 Background

Here we introduce the basics of GP models (Rasmussen and Williams, 2006) as well as the two main techniques for dealing with large datasets: sparse variational inference and inference in the state-space formulation. They both lead to sparse algorithms, albeit in a different way.

2.1 Gaussian processes

In the classic GP regression setting we are given a dataset $\mathcal{D} = \{\mathbf{x}_n, y_n\}_{n=1\dots N} \in (\mathcal{X}, \mathbb{R})^N$, where y_n corresponds to the evaluation of a latent function corrupted with observation noise $y_n = f(\mathbf{x}_n) + \epsilon_n$ and the task is to predict $f(\mathbf{x}_*)$ for some input location $\mathbf{x}_* \in \mathcal{X}$. Under the assumptions that f is a Gaussian process $f(\cdot) \sim \mathcal{GP}(0, k(\cdot, \cdot))$ and that ϵ also follows a multivariate normal distribution $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, the prediction problem can be solved analytically and results in a Gaussian distribution:

$$\mathbb{E}_{f|\mathcal{D}}[f(\mathbf{x}_*)] = \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \quad (1)$$

$$\mathbb{V}_{f|\mathcal{D}}[f(\mathbf{x}_*)] = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (2)$$

where the matrix \mathbf{K} is defined as $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ and \mathbf{k}_* is vector such that $\mathbf{k}_{*i} = k(\mathbf{x}_*, \mathbf{x}_i)$. The matrix inverse required in the prediction yields a computational complexity of $\mathcal{O}(N^3)$ which is prohibitively expensive for large datasets.

2.2 Variational inter-domain approximations

Variational inference in probabilistic models turns inference into an optimisation problem (Jordan et al., 1999). The sparse variational Gaussian process, originally introduced by Titsias (2009) and more rigorously defined by Matthews et al. (2016), approximates the posterior distribution $p(f(\cdot)|\mathcal{D})$ by a distribution $q(f)$ that depends on ‘inducing points’:

$$q(f(\cdot)) = \int p(f(\cdot)|\{f(\mathbf{z}_i) = \mathbf{u}_i\}_{i=1}^M) q_{\mathbf{u}}(\mathbf{u}) d\mathbf{u}. \quad (3)$$

The inducing variables \mathbf{u} are typically assumed to be normally distributed (i.e., $q_{\mathbf{u}} = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}}, \boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}})$) with the moments corresponding to the variational parameters.

The above approach has been generalised by introducing the idea of inter-domain features (Alvarez and Lawrence, 2009; Lázaro-Gredilla and Figueiras-Vidal, 2009) which replaces the conditioning $\{f(\mathbf{z}_i) = \mathbf{u}_i\}$ in Eq. (3) by $\{\Psi_i[f(\cdot)] = \mathbf{u}_i\}$ with Ψ a linear operator $\mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}$. Choosing $\Psi_i : f \rightarrow f(\mathbf{z}_i)$ allows to recover the classic inducing points but more interesting behaviour can be obtained by using other operators such as the integrals or convolutions (Hensman et al., 2018; van der Wilk et al., 2017). To enhance readability, we denote by $\Psi[f]$ the vector of size M with entries $\Psi_i[f]$, and by p_{Ψ} its distribution.

The distribution of the conditioned GP is

$$f(\cdot)|\{\Psi[f] = \mathbf{u}\} \sim \mathcal{N}(\mathbf{k}_{\Psi}^\top(\cdot) \mathbf{K}_{\Psi\Psi}^{-1} \mathbf{u}, k(\cdot, \cdot) - \mathbf{k}_{\Psi}^\top(\cdot) \mathbf{K}_{\Psi\Psi}^{-1} \mathbf{k}_{\Psi}(\cdot)), \quad (4)$$

with $(\mathbf{k}_{\Psi}(x))_i = \text{cov}(f(x), \Psi_i[f])$ and $(\mathbf{K}_{\Psi\Psi})_{i,j} = \text{cov}(\Psi_i[f], \Psi_j[f])$. Taking the expectation of Eq. (4) under $q_{\mathbf{u}}(\mathbf{u})$ gives a closed form expression for $q(\cdot)$:

$$q(\cdot) = \mathcal{GP}(\mathbf{k}_{\Psi}^\top(\cdot) \mathbf{K}_{\Psi\Psi}^{-1} \boldsymbol{\mu}_{\mathbf{u}}, k_{ff}(\cdot, \cdot) - \mathbf{k}_{\Psi}^\top(\cdot) \mathbf{K}_{\Psi\Psi}^{-1} (\mathbf{K}_{\Psi\Psi} - \boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}) \mathbf{K}_{\Psi\Psi}^{-1} \mathbf{k}_{\Psi}(\cdot)). \quad (5)$$

The variational lower bound to the log-marginal likelihood is then given by:

$$\mathcal{L}(q) = \mathbb{E}_q[\log p(\mathbf{y}|f(\mathbf{X}))] - \text{KL}[q_{\mathbf{u}} \| p_{\Psi}]. \quad (6)$$

The evaluation of \mathcal{L} can be shown to scale as $\mathcal{O}(NM^2 + M^3)$. The linear scaling with the size of the training set allows to apply this approximation to large datasets (Hensman et al., 2013). In some cases, it is shown to approximate the posterior process with high accuracy at a low computational cost in the large data regime (Burt et al., 2019). Finally, it provides a well-defined objective, amenable to gradient-based optimisation that works well in practice (Bauer et al., 2016).

2.3 Inference in the state-space formulation

A large class of Gaussian processes defined on $\mathcal{X} \subseteq \mathbb{R}$ can be written as linear stochastic differential equations

(SDEs). This class of kernels is described in depth in Chapter 12 of Särkkä and Solin (2019) and includes, for example, all Matérn $k/2$ (k odd), harmonic oscillators, and all sum and products of such kernels. Also, many kernels can be approximated by an element of this class, see for example Särkkä and Piché (2014) for the RBF kernel.

Here we focus on the class of GPs with stationary kernels which can be represented as a linear time-invariant (LTI) SDE of the form:

$$\dot{\mathbf{s}}(t) = \mathbf{F}\mathbf{s}(t) + \mathbf{L}\boldsymbol{\varepsilon}(t), \quad f(t) = \mathbf{H}\mathbf{s}(t). \quad (7)$$

The state-space vector $\mathbf{s}(t) \in \mathbb{R}^d$ is given by evaluations of the process and its derivatives $\mathbf{s}(t) = [f(t), f^{(1)}(t), \dots, f^{(d-1)}(t)]^\top$. $\boldsymbol{\varepsilon}(t) \in \mathbb{R}^r$ is a white noise process with spectral density \mathbf{Q}_c . $\mathbf{F} \in \mathbb{R}^{d \times d}$, $\mathbf{L} \in \mathbb{R}^{d \times r}$, $\mathbf{H} \in \mathbb{R}^{1 \times d}$ are the feedback, noise effect and observation matrix of the system.

The marginal distribution of the solution of this LTI-SDE evaluated at any ordered set $\mathbf{x} = [x_1, \dots, x_N]^\top \in \mathbb{R}^N$ follows a discrete-time linear system:

$$\begin{aligned} \mathbf{s}(x_{n+1}) &= \mathbf{A}_{n,n+1}\mathbf{s}(x_n) + \mathbf{q}_n, \quad \mathbf{q}_n \sim \mathcal{N}(0, \mathbf{Q}_{n,n+1}) \\ \mathbf{s}(x_0) &\sim \mathcal{N}(0, \mathbf{P}_0), \quad f(x_n) = \mathbf{H}\mathbf{s}(x_n) \end{aligned} \quad (8)$$

where the state transition matrices $\mathbf{A}_{n,n+1} \in \mathbb{R}^{d \times d}$, noise covariance matrices $\mathbf{Q}_{n,n+1} \in \mathbb{R}^{d \times d}$, and state stationary covariance matrix \mathbf{P}_0 can be computed analytically. Denoting Φ the matrix exponential and $\Delta_n = x_{n+1} - x_n$, we have

$$\begin{aligned} \mathbf{A}_{n,n+1} &= \Phi(\mathbf{F}\Delta_n), \\ \mathbf{Q}_{n,n+1} &= \int_0^{\Delta_n} \Phi(\Delta_n - \tau)\mathbf{L}\mathbf{Q}_c\mathbf{L}^\top\Phi(\Delta_n - \tau)^\top d\tau. \end{aligned} \quad (9)$$

$$(10)$$

By following the standard Kalman recursions (Särkkä, 2013), inference in conjugate models using the state-space formulation of the GP prior scales linearly with the number of data points and cubically with the state-space dimension, i.e., $\mathcal{O}(Nd^3)$. Equally efficient approximations have been derived in the non-conjugate case (Durrande et al., 2019; Nickisch et al., 2018), by exploiting the block-tridiagonal structure of the precision (Grigorievskiy et al., 2017).

3 Doubly sparse inference

In this section we introduce the idea of combining the inter-domain inducing features with the state-space GP formulations, which results in what we dub a “doubly sparse variational GP” approximation (S²VGP). Although the scope of our approach is broader, we formulate our idea under the classic GP regression setting with factorising likelihood: $p(\mathbf{y}, f(\cdot)|\mathbf{x}) = p(f(\cdot)) \prod_n p(y_n|f(x_n))$.

3.1 State-space inducing features

We restrict our analysis to GPs with an LTI-SDE representation and we choose $\Psi_i : f \rightarrow \mathbf{s}(z_i) = [f(z_i), \dots, f^{(d-1)}(z_i)]^\top$ as our inter-domain features evaluated at ordered inducing inputs $\mathbf{z} = [z_1, \dots, z_M]^\top$. This has two immediate desirable consequences:

Property 1: The sequence of inducing states $\mathbf{u} = \{\Psi_i[f]\}_{i=1}^M$ is Markovian and its distribution is multivariate normal $p_\Psi = \mathcal{N}(0, \mathbf{Q}_\Psi^{-1})$. This means that \mathbf{Q}_Ψ has a block-band structure and that the sufficient statistics are $\mathbf{t}(\mathbf{u}) = [\mathbf{u}, \text{btd}[\mathbf{u}\mathbf{u}^\top]]$, where btd extracts the block-tridiagonal elements (with block size d) and returns them as a vector. This property can be summarised as:

$$p_\Psi(\mathbf{u}) = p(\mathbf{u}_1) \prod_m p(\mathbf{u}_{m+1}|\mathbf{u}_m) \propto \exp(\boldsymbol{\theta}^\top \mathbf{t}(\mathbf{u})), \quad (11)$$

where $\boldsymbol{\theta} = [\mathbf{Q}_\Psi \boldsymbol{\mu}_\Psi, -1/2 \text{btd}[\mathbf{Q}_\Psi]]$ are the natural parameters of the distribution.

Property 2: The posterior of the function evaluation at a point x_n conditioned on the inducing variables, depends only on the closest left and right inducing states:

$$\begin{aligned} p(f(x_n)|\mathbf{u}) &= \frac{p(\mathbf{u}_{1:n_-})p(f(x_n)|\mathbf{u}_{n_-})p(\mathbf{u}_{n_+}|f(x_n))p(\mathbf{u}_{n_+ : M}|\mathbf{u}_{n_+})}{p(\mathbf{u}_{1:n_-})p(\mathbf{u}_{n_+}|\mathbf{u}_{n_-})p(\mathbf{u}_{n_+ : M}|\mathbf{u}_{n_+})} \\ &= p(f(x_n)|\mathbf{u}_{n_-}, \mathbf{u}_{n_+}), \end{aligned} \quad (12)$$

where the indices $n_- \in \llbracket 1, M-1 \rrbracket$ and $n_+ = n_- + 1$ correspond to the indices of the closest lower and upper neighbor of x_n in \mathbf{z} , i.e., $z_1 < \dots < z_{n_-} < x_n < z_{n_+} < \dots < z_M$. This follows directly from $f(x_n) = \mathbf{H}\mathbf{s}(x_n)$ and the Markovian property of $\{\mathbf{u}_1, \dots, \mathbf{u}_{n_-}, \mathbf{s}(x_n), \mathbf{u}_{n_+}, \dots, \mathbf{u}_M\}$ which ensures that $p(\mathbf{s}(x_n)|\mathbf{u}) = p(\mathbf{s}(x_n)|\mathbf{u}_{n_-}, \mathbf{u}_{n_+})$. The graphical model associated with these properties is given in Figure 2. In practice, we get the statistics of $p(\mathbf{s}(x_n)|\mathbf{u}_{n_-}, \mathbf{u}_{n_+})$ in time $\mathcal{O}(d^3)$ using the state-space parameters:

$$\begin{aligned} p(\mathbf{s}(x_n)|\mathbf{u}_{n_-}, \mathbf{u}_{n_+}) &\propto \mathcal{N}(\mathbf{s}(x_n); \mathbf{A}_{n_-,n}\mathbf{u}_{n_-}, \mathbf{Q}_{n_-,n}) \\ &\quad \times \mathcal{N}(\mathbf{u}_{n_+}|\mathbf{A}_{n,n_+}\mathbf{s}(x_n), \mathbf{Q}_{n,n_+}) \\ &= \mathcal{N}(\mathbf{s}(x_n); \mathbf{P}_n\mathbf{v}_n, \mathbf{T}_n), \end{aligned} \quad (13)$$

where $\mathbf{v}_n = [\mathbf{u}_{n_-}; \mathbf{u}_{n_+}]$. The matrices $\mathbf{P}_n, \mathbf{T}_n$ depend on the statistics of the prior state transitions between time points triplet (z_{n_-}, x_n, z_{n_+}) and are given in Appendix A.1.

Finally, we define the marginal posterior on \mathbf{v}_n as $q(\mathbf{v}_n) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{v}_n}, \boldsymbol{\Sigma}_{\mathbf{v}_n}\mathbf{v}_n)$, and obtain the posterior predictions analytically as

$$q(\mathbf{s}(x_n)) = \mathcal{N}(\mathbf{s}(x_n); \mathbf{P}_n\boldsymbol{\mu}_{\mathbf{v}_n}, \mathbf{T}_n + \mathbf{P}_n\boldsymbol{\Sigma}_{\mathbf{v}_n}\mathbf{v}_n\mathbf{P}_n^\top). \quad (14)$$

3.2 Optimal variational distribution for inducing state-space features




Figure 2: (a) Graphical model representing the Markovian joint prior of states indexed at $\{x_n, z_m\}$ (black arrows). (b) Graphical model for marginal posterior $q(s_2, \mathbf{u}) = p(s_2|\mathbf{u})q(\mathbf{u})$ highlighting the statistical properties of the variational posterior: (1) $q_{\mathbf{u}}(\mathbf{u})$ is a Markovian sharing the same structure as $p_{\Psi}(\mathbf{u})$ (red arrows); (2) s_2 only depends on the two nearest inducing states, i.e., $p(s_2|\mathbf{u}) = p(s_2|u_1, u_2)$ (blue arrows). The conditional dependencies $p(s_i|\mathbf{u})$, are shown in light blue.

We continue our analysis by investigating the form of the approximating distribution on the inducing states $q_{\mathbf{u}} = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}}, \boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}})$. By following the approach of (Opper and Archambeau, 2009), we show that at the optimum the variational distribution $q_{\mathbf{u}}$ has a precision with the same block-tridiagonal structure as the prior precision \mathbf{Q}_{Ψ} . More specifically, we start from the variational loss in Eq. (6) and expand the terms of the KL that depend on the posterior covariance $\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}$:

$$\begin{aligned} \mathcal{L}(q) &= \sum_n \mathbb{E}_q [\log p(y_n|f(x_n))] - \text{KL}[q_{\mathbf{u}} \| p_{\Psi}] \\ &= \sum_n \mathbb{E}_q [\log p(y_n|f(x_n))] - \frac{1}{2} \text{tr}(\mathbf{Q}_{\Psi} \boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}) \\ &\quad + \frac{1}{2} |\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}| + c(\boldsymbol{\mu}_{\mathbf{u}}, p_{\Psi}), \end{aligned} \quad (15)$$

where $c(\boldsymbol{\mu}_{\mathbf{u}}, p_{\Psi})$ contains the terms in KL that do not depend on $\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}$. At the optimal covariance $\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}^*$, the gradient of the loss w.r.t. the variational parameters is zero, i.e. $\nabla_{\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}} \mathcal{L}(q)|_{\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}^*} = 0$, which leads to:

$$\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}^{*-1} = \mathbf{Q}_{\Psi} - 2 \sum_n \nabla_{\boldsymbol{\Sigma}_{\mathbf{u}\mathbf{u}}} \mathbb{E}_q [\log p(y_n|f(x_n))]. \quad (16)$$

The term \mathbf{Q}_{Ψ} contributes only to the block-tridiagonal band of the precision. The posterior prediction for each data point $q(f(x_n))$ only depends on the marginal covariance $\boldsymbol{\Sigma}_{v_n v_n}$ of the two neighboring inducing states from Eq. (14). So, each likelihood term for data y_n contributes to the posterior precision at the location z_{n-}, z_{n+} , on which $q(f(x_n))$ depends, which is also in the band. That is, the optimal $q_{\mathbf{u}}$ has the same sparsity pattern in the precision as p_{Ψ} . So, we choose to parameterise $q_{\mathbf{u}} = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{u}}, \mathbf{Q}_{\mathbf{u}}^{-1})$, by its mean $\boldsymbol{\mu}_{\mathbf{u}}$ and by a lower triangular matrix $\mathbf{L}_{\mathbf{u}}$ with two block-diagonals which can be interpreted as the Cholesky factor of the block-tridiagonal $\mathbf{Q}_{\mathbf{u}}$. This parameterisation results in a $\mathcal{O}(Md^2)$ storage footprint instead

of $\mathcal{O}(M^2d^2)$ had we chosen the more general *mean, covariance* parameterisation.

3.3 Doubly sparse variational GP inference: S²VGP

We denote by S²VGP the algorithm that performs sparse variational inference with inducing state-space features while restricting the variational distribution to be in the class of multivariate normal distributions with block-tridiagonal banded precisions. Graphical models summarising the corresponding prior and approximate posterior assumptions are given in Figure 2.

S²VGP has the following computational advantages: (1) both the KL divergence and the pairwise marginal posterior predictions of contiguous inducing states (a pairwise Kalman-like smoothing of $q_{\mathbf{u}}$) can be evaluated in $\mathcal{O}(Md^3)$; (2) Given these pairwise marginals on inducing states, the marginal posterior predictions of function evaluations at the data points $q(f(x_n))$ can be evaluated in parallel in $\mathcal{O}(Nd^3)$. Overall, the evaluation of the variational loss (and of its gradient) has complexity $\mathcal{O}((N+M)d^3)$ ¹, which compares favourably against alternative variational methods as shown in Table 1. For efficient implementation of the variational loss and the gradients based on banded precision parameterised Gaussian distributions we refer to (Durrande et al., 2019).

S²VGP further inherits two properties of the SVGP approximation. First, marginal posterior predictions $q(f(x_n))$ for each data point can be computed independently which allows to perform stochastic optimisation of an estimator of the loss evaluated on random mini-batches of size N_b of the data, reducing the complexity of an evaluation of the objective to $\mathcal{O}((N_b + M)d^3)$. Second because the inducing inputs \mathbf{z} are decoupled from the data inputs \mathbf{x} , S²VGP can be used as a GP-layer in a deep (compositional) architecture as in (Salimbeni and Deisenroth, 2017). Both properties are not available to alternative state-space approaches to GP models.

3.4 Natural gradient updates

To learn the variational parameters and the parameters of the model we resort to gradient based optimisation. The gradient of the objective w.r.t. a parameter $\boldsymbol{\xi}$ is defined as $\nabla_{\boldsymbol{\xi}} \mathcal{L} = \lim_{\epsilon \rightarrow 0} \arg \min_{\boldsymbol{\delta}} \frac{1}{\epsilon} \mathcal{L}(\boldsymbol{\xi} + \boldsymbol{\delta})$ subject to constraint $\|\boldsymbol{\delta}\| = \epsilon$. Intuitively, it is the direction of steepest descent with respect to the Euclidean norm of $\boldsymbol{\delta}$, that maximally reduces the loss.

¹Even in this full batch form, this is not always slower than the $\mathcal{O}(Nd^3)$ Kalman smoothing hiding a large constant factor.

Table 1: Complexity and capabilities of variational inference algorithms for GP regression.

Algorithm	Complexity	Storage	deep \mathcal{GP}	minibatch
VGP (GPflow)	$\mathcal{O}(N^3)$	$\mathcal{O}(N^2)$	✗	✗
SVIGP (Hensman et al., 2013)	$\mathcal{O}(N_b M^2 + M^3)$	$\mathcal{O}(M^2)$	✓	✓
VGP (banded) (Durrande et al., 2019)	$\mathcal{O}(Nd^3)$	$\mathcal{O}(Nd^2)$	✗	✗
S ² VGP [this work]	$\mathcal{O}((N_b + M)d^3)$	$\mathcal{O}(Md^2)$	✓	✓

However, the Euclidean norm can be deceiving when optimising over distributions: small changes in parameters can induce large changes in the distributions, and changing the parameterisation of the distribution typically leads to different optimisation performance. Natural gradients solve this problems by substituting the constraint on the Euclidean norm by $\text{KL}[q_{\mathbf{u}}(\boldsymbol{\xi}) \parallel q_{\mathbf{u}}(\boldsymbol{\xi} + \boldsymbol{\delta})] = \epsilon$ (Amari, 1998). Such a constraint can be shown to induce a quadratic norm in the parameter space with curvature given by the Fisher information matrix $\mathbf{F}_{\boldsymbol{\xi}}$. The direction of steepest descent w.r.t. this norm is given by $\tilde{\nabla}_{\boldsymbol{\xi}} \mathcal{L} = (\nabla_{\boldsymbol{\xi}} \mathcal{L}) \mathbf{F}_{\boldsymbol{\xi}}^{-1}$.

Conveniently, for distributions in the exponential family, the Fisher matrix takes a rather simple form for any parameterisation $\boldsymbol{\xi}$ (Malagò and Pistone, 2015):

$$\mathbf{F}_{\boldsymbol{\xi}} = \left(\frac{d\boldsymbol{\theta}}{d\boldsymbol{\xi}} \right)^{\top} \frac{d\boldsymbol{\eta}}{d\boldsymbol{\theta}} \frac{d\boldsymbol{\theta}}{d\boldsymbol{\xi}}, \quad (17)$$

where $\boldsymbol{\theta}$ are the natural parameters, $\boldsymbol{\eta}$ the expectation parameters and $\boldsymbol{\xi}$ the parameterisation of our choice. For the variational distribution $q_{\mathbf{u}}$ defined in Section 3.2 these parameters are equivalent to:

$$\boldsymbol{\theta} = [\mathbf{L}_{\mathbf{u}} \mathbf{L}_{\mathbf{u}}^{\top} \boldsymbol{\mu}_{\mathbf{u}}, -1/2 \text{btd}[\mathbf{L}_{\mathbf{u}} \mathbf{L}_{\mathbf{u}}^{\top}]], \quad (18)$$

$$\boldsymbol{\eta} = \mathbb{E}_{q_{\mathbf{u}}}[\mathbf{t}(\mathbf{u})] = [\boldsymbol{\mu}_{\mathbf{u}}, \text{btd}[\boldsymbol{\mu}_{\mathbf{u}} \boldsymbol{\mu}_{\mathbf{u}}^{\top} + \mathbf{L}_{\mathbf{u}}^{-\top} \mathbf{L}_{\mathbf{u}}^{-1}]]. \quad (19)$$

It is clear that the banded structure of the sufficient statistics is reflected in both natural and expectation parameters. This allows us to derive efficient updates as in (Salimbeni et al., 2018) using the banded operators introduced by Durrande et al. (2019):

$$\tilde{\nabla}_{\boldsymbol{\xi}_{[band]}^{\top}} \mathcal{L} = \frac{d\boldsymbol{\xi}_{[band]}}{d\boldsymbol{\theta}_{[band]}} \frac{d\mathcal{L}}{d\boldsymbol{\eta}_{[band]}}. \quad (20)$$

There is one caveat: the covariance term $(\mathbf{L}_{\mathbf{u}}^{-\top} \mathbf{L}_{\mathbf{u}}^{-1})$ in $\boldsymbol{\eta}$ is a full matrix, but for an efficient update we need to transfer between $\boldsymbol{\xi}_{[band]} \rightleftharpoons \boldsymbol{\eta}_{[band]}$ using only the elements in the band. This requires a novel operator (the reverse of $\mathbb{I}[\cdot]$ in (Durrande et al., 2019, Sec. 4.1)), which maps the band of a covariance to the Cholesky factor of its banded precision. Detailed transformations between $\boldsymbol{\theta}_{[band]} \rightleftharpoons \boldsymbol{\xi}_{[band]} \rightleftharpoons \boldsymbol{\eta}_{[band]}$, along with the algorithm for the operator can be found in the supplementary materials.

In practice, the natural gradient update is recommended when the likelihood is Gaussian since it en-

ures convergence in very few steps (typically one step, see (Salimbeni et al., 2018)). We also observed that it performs well in the non-conjugate case, especially in the first iterations of the optimisation where it can very quickly move to areas of interest. These two properties are illustrated on simple examples in Appendix C.1.

4 Experiments

4.1 Solar irradiance

The aim of this first experiment is to visually illustrate the predictive power of the proposed methodology on a simple GP regression example. We consider here the solar irradiance dataset² and choose the model $y(x) = f(x) + \varepsilon$ where f is a centered GP with Matérn $3/2$ covariance and ε is i.i.d. $\mathcal{N}(0, \sigma^2)$. We compare three (approximate) posteriors for this model: (a) the classic SVGP, (b) the proposed S²VGP, and (c) the exact GPR posterior. For (a) and (b), a grid of 60 inducing locations is used, and the kernel parameters, the noise variance σ^2 and the variational parameters are estimated by maximising the ELBO. For (c) the kernel parameters and the noise variance are estimated by maximising the model log-marginal likelihood.

As shown in Figure 3, the proposed method is much more accurate than SVGP on this example, and its predictions are extremely similar to the exact GPR model. On the one hand, this experimental setting may be seen to the advantage of our method since having the same number of inducing inputs means that there are three times more inducing variables in the S²VGP model. On the other, the computational burden is much smaller for S²VGP and this added flexibility actually comes with a reduction of the computational cost. This calls for a more thorough investigation that explores how the ELBO of SVGP and S²VGP compare with respect to the numbers of inducing variables, variational parameters and the execution time. This is what we do in the next section on a larger dataset.

4.2 Conjugate regression on time series

In this section we illustrate the computational and storage savings our method entails against the classical

²<https://github.com/jameshensman/VFF>.

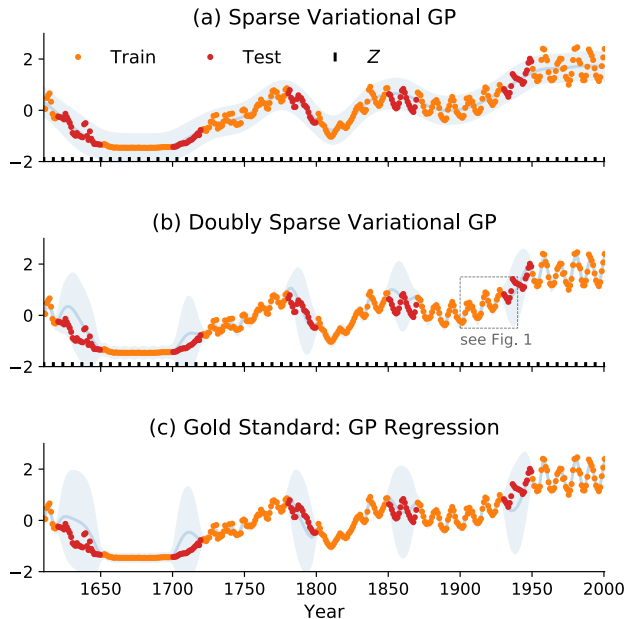


Figure 3: Model comparison on the solar dataset. The SVGP (panel a) fails at capturing the high frequency variations, whereas the proposed S²VGP (panel b) results in predictions that are extremely similar to the ones obtained with exact inference (panel c).

SVGP algorithm on a conjugate regression problem (see Appendix D for a non-conjugate example). The data consists of an uttered vowel from a female speaker (Hillenbrand et al., 1995, file *w01ae.wav*) of length $N = 4879$, sampled at 16kHz. A vowel is a typical quasi-periodic signal where the pitch (or fundamental frequency) is fixed but the repeated pattern varies through time.

We encode our assumptions about this sound by constructing quasi-periodic kernels. Such kernels can be obtained as the product of a periodic kernel k_p and a Matérn $1/2$ kernel $k_{1/2}$ whose lengthscale controls the rate of change of the periodic pattern (Solin and Särkkä, 2014). We construct periodic kernels of varying complexity as weighted sums of cosine kernels in harmonic ratio of frequencies $k_p^J(\tau) = \sum_{j=1}^J \gamma_j^2 \cos(2\pi f_0 j \tau)$ where f_0 is the fundamental frequency and γ_j controls the magnitude of each harmonic component. Each harmonic increases the state dimension by 2, so the resulting kernel $k^J(\tau) = k_p^J(\tau)k_{1/2}(\tau)$ has a state dimension of $d = 2J$.

We then perform approximate inference in settings where we both vary the model complexity (using 1 to 4 harmonics) and the flexibility of the variational distribution by increasing M in powers of 2 (from 16 to 512). For each setting, we optimise the variational parameters and use the divergence $KL[q(\mathbf{s}(\cdot))|p(\mathbf{s}(\cdot)|\mathbf{y})]$ as a

performance metric. We also compare the execution time of the evaluation of the gradient of \mathcal{L} with respect to the variational parameters. We use similar implementations of SVGP and S²VGP as in the previous section.

Results are displayed in Figure 4. As a function of the number of inducing points M , the KL decreases much faster for S²VGP (a). This is because each inducing state contains more information about the process than an inducing evaluation. Both methods reduce the KL following a similar trend, with a small advantage for SVGP when compared against the actual number of inducing variables (b). However, as summarised in Table 1, storage and computational complexities of S²VGP grow linearly with M but are respectively quadratic and cubic in M for SVGP. Strikingly, the scaling $\mathcal{O}((N+M)d^3)$ for the gradient evaluation means the cost of adding extra inducing points in a S²VGP model is independent of the size of the dataset. For large N , it is thus possible to increase the number of inducing variable with very little impact on the computational time since the latter is dominated by the $\mathcal{O}(Nd^3)$ term as demonstrated by the vertical lines in (d). Note that this is not the case for SVGP which scales as $\mathcal{O}(NM^2 + M^3)$.

A visual illustration of an inference is given in Figure 5 ($N = 4541, M = 318$), where we have also removed part of the signal to demonstrate the out-of-sample predictive ability of S²VGP. Our method correctly interpolates between the periodic patterns at both ends of the missing data region with highest predictive uncertainty in the middle of this region.

4.3 Additive regression

This experiment illustrates three core capabilities of the proposed S²VGP algorithm: (i) it can deal with large datasets; (ii) it allows minibatching with rather large batch-size; (iii) it is not restricted to problems with 1-dimensional inputs. We also show that the proposed method is competitive in terms of performance.

The airline delay dataset consists of flight details (route distance, airtime, aircraft age, etc.) for every commercial flight in the USA for the year 2008. We use the same $c = 8$ covariates \mathbf{x} as in (Hensman et al., 2013) to predict the delay y of the aircraft at landing.

We perform regression from $\mathbb{R}^c \rightarrow \mathbb{R}$ under the modelling assumption that the delay is additive, i.e. $f(\mathbf{x}) = \sum_i f_i(x^{(i)})$, where $x^{(i)} \in \mathbb{R}$. We set GP priors over each function $f_i \sim \mathcal{GP}(0, k_i)$, where k_i are Matérn $3/2$ kernels. We propose a mean-field approximation to the posterior over processes $q(f_1, \dots, f_c) = \prod_i q^{(i)}(f_i)$, where each process is approximated using our S²VGP parameterisation. Details are given in Appendix C.3.

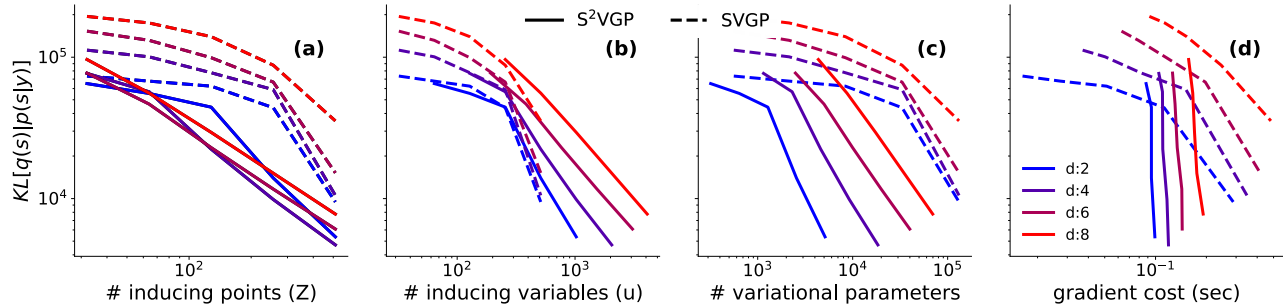


Figure 4: Comparison of sparse variational GP regression using classic inducing points and state-space features on an audio time series (see Fig. 5). Each line corresponds to a given inference method (line style) and model complexity (color) and shows the KL divergence for a an increasing number of inducing points. The proposed method is more accurate for a given number of inducing points, requires less variational parameters to reach a good accuracy, and is extremely fast (especially for large M).

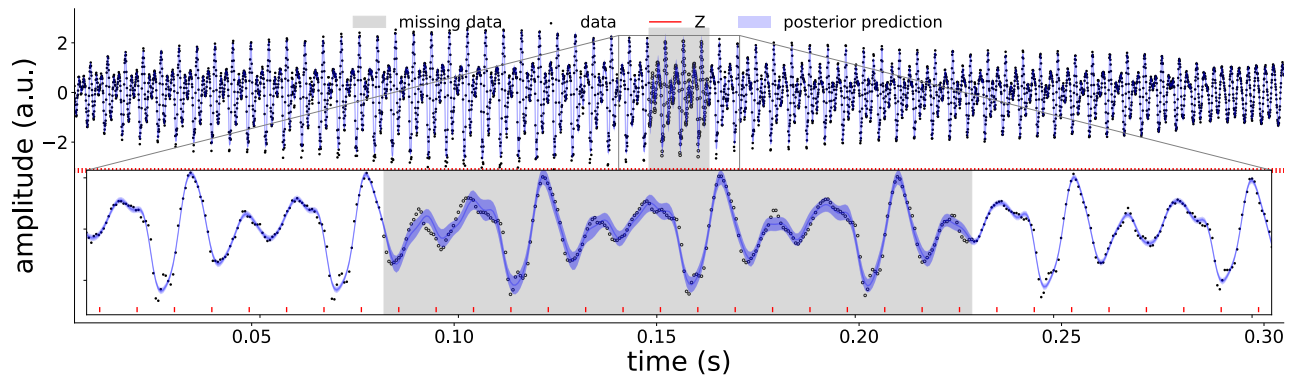


Figure 5: Vowel waveform with missing data and fit using S^2VGP . Black dots are the data. Blue line and shaded area correspond to the posterior predictions. Red vertical lines are the locations of the inducing points.

We compare the MSE obtained with two different optimisation schemes: the first one optimises all parameters with Adam (Kingma and Ba, 2015), while the second uses natural gradient for the variational parameters and Adam for the remaining ones. All learning rates are set to a constant: $\gamma_{matgrads} = 0.01$ and $\gamma_{Adam} = 0.01$. Given the size of the dataset, we use minibatches of 10k points for the two optimisers when $N \geq 10^6$.

As illustrated in Table 2, natural gradient provides the best performances when the number of observations is small, but as expected it suffers from the minibatching on larger datasets. On the other hand, Adam performs similarly, even when it only has access to sub-samples of the data. The proposed approach has similar accuracy to the state of the art (Hensman et al., 2018). A graphical version of the results in this table is given in Appendix (Figure 9).

4.4 Time warping with deep GPs

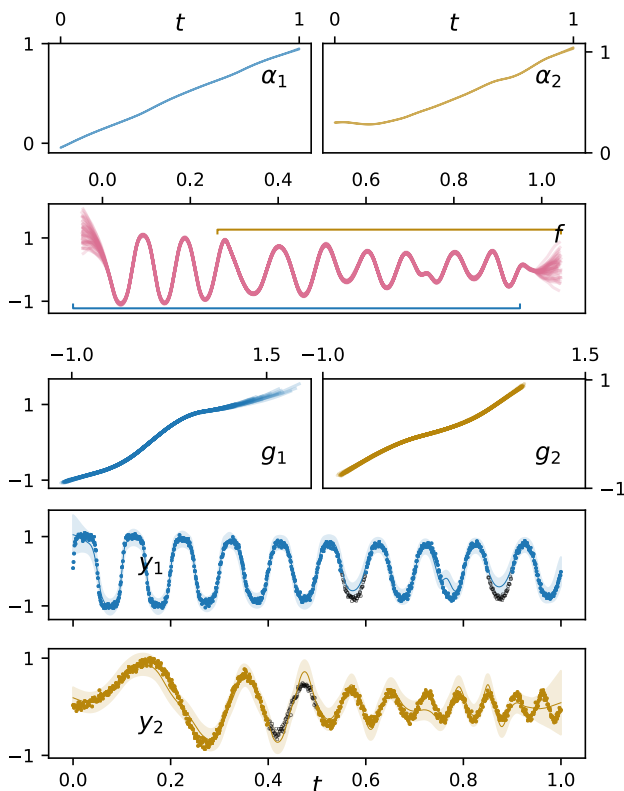
In this section we demonstrate the ability of S^2VGP to perform variational inference in a deep-GP

model (Damianou and Lawrence, 2013), where inference is performed by following the approach presented in Salimbeni and Deisenroth (2017). We consider the problem of data alignment and focus on reproducing the results from (Kaiser et al. (2018, Sec. 4.1), see Appendix C.4). The dataset consists of two times series generated by a three layer model and further corrupted by additive Gaussian noise so that $\mathbf{y}_k = g_k(f(a_k(t))) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. The function f is a sine wave shared across the two observed series indexed by $k = \{1, 2\}$. The functions a_k are time-warping functions, with a_1 being the identity and a_2 a quadratic function; while g_k are output distortions applied to f , with g_1 the hyperbolic tangent and g_2 the identity. Parts of the observed sequences have been removed at different locations to assess the generalisation performance of the model. Compared to the setting of Kaiser et al. (2018), we double both the frequency of the true f and the number of observations.

The goal here is to infer all five functions a_1, a_2, f, g_1, g_2 under the true model structure. We place GP priors on all five functions with Matérn $3/2$ kernels. We fur-

Table 2: Predictive mean squared errors (MSEs) and negative log predictive densities (NLPDs) with one standard deviation on the airline arrival delays experiment.

N	10k		100k		1m		$\approx 6m$	
	MSE	NLPD	MSE	NLPD	MSE	NLPD	MSE	NLPD
VFF (Hensman et al., 2018)	0.89 ± 0.15	1.36 ± 0.09	0.82 ± 0.05	1.32 ± 0.03	0.83 ± 0.01	1.336 ± 0.008	0.827 ± 0.004	1.324 ± 0.003
S ² VGP (Adam)	0.96 ± 0.13	1.40 ± 0.07	0.83 ± 0.07	1.33 ± 0.04	0.81 ± 0.02	1.316 ± 0.010	0.822 ± 0.006	9.22 ± 20.4
S ² VGP (Adam+Natgrads)	0.90 ± 0.12	1.36 ± 0.07	0.81 ± 0.07	1.32 ± 0.04	0.82 ± 0.02	1.321 ± 0.010	0.829 ± 0.006	1.325 ± 0.004


 Figure 6: Data alignment with S²VGP layers. The top three panels show samples from the inferred functions; the bottom two show the observed data \mathbf{y} (coloured dots) and posterior predictions for the whole sequence for each output. Black circles indicate missing data.

ther introduce linear mean functions on all priors to avoid pathologies while propagating samples through the layers (Salimbeni and Deisenroth, 2017). We use $M = [50, 100, 50]$ inducing points at each layer with the inducing inputs placed on a linear grid and kept fixed (all functions within a layer share the inducing input locations). We use natural gradients to learn the variational parameters of the approximate distributions in each layer, with the learning rate set to $\gamma_{natgrads} = 0.001$. The hyper-parameters of the kernel, the mean functions and the likelihood noise are learnt using Adam (Kingma and Ba, 2015) with exponential decaying learning rate (initialised from $\gamma_{Adam} = 0.001$).

Results of the inference are shown in Figure 6. The inferred a_k and g_k functions (first and third row) share the same characteristics as the corresponding ground truth functions. The function f is also recovered with the correct dampening (second row). More interestingly, we see that due to the time warping of the first layer, the model has successfully learnt to reconstruct each observed sequence from different sections of f . Finally, in the lower two panels we report the model predictions highlighting how the learnt model is rightfully more uncertain in the missing data region. Uncertainty decomposition across layers is discussed in Appendix C.4.

5 Discussion and conclusions

The proposed doubly sparse variational Gaussian processes combines the variational sparse approximations for Gaussian processes with the state-space representation of the process. It inherits its appealing tractability from the variational approach, but has the representative power and computational scalability of state-space representations. Unlike other state-space GP methods, it is readily applicable to deep GP settings and supports mini-batch stochastic training. We showed that our framework can be used to approximate functions with more than one input variable while preserving the computational gain of state-space models.

To ease the optimisation of our variational objective, we derived natural gradient updates for the class of multivariate normal distribution with banded precisions. Although the objective is non-convex, this leads to few shot inference in the conjugate setting and empirically improves optimisation in non-conjugate settings. To further improve the applicability of S²VGP, different sub-optimal variational parameterisations could be used for $q_{\mathbf{u}}$ leading to better behaved objectives or additional scalability improvement, albeit at the cost of reduced expressivity. Another route of improvement could consist in making a further steady-state approximation to the posterior to reduce the scaling with the GP state dimension from cubic to quadratic as in (Solin et al., 2018). As in (Nickisch et al., 2018), further computational gains could be achieved by using interpolations when we compute the SSM parameters.

References

- Mauricio Alvarez and Neil D. Lawrence. Sparse convolved Gaussian processes for multi-output regression. In Advances in Neural Information Processing Systems, pages 57–64, 2009.
- Shun-Ichi Amari. Natural gradient works efficiently in learning. Neural computation, 10(2):251–276, 1998.
- Matthias Bauer, Mark van der Wilk, and Carl E. Rasmussen. Understanding probabilistic sparse Gaussian process approximations. In Advances in Neural Information Processing Systems, pages 1533–1541, 2016.
- David R. Burt, Carl E. Rasmussen, and Mark van der Wilk. Rates of convergence for sparse variational Gaussian process regression. In International Conference on Machine Learning, 2019.
- Andreas Damianou and Neil D. Lawrence. Deep Gaussian processes. In Artificial Intelligence and Statistics, pages 207–215, 2013.
- Nicolas Durrande, Vincent Adam, Lucas Bordeaux, Stefanos Eleftheriadis, and James Hensman. Banded matrix operators for Gaussian Markov models in the automatic differentiation era. In Artificial Intelligence and Statistics, pages 2780–2789, 2019.
- GPflow. VGP model in gpflow. <https://github.com/GPflow/GPflow/blob/develop/gpflow/model>.
- Alexander Grigorievskiy, Neil D. Lawrence, and Simo Särkkä. Parallelizable sparse inverse formulation Gaussian processes (SpInGP). In International Workshop on Machine Learning for Signal Processing, pages 1–6, 2017.
- James Hensman, Nicolò Fusi, and Neil D. Lawrence. Gaussian processes for big data. In Uncertainty in Artificial Intelligence, pages 282–290, 2013.
- James Hensman, Nicolas Durrande, and Arno Solin. Variational Fourier features for Gaussian processes. Journal of Machine Learning Research, 18(151):1–52, 2018.
- James Hillenbrand, Laura A. Getty, Michael J. Clark, and Kimberlee Wheeler. American English vowels dataset, 1995.
- Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. Machine learning, 37(2):183–233, 1999.
- Markus Kaiser, Clemens Otte, Thomas Runkler, and Carl Henrik Ek. Bayesian alignments of warped multi-output Gaussian processes. In Advances in Neural Information Processing Systems, pages 6995–7004, 2018.
- Rudolph E. Kalman. A new approach to linear filtering and prediction problems. Transactions of the ASME–Journal of Basic Engineering, 82:35–45, 1960.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In International Conference on Learning Representations, 2015.
- Miguel Lázaro-Gredilla and Aníbal Figueiras-Vidal. Inter-domain Gaussian processes for sparse inference using inducing features. In Neural Information Processing Systems, 2009.
- Luigi Malagò and Giovanni Pistone. Information geometry of the Gaussian distribution in view of stochastic optimization. In Foundations of Genetic Algorithms, pages 150–162, 2015.
- Alexander G. de G. Matthews, James Hensman, Richard E. Turner, and Zoubin Ghahramani. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. Journal of Machine Learning Research, 51:231–239, 2016.
- Hannes Nickisch, Arno Solin, and Alexander Grigorevskiy. State space Gaussian processes with non-Gaussian likelihood. In International Conference on Machine Learning, pages 3789–3798, 2018.
- Manfred Opper and Cédric Archambeau. The variational Gaussian approximation revisited. Neural computation, 21(3):786–792, 2009.
- Carl E. Rasmussen and Christopher K. I. Williams. Gaussian processes for machine learning. The MIT Press, Cambridge, MA, USA, 2006.
- Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep Gaussian processes. In Advances in Neural Information Processing Systems, pages 4588–4599, 2017.
- Hugh Salimbeni, Stefanos Eleftheriadis, and James Hensman. Natural gradients in practice: Non-conjugate variational inference in Gaussian process models. In Artificial Intelligence and Statistics, pages 689–697, 2018.
- Simo Särkkä. Bayesian filtering and smoothing, volume 3. Cambridge University Press, 2013.
- Simo Särkkä and Robert Piché. On convergence and accuracy of state-space approximations of squared exponential covariance functions. In 2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP), pages 1–6. IEEE, 2014.
- Simo Särkkä and Arno Solin. Applied stochastic differential equations. Cambridge University Press, 2019.
- Arno Solin and Simo Särkkä. Explicit link between periodic covariance functions and state space models.

In Artificial Intelligence and Statistics, pages 904–912, 2014.

Arno Solin, James Hensman, and Richard E. Turner. Infinite-horizon Gaussian processes. In Advances in Neural Information Processing Systems, pages 3486–3495, 2018.

Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In Artificial Intelligence and Statistics, pages 567–574, 2009.

Mark van der Wilk, Carl E. Rasmussen, and James Hensman. Convolutional Gaussian processes. In Advances in Neural Information Processing Systems, pages 2849–2858, 2017.