
Free-Form Variational Inference for Gaussian Process State-Space Models

Xuhui Fan¹ Edwin V. Bonilla² Terence J. O’Kane³ Scott A. Sisson⁴

Abstract

Gaussian process state-space models (GPSSMs) provide a principled and flexible approach to modeling the dynamics of a latent state, which is observed at discrete-time points via a likelihood model. However, inference in GPSSMs is computationally and statistically challenging due to the large number of latent variables in the model and the strong temporal dependencies between them. In this paper, we propose a new method for inference in Bayesian GPSSMs, which overcomes the drawbacks of previous approaches, namely over-simplified assumptions, and high computational requirements. Our method is based on free-form variational inference via stochastic gradient Hamiltonian Monte Carlo within the inducing-variable formalism. Furthermore, by exploiting our proposed variational distribution, we provide a collapsed extension of our method where the inducing variables are marginalized analytically. We also showcase results when combining our framework with particle MCMC methods. We show that, on six real-world datasets, our approach can learn transition dynamics and latent states more accurately than competing methods.

1. Introduction

State-space models (SSMs; [Murphy, 2023](#), Ch. 29) characterize the underlying dynamics of a latent state given a set of observations via a transition (or evolution) function and an observation model. As a modeling framework, they provide a general approach for understanding time-series data ([Kitagawa, 1987](#)) and for data assimilation problems ([Katzfuss et al., 2016](#)). Applications of SSMs abound and span diverse areas such as econometrics ([Tsay, 2005](#)), meteorology ([Hernandez et al., 1991](#)), control engineering ([Ogata et al.,](#)

2010) and neuroscience ([Brown et al., 1998](#)).

In this paper, we focus on Bayesian SSMS, where the transition function describing the dynamics of the system is given a prior distribution. A paradigmatic example of Bayesian SSMS are Gaussian process state-space models (GPSSMs; [Frigola, 2015](#)), where this prior distribution is a Gaussian process (GP; [Williams & Rasmussen, 2006](#)). Due to their Bayesian non-parametric nature, GPSSMs represent a principled and flexible approach to Bayesian SSMS.

However, the flexibility of Gaussian processes (GPs) adds significant computational and statistical challenges to the already difficult problem of inference in Bayesian SSMS. Indeed, even for non-Bayesian SSMS, standard problems such as filtering, smoothing, and prediction are, in general, analytically intractable¹. Having a GP prior over the transition function in SSMS increases the number of latent variables significantly; incorporates strong (and potentially long-term) dependencies across states; and introduces a cubic time complexity as a function of the number of observations.

Within the GP community, significant advances have been made addressing the computational issues in GP regression and classification problems, most notably using inducing-variable approximations ([Titsias, 2009](#); [Hensman et al., 2015](#); [Rossi et al., 2021](#)) but also random-feature expansions ([Cutajar et al., 2017](#); [Marmin & Filippone, 2022](#)) and, more recent innovative approaches such as the Vecchia approximation ([Sauer et al., 2022](#)). Although these approximations are applicable to GP-based dynamic models, the challenges above remain prevalent within the context of GPSSMs.

Nevertheless, previous approaches have developed insightful and practical algorithms for inference in GPSSMs, mainly based on inducing-variable approximations, which is also our main underpinning methodology for scalable GPs. In outlining the most relevant approaches, our main object of interest is the approximate joint posterior over state trajectories $\mathbf{x}_{0:T}$ and inducing variables \mathbf{u} , $q(\mathbf{x}_{0:T}, \mathbf{u})$, where $T + 1$ is the length of the trajectory. We consider two main aspects of this joint distribution: (i) whether the dependencies between state trajectories and inducing variables are

¹With the notable exception of linear state transition and observation models with additive Gaussian noise, for which the solution to the optimal filtering problem is given by the Kalman filter ([Kalman, 1960](#)).

¹University of Newcastle, Australia ²CSIRO’s Data61, Australia ³CSIRO’s Environment, Australia ⁴University of New South Wales, Australia. Correspondence to: Xuhui Fan <xhfan.ml@gmail.com>.

captured and (ii) whether their corresponding distributions are unconstrained, i.e., not restricted to a sub-optimal parametric form. The seminal variational Gaussian process state space model (VGPSSM) proposed by Frigola et al. (2014) as well as the subsequent identifiable Gaussian process state space model (IGPSSM) of Eleftheriadis et al. (2017) use mean-field approaches, therefore, ignoring the posterior dependencies between state trajectories and inducing variables. The more recent methods, namely the probabilistic recurrent state space model (PRSSM) of Doerr et al. (2018) and the variationally coupled dynamic trajectories (VCDT) of Ialongo et al. (2019) introduce couplings across state trajectories and inducing variables. However, their posteriors are constrained to be Gaussians. Thus, these previous works have either assumed independence between state trajectories and inducing variables or imposed strong parametric constraints in their corresponding posteriors or both.

As shown by Ialongo et al. (2019), a mean-field posterior can yield poor practical performance. Similarly, a Gaussian assumption on the state posterior or the inducing variable posterior is also very strong and, by definition, will not generally capture the true posterior even in the limit of infinite computation. To address these issues, we propose a free-form variational inference approach to posterior estimation in GPSSMs that models the full joint distribution over states and inducing variables, $q(\mathbf{x}_{0:T}, \mathbf{u})$, without any mean-field or parametric assumptions. We refer to our method as free-form variational dynamics (FFVD) and summarize the major differences between its posterior assumptions and those of previous approaches in Table 1. Below we describe our contributions in more detail.

(i) Flexible posterior: We develop an inference algorithm for GPSSMs based on stochastic gradient Hamiltonian Monte Carlo (SGHMC; Chen et al., 2014; Havasi et al., 2018), which represents the posterior over states and inducing variables using samples. Further, our approach lifts the limitations of previous variational approaches to GPSSMs, as they have ignored couplings in the posterior between the inducing variables and the state trajectories or have assumed a constrained parametric form for this posterior. Regarding this, our approach naturally captures the posterior correlations between states and inducing variables; does not constrain this posterior to any parametric form, and, is scalable to a large number of observations.

(ii) Collapsed inference that accelerates convergence: we show that (i) our formulation allows us to collapse the inducing variables \mathbf{u} , (ii) sample from the lower-dimensional marginal $q(\mathbf{x}_{0:T})$ and, at the end of the sampling procedure, (iii) obtain samples from the conditional $q(\mathbf{u}|\mathbf{x}_{0:T})$, for which we derive a closed-form expression. We show that collapsing accelerates convergence significantly.

(iii) Extensions with particle Markov chain Monte Carlo

(PMCMC): we further investigate whether more elaborate inference algorithms, such as PMCMC (Andrieu et al., 2010), that account for the sequential nature of the problem can provide more accurate posteriors.

(iv) State-of-the-art performance: we showcase the properties and benefits of our approach compared to previous methods such as VGPSSM, PRSSM and VCDT in a synthetic example and six system identification benchmarks. Overall, our method provides state-of-the-art performance when evaluated on these problems, while having comparable computational requirements to previous approaches.

We open-source our proposed FFVD approach, together with baseline methods reported in the experiments: <https://github.com/xuhui-fan/FFVD>.

2. Inference in Gaussian Process Models

Gaussian processes (GPs) are priors over functions where every subset of function values follow a Gaussian distribution. We use $f(\mathbf{x}) \sim GP(m_f(\mathbf{x}), \kappa_f(\mathbf{x}, \mathbf{x}^l; \boldsymbol{\theta}))$ to denote that f is distributed according to a GP with mean function $m_f(\cdot)$ and covariance function $\kappa_f(\cdot, \cdot; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ are referred to as the GP hyper-parameters². By definition, a GP prior over functions implies a finite prior over T function values $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_T)]^\top$, i.e., $\mathbf{f} \sim \mathcal{N}(\mathbf{f}; \mathbf{m}, \mathbf{K})$, where \mathbf{m} and \mathbf{K} are obtained by evaluating the mean function and covariance function at all the inputs $\mathbf{X} := \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$.

In supervised learning settings, we are given input-output observations $f_{\mathbf{x}_t}, y_t, g_{t=1}^T$ and a conditional likelihood model $p(\mathbf{y}|\mathbf{f})$. Inference involves estimating the posterior distribution $p(\mathbf{f}|\mathbf{y}, \mathbf{X})$ and the hyper-parameters $\boldsymbol{\theta}$ from data. Consequently, we can use these to estimate the posterior predictive distribution at a new point \mathbf{x}_* , i.e., $p(f(\mathbf{x}_*)|\mathbf{X}, \mathbf{y}, \mathbf{x}_*)$. Notoriously, these tasks have cubic time complexity as a function of the number of training observations, i.e., $O(T^3)$, arising from algebraic operations involving the computation of the inverse covariance and its log determinant. This motivates the need for sparse approximations.

2.1. Sparse GP Approximations via Inducing Variables

Various approximations have been proposed to deal with the cubic computational complexity of inference in GP models. In this paper, we focus on inducing-variable approximations based on variational inference, as originally proposed by Titsias (2009) and made scalable to very large datasets by Hensman et al. (2013). The main idea of these approximations is to augment the space of function values with a set of M inducing variables $\mathbf{u} := \{u_i, g\}$ and their correspond-

²We note here that the mean function can also have additional parameters but, in this work, we have simply assumed identity mean functions.

Table 1. Comparison across methods in terms of their assumptions on the variational distribution. The rows refer to whether the variational posterior captures the dependencies between the state trajectories and the inducing variables (coupled $q(\mathbf{x}_{0:T}, \mathbf{u})$); the distribution over states is unconstrained (unconstrained $q(\mathbf{x}_{0:T} | \mathbf{u})$ or $q(\mathbf{x}_{0:T})$), i.e., not restricted to a sub-optimal parametric form; and whether the distribution over the inducing variables is also unconstrained (unconstrained $q(\mathbf{u})$). Our method is referred to as FFVD.

	VGPSSM	IGPSSM	PRSSM	VCDDT	FFVD
Coupled $q(\mathbf{x}_{0:T}, \mathbf{u})$	✗	✗	✓	✓	✓
Unconstrained $q(\mathbf{x}_{0:T} \mathbf{u})$ or $q(\mathbf{x}_{0:T})$	✓	✗	✗	✗	✓
Unconstrained $q(\mathbf{u})$	✓	✗	✗	✗	✓

ing inducing inputs $\mathbf{Z} := \tilde{f}_{\mathbf{z}_i} g$. Thus, inference involves estimating the posterior $q(\mathbf{u}, \mathbf{f}) = p(\mathbf{u}, \mathbf{f} | \mathbf{X}, \mathbf{y})$ and the GP hyper-parameters via variational inference. Under the assumption that $q(\mathbf{u}, \mathbf{f}) := q(\mathbf{u})p(\mathbf{f} | \mathbf{u})$ where $p(\mathbf{f} | \mathbf{u})$ is the conditional prior, the variational objective, the so-called evidence lower bound (ELBO), decomposes over the observations, and inference can be carried out with time complexity of $O(M^3)$, providing significant advantages when $M \gg T$.

3. Gaussian Process State-Space Models

Let us assume we are given a time series of T multi-dimensional observations $\mathbf{y}_{1:T}$ and denote their corresponding latent states with $\mathbf{x}_{0:T}$, where \mathbf{x}_0 is the initial state. Here we denote the time series $\mathbf{x}_{t_1:t_2} := \tilde{f}_{\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_2}} g$, and similarly for $\mathbf{y}_{1:T}$. In general, $\mathbf{x}_t \in \mathbb{R}^{d_x}$ and $\mathbf{y}_t \in \mathbb{R}^{d_y}$. Gaussian process state-space models (GPSSMs) formulate a discrete-time state-space model (SSM) where the transition dynamics is given by a GP. The full generative process is:

$$\mathbf{x}_0 \sim p(\mathbf{x}_0), f(\mathbf{x}) \sim GP(m_f(\mathbf{x}), \kappa_f(\mathbf{x}, \mathbf{x}^0; \boldsymbol{\theta})), \quad (1)$$

$$\mathbf{f}_t := f(\mathbf{x}_{t-1}), \quad \mathbf{x}_t | \mathbf{f}_t \sim N(\mathbf{x}_t; \mathbf{f}_t, \mathbf{Q}), \quad (2)$$

$$\mathbf{y}_t | \mathbf{x}_t \sim p(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}), \quad (3)$$

where \mathbf{Q} is the transition process covariance, and $\boldsymbol{\theta}$ is the vector of parameters of the conditional likelihood $p(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta})$. Although our framework does not make any parametric assumptions about this conditional likelihood, in our experiments in § 7, we adopt the same setting as in previous works (Ialongo et al., 2019; Doerr et al., 2018) and set $p(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) := N(\mathbf{y}_t; \mathbf{C}\mathbf{x}_t + \mathbf{d}, \mathbf{R})$, with $\boldsymbol{\theta} = \{f\mathbf{C}, \mathbf{d}, \mathbf{R}\}$, where \mathbf{C}, \mathbf{d} are the weights and bias of the linear transformation on \mathbf{x}_t and \mathbf{R} is the observation covariance.

3.1. Joint Distribution

As shown by Frigola (2015), to sample \mathbf{f}_t , instead of conditioning on an infinite-dimensional function, we can condition only on the transitions seen up to (but not including) time t , i.e., $\tilde{f}(\mathbf{x}_{i-1}, \mathbf{f}_i) g_{i=1}^t$. We can then write the joint

distribution over latent variables and observations as:

$$p(\mathbf{y}_{1:T}, \mathbf{x}_{0:T}, \mathbf{f}_{1:T}) = \prod_{t=1}^T p(\mathbf{x}_0) p(\mathbf{f}_t | \mathbf{f}_{1:t-1}, \mathbf{x}_{0:t-1}) p(\mathbf{x}_t | \mathbf{f}_t) p(\mathbf{y}_t | \mathbf{x}_t), \quad (4)$$

where $p(\mathbf{x}_0), p(\mathbf{x}_t | \mathbf{f}_t), p(\mathbf{y}_t | \mathbf{x}_t)$ are defined as above and, for the edge case of $t = 1$, we have: $p(\mathbf{f}_1 | \mathbf{x}_0) = N(\mathbf{f}_1; m_f(\mathbf{x}_0), \kappa_f(\mathbf{x}_0, \mathbf{x}_0; \boldsymbol{\theta}))$. Furthermore, we recognize each conditional distribution $p(\mathbf{f}_t | \mathbf{f}_{1:t-1}, \mathbf{x}_{0:t-1})$ in Eq. (4) as the GP prediction at a single point \mathbf{x}_{t-1} using noiseless outputs $\mathbf{f}_{1:t-1}$ and inputs $\mathbf{x}_{0:t-2}$,

$$p(\mathbf{f}_t | \mathbf{f}_{1:t-1}, \mathbf{x}_{0:t-1}) = p(\mathbf{f}_t | \mathbf{x}_{t-1}, \mathbf{f}_{1:t-1}, \mathbf{x}_{0:t-2}) \quad (5)$$

$$= N(\mathbf{f}_t; \mathbf{m}_f, \boldsymbol{\Sigma}_f), \quad (6)$$

with conditional mean and covariance given by

$$\mathbf{m}_f = \mathbf{m}_{t-1} + \mathbf{K}_{t-1,0:t-2} \mathbf{K}_{0:t-2}^{-1} (\mathbf{f}_{1:t-1} - \mathbf{m}_{0:T-2}) \quad (7)$$

$$\boldsymbol{\Sigma}_f = \mathbf{K}_{t-1} - \mathbf{K}_{t-1,0:t-2} \mathbf{K}_{0:t-2}^{-1} \mathbf{K}_{0:t-2,t-1}, \quad (8)$$

where the subscript notation indicates the mean vectors and covariance matrices obtained from evaluating the mean function and covariance function, respectively, at the corresponding ranges, $\mathbf{m}_{t_1:t_2} := m_f(\mathbf{x}_{t_1:t_2})$, $\mathbf{m}_t := \mathbf{m}_{t:t}$, $\mathbf{K}_{t_1:t_2,t_3:t_4} := \kappa_f(\mathbf{x}_{t_1:t_2}, \mathbf{x}_{t_3:t_4}; \boldsymbol{\theta})$, $\mathbf{K}_{t_1:t_2} := \mathbf{K}_{t_1:t_2,t_1:t_2}$ and $\mathbf{K}_t := \mathbf{K}_{t:t}$.

3.2. Multidimensional Latent States & Control Inputs

In the case of multidimensional latent states, i.e., $d_x > 1$, we assume independent GPs on each dimension, each with its own mean function and covariance function. Because of this independence assumption, each GP only has to condition on its own function evaluations. Therefore, for simplicity in the notation, we do not index the means and covariances with respect to their dimension d and consider the underlying GPs, their means, and covariance functions as multi-dimensional. Furthermore, as we shall see in § 7, our experiments consider additional control inputs (\mathbf{a}_t) that affect the transitions in a Markovian way. This is easy to incorporate in our framework by augmenting our input space and, therefore, indexing the GPs in the higher-dimensional space given by

the concatenation $[x^> ; a^>]^>$. However, since the control inputs are fixed and deterministic, we do not need to include them as part of our inference method and do not make them explicit in our subsequent mathematical development.

3.3. Sparse GPSSM

In order to make the GPSSM inherently scalable, similarly to the standard supervised regression setting described in § 2, we augment the full GPSSM model with M inducing variables and corresponding inducing inputs $u := f u_i g_{i=1}^M$ and $Z := f z_i g_{i=1}^M$. This gives rise to the sparse GPSSM model:

$$p(y_{1:T}; x_{0:T}; f_{1:T}; u; j Z) = p(u; j Z) p(x_0) \prod_{t=1}^T p(f_t; j f_{1:t-1}; x_{0:t-1}; u; Z) p(x_t; j f_t) p(y_t; j x_t); \quad (9)$$

where the prior over the inducing variables is determined by the GP prior, i.e., $p(u; j Z) = N(m_Z; K_Z)$ with $m_Z := m_f(Z)$ and $K_Z := k_f(Z; Z; \cdot)$. Besides having a prior over the inducing variables, the main difference with our previous full-model formulation of § 3.1 is that the conditional distributions over f_t have been augmented with inducing variables u and corresponding inducing inputs Z . Each conditional $p(f_t; j f_{1:t-1}; x_{0:t-1}; u; Z) = p(f_t; j x_{t-1}; f_{1:t-1}; u; x_{0:t-2}; Z)$, is the predictive GP regression distribution of f_t at test input x_{t-1} when observing $(f_{1:t-1}; u)$ at their respective locations $(x_{0:t-2}; Z)$.

4. Free-Form Variational Inference

In this section, we develop a posterior estimation method using variational inference (VI). We show that under a specific form of the approximate posterior, we can estimate the joint posterior over $(x_{0:T}; u; g)$ in free-form, i.e., optimally, without making any assumptions such as independence between state trajectories and inducing variables, typical of mean-field approaches (Frigola et al., 2014; Eleftheriadis et al., 2017), or constraining the form of the posterior to sub-optimal parametric forms (Ialongo et al., 2019; Doerr et al., 2018).

4.1. Variational Family and Evidence Lower Bound

Variational inference is underpinned by the maximization of the evidence lower bound (ELBO), which is equivalent to minimizing the Kullback-Leibler (KL) divergence between the approximate posterior and the true posterior. This objective is given by

$$L_{ELBO}(q) := E_{q(x_{0:T}; u; f_{1:T})} [\log p(y_{1:T}; x_{0:T}; f_{1:T}; u; j Z) \log q(x_{0:T}; u; f_{1:T})]; \quad (10)$$

where $p(y_{1:T}; x_{0:T}; f_{1:T}; u; j Z)$ is the joint distribution in Eq. (9) and $q(x_{0:T}; u; f_{1:T})$ is our proposed approximate joint posterior, which we define as

$$q(f_{1:T}; j x_{0:T}; u) := \prod_{t=1}^T p(f_t; j f_{1:t-1}; x_{0:t-1}; u; Z); \quad (11)$$

$$q(x_{0:T}; u; f_{1:T}) := q(x_{0:T}; u) q(f_{1:T}; j x_{0:T}; u); \quad (12)$$

It is easy to show that the ELBO, as defined in Eq. (10), is a lower bound on the log marginal likelihood, i.e., $L_{ELBO}(q) \leq \log p(y; j Z)$. We use the prior distribution of $f_{1:T}$ as its variational distribution in Eq. (11), which follows the same settings as all previous scalable variational approaches (that we are aware of) VGPSSM, IGPSSM, PRSSM and VCDT. It is necessary to make the variational framework scalable as a function of the number of observations, which will yield an inference algorithm that alleviates the cubic time complexity on the number of observations (T) and, therefore, scales up to large datasets. Crucially, our proposed joint variational distribution in Eq. (12), will allow us to derive an optimal variational distribution $q(x_{0:T}; u)$ in free-form, without imposing any parametric constraints over it and, instead, represent it via samples. We will describe this in the next section.

4.2. Evidence Lower Bound Maximization

Our first step is to expand the expression for the ELBO in Eq. (10) using our joint model distribution in Eq. (9) and our proposed variational distribution in Eqs. (11) and (12). We first note that our definition of the variational distribution in Eq. (12) uses the same conditional prior as in the joint distribution in Eq. (9). Therefore, this term cancels out, avoiding the computation of operations on fully-coupled high-dimensional distributions over latent functions. Thus, we have that:

$$L_{ELBO}(q) = \int q(x_{0:T}; u) \left(\log q(x_{0:T}; u) + \log p(x_0) p(u; j Z) \prod_{t=1}^T p(y_t; j x_t) + E_{q(f_{1:T}; j x_{0:T}; u)} \left[\log \prod_{t=1}^T p(x_t; j f_t) \right] \right) dx_{0:T} du; \quad (13)$$

4.2.1. ELBO MAXIMIZATION

Next we aim to maximize the ELBO functional above with respect to $q(x_{0:T-1}; u)$ subject to the constraint $\int q(x_{0:T-1}; u) dx_{0:T-1} du = 1$. We can do this by solving

the corresponding Euler-Lagrange equation:

$$\begin{aligned} & \frac{\partial}{\partial (x_{0:T}; u)} \left(\log q(x_{0:T}; u) + \sum_{t=1}^h \log p(x_0) p(u | Z) \prod_{i=1}^I p(y_t | x_t) \right) \\ & + \mathbb{E}_{q(f_{1:T} | x_{0:T}; u)} \sum_{t=1}^h \log p(x_t | f_t) = 0: \end{aligned} \quad (14)$$

By doing the corresponding derivatives we obtain

$$\begin{aligned} & \log q(x_{0:T}; u) + \sum_{t=1}^h \log p(x_0) p(u | Z) \prod_{i=1}^I p(y_t | x_t) \\ & + \mathbb{E}_{q(f_{1:T} | x_{0:T}; u)} \sum_{t=1}^h \log p(x_t | f_t) = 0: \end{aligned} \quad (15)$$

Now we note that the expectation above can be solved in closed form:

$$\begin{aligned} & \mathbb{E}_{q(f_{1:T} | x_{0:T}; u)} \sum_{t=1}^h \log p(x_t | f_t) \\ & = \sum_{t=1}^h \mathbb{E}_{p(f_t | x_{t-1}; u; Z)} \log p(x_t | f_t) \\ & = \sum_{t=1}^h \log N(x_t; u_{x_t}; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}); \end{aligned} \quad (16)$$

where $p(f_t | x_{t-1}; u; Z)$ is the GP predictive distribution over the function values f_t at locations x_{t-1} given the inducing variables u at inducing inputs Z , i.e.,

$$p(f_t | x_{t-1}; u; Z) = N(f_t; \mu_{x_t}; B_{t-1}); \quad (17)$$

$$\mu_{x_t} := m_{t-1} + A_{t-1}(u - m_Z); \quad (18)$$

Here we have defined

$$A_{t-1} := K_{t-1;Z} K_Z^{-1}; \quad (19)$$

$$B_{t-1} := K_{t-1} - K_{t-1;Z} K_Z^{-1} K_{Z;t-1}; \quad (20)$$

and the cross-covariance term $K_{t-1;Z} := k_f(x_{t-1}; Z; \theta)$ and similarly for $K_{Z;t-1}$. Finally, $\text{Tr}(\cdot)$ is the trace operator and, as defined at the beginning of §3, Q is the transition noise covariance.

4.2.2. OPTIMAL VARIATIONAL POSTERIOR

With this, we obtain the form of the optimal variational distribution $q(u; x_{0:T})$ up to a normalizing constant Z_q as:

$$\begin{aligned} \log q(u; x_{0:T}) &= \log p(u | Z) + \log p(x_0) \\ &+ \sum_{t=1}^h \log p(y_t | x_t) + \log N(x_t; \mu_{x_t}; Q) \\ &- \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}) + \log Z_q; \end{aligned} \quad (21)$$

Here we note that the function values $f_{1:T}$ have, effectively, been marginalized variationally. The optimal joint posterior over inducing variables and state trajectories depends on the prior over the inducing variables $p(u | Z)$ stemming from the GP functional prior, the prior over the initial state $p(x_0)$ and the conditional likelihood terms $p(y_t | x_t)$. It also depends on the resulting transitions mapping x_{t-1} to x_t via the densities $N(x_t; \mu_{x_t}; Q)$, where μ_{x_t} depends on x_{t-1} in a nonlinear way, as specified by Eq. (19). The final trace term, $\text{Tr}(Q^{-1} B_{t-1})$, can be seen as a regularization term acting on state transitions, encouraging higher transition variances and, therefore, helping prevent overfitting.

4.2.3. ALTERNATIVE PERSPECTIVE

An alternative way to obtain the optimal joint posterior over state trajectories and inducing variable is by bounding the true log joint marginal $\log p(y_{1:T}; x_{0:T}; u | Z)$ using Jensen's inequality. We give details of such an approach in Appendix E. This has been used by previous work in standard regression settings (see, e.g., Rossi et al., 2021, and references therein). However, our setting considers the more complex case of GPSSMs. Additionally, our development in this section shows much more clearly the optimal nature of the variational posterior, as we have obtained it via calculus of variations.

4.3. Posterior Sampling

Having the form of the optimal posterior in Eq. (21), we can then set the latent variables $\theta := f(u; x_{0:T}; g)$ and have $q(\theta) / q(u) = p(\theta | y_{1:T})$. Thus, we can draw samples from our approximate posterior using stochastic gradient Hamiltonian Monte Carlo (SGHMC; Chen et al., 2014; Havasi et al., 2018) and the energy function $U(\theta) = -\log p(\theta; y_{1:T}) = -\log p(\theta | y_{1:T}) + \log Z_q - \log q(\theta) + \log Z_q$. Using this procedure, samples from the target distribution can be obtained even with noisy gradients (e.g., with mini-batches) without requiring the evaluation of Metropolis ratios. Importantly, other variables such as hyper-parameters and inducing locations Z can be easily included in θ using suitable priors and incorporating them in our objective in Eq. (21).

4.3.1. COMPUTATIONAL COST & PRIOR WHITENING

An interesting aspect of GPSSM models is that, despite their apparent Markovian nature, sampling at time t requires conditioning on all the previous $t-1$ points. This is due to the non-parametric coupled prior over the transition function, making inference in the full model $\mathcal{O}(T^3)$ in time. Sparse variational inference approaches, such as those based on inducing variable approximations, still require expectations over entire trajectories and, unlike standard supervised i.i.d settings, their time complexity is inherently dependent on T . Evaluation of the stochastic gradient Hamiltonian Monte Carlo (SGHMC) objective in Eq. (21) for sampling in our free-form variational dynamic SFFVD algorithm is $\mathcal{O}(M^2T)$. this is the same cost as that attained for ELBO evaluation in variationally coupled dynamic trajectories (VCDT), which like our FFVD, models dependencies between state trajectories and inducing variables.

We may have a further comparison between our FFVD and the VCDT. Our SGHMC objective in Eq. (21) is very similar to the ELBO used in VCDT. The ELBO in VCDT requires expectations over $q(u)$ and $q(x_t; x_{t-1}|u)$. Given the factorization assumptions and the Gaussian constraints on these distributions, these expectations are estimated straightforwardly via Monte Carlo samples. The time complexity of evaluating Eq. (21) or ELBO in VCDT once (using one sample) is the same, i.e. $\mathcal{O}(M^3 + TM^2)$.

The overall time complexity of both algorithms, FFVD and VCDT, depends on (i) the number of samples (noting, again, that VCDT also requires samples from the approximate posterior to estimate the gradients of the ELBO) and (ii) the number of iterations (either the length of the SGHMC chain in FFVD or the number of epochs for gradient-based optimization in VCDT). As described in the appendix, our experimental setting followed closely that of the original VCDT paper, which uses $S = 100$ samples for training and $S = 10^5$ for predictions. Similarly, we use $S = 100$ samples for FFVD. Remarkably, the number of iterations in our experiments for FFVD was 50,000 while for VCDT was 200,000 to achieve convergence. Furthermore, our analysis in the appendix shows that, in fact, our FFVD algorithm converges in less than 10,000 iterations.

As in previous work (see, e.g., Hensman et al., 2015), we have observed that whitening the prior over the inducing variables improves the performance of our algorithm. See details of our whitening procedure and the resulting unnormalized log posterior in Appendices D and E.1.

4.4. Smoothing and Predictive Distributions

We are interested in estimating the smoothing distribution $p(x_{0:T} | y_{1:T})$ and the predictive distribution $p(y_{T+1:T} | y_{1:T})$ for $T^0 \leq T$. At the end of our SGHMC

procedure, we have S samples from our joint approximate posterior $q(x_{0:T}; u) = p(x_{0:T} | y_{1:T})$, i.e., $\{x_{0:T}^{(s)}; u^{(s)}\}_{s=1}^S$ and, therefore, the smoothing distribution (i.e., the marginal posterior over $x_{0:T}$) is readily available through this Monte Carlo approximation.

We can also make one-step-ahead predictions using our posterior samples. In particular, using Eqs. (2) and (17) we have that:

$$p(x_t | x_{t-1}; u; Z) = N(x_t; m_{t-1} + A_{t-1}(u - m_Z); B_{t-1} + Q); \quad (22)$$

for $t > T$. Thus, replacing the values of $A_{t-1}; B_{t-1}$ using Eqs. (19) and (20) we can make predictions for the next state using samples as

$$x_t^{(s)} | x_{t-1}^{(s)}; u^{(s)} \sim N(x_t; \begin{pmatrix} m_{t-1}^{(s)} \\ m_Z^{(s)} \end{pmatrix}; \begin{pmatrix} B_{t-1}^{(s)} + Q & K_{t-1,Z}^{(s)} \\ K_{t-1,Z}^{(s)} & K_Z^{(s)} \end{pmatrix}) \quad (23)$$

For a general likelihood model, we can sample the (noisy) targets using $y_t^{(s)} | x_t^{(s)} \sim p(y_t | x_t^{(s)}; \theta)$. In the case of a Gaussian conditional likelihood, as described in § 4.1, we can see that given samples from the latent state, the predictive distribution is a Gaussian

$$y_t^{(s)} | x_t^{(s)}; u^{(s)} \sim N(y_t; C_{t-1}^{(s)} x_t^{(s)} + d; C_{t-1}^{(s)} C^> + R); \quad (24)$$

where the samples of $x_t^{(s)}; u^{(s)}$ are readily available after running SGHMC on their joint space.

5. Collapsing Inducing Variables

So far we have described our method to obtain samples from the optimal variational posterior over the joint distribution of state trajectories and inducing variables using Eq. (21) and SGHMC. In this section we show that we can, in fact, integrate out the inducing variables from our joint variational distribution and obtain the optimal marginal distribution for latent states $x_{0:T}$. We start by retaking Eq. (21) and isolating the terms that depend on the inducing variables,

$$\begin{aligned} q(x_{0:T}) &= \int_u q(u; x_{0:T}) du \\ &= p(x_0) \prod_{t=1}^T p(y_t | x_t) \exp\left(-\frac{1}{2} \text{Tr}(Q^{-1} B_{t-1})\right) \\ &\quad \int_u p(u | Z) \prod_{t=1}^T N(x_t; x_{t-1}; Q) du; \quad (25) \end{aligned}$$

where we note that x_t , as defined in Eq. (18), depends on the inducing variables u . With this, we can complete the square and identify the terms in the integral as products of Gaussian distributions, whose normalization constant is Gaussian. Therefore, we have

$$q(x_{0:T}) = p(x_0) \prod_{t=1}^T p(y_t | x_t) \exp\left(-\frac{1}{2} \text{Tr}(Q^{-1} B_{t-1})\right) N(x_t; m_{t-1}; Q) = N(x; 0; \tilde{\Sigma}_x); \quad (26)$$

where the Gaussian density in the denominator is determined by

$$x = \sum_{t=1}^T A_{t-1}^> Q^{-1} (x_t - m_{t-1}); \quad (27)$$

$$\tilde{\Sigma}_x = I + \sum_{t=1}^T A_{t-1}^> Q^{-1} A_{t-1}; \quad (28)$$

$$A_{t-1} = K_{z,t-1} (L_Z^>)^{-1}; \quad (29)$$

where L_Z is the Cholesky decomposition of K_Z , i.e., $K_Z = L_Z L_Z^>$. We note here that it is actually a projection of the cumulative uncorrelated input $Q^{-1}(x_t - m_{t-1})$ via the projection matrix $L_Z^{-1} K_{z,t-1}$ and, therefore, $x \in \mathbb{R}^M$.

Furthermore, in this collapsed version, the terms unrelated to u are kept the same as those in the original optimal joint distribution in Eq. (21). The individual Gaussian transitions for the latent states are now $N(x_t; m_{t-1}; Q)$ since u has been integrated out and we recall that $m_{t-1} = m_t(x_{t-1})$. The additional term, which is the inverse of a multivariate Gaussian distribution, records the cumulative projected trajectory with the corresponding projected variances. Maximizing $q(x_{0:T})$ would tend to minimize this density function, which pushes the latent states away from 0 and also decreases the values of Q -regulated variance.

In the prediction scenario, we run SGHMC to obtain samples $x_{0:T}^{(s)}$ from the marginal $q(x_{0:T})$ and then use the closed-form expression for the conditional (Section G in the Supplementary Material) to obtain $u^{(s)} | x_{0:T}^{(s)}$.

5.1. Advantages of Collapsed Algorithm

Collapsing the inducing variables will generally tend to improve the convergence of our algorithm, as we are required to sample from a significantly lower number of latent variables. The computational cost is similar to that of the uncollapsed algorithm since despite avoiding the computation of K_Z^{-1} in $p(u | z)$, we require a similar term L_Z^{-1} , in the evaluation of the projected Gaussian distributions.

It is important to emphasize one particular difference in our approach with respect to the closely related VCDT algorithm of [Ialongo et al. \(2019\)](#). As described before, [Ialongo et al.](#)

[\(2019\)](#) also propose a coupled joint posterior between state trajectories and inducing variables. Their factorization is $q_{\text{VCDT}} = q(u)q(x_{0:T} | u)$, and they impose additional Gaussian constraints on these densities. Our implicit assumed factorization is $q_{\text{FFVD}} = q(x_{0:T})q(u | x_{0:T})$, which allows us to obtain the optimal variational distribution without imposing any additional parametric constraints, integrate out the inducing variables analytically and get a Monte Carlo approximation to the optimal marginal $q(x_{0:T})$ via samples. We also provide an expression for the conditional $q(u | x_{0:T})$ in closed-form. Details can be found in [Appendix G](#).

5.2. Particle Markov chain Monte Carlo (PMCMC)

It is clear that the latent states are constructed in a Markovian manner when the transition function $m_t(\cdot)$ is given, as the value of the current latent state is dependent on the previous latent state's value. Therefore, we can use PMCMC methods ([Andrieu et al., 2010](#)) to infer the posterior distribution of the Markov structured latent states $s_{0:T}$. This Bayesian treatment might improve performance over SGHMC, as it incorporates the sequential nature of the problem into the sampling algorithm. The advantages of PMCMC over standard sequential Monte Carlo approaches have been documented previously, see, e.g., [Andrieu et al. \(2010\)](#). Here we note that [Frigola et al. \(2013\)](#) also proposed PMCMC treatment for $x_{0:T}$. However, their algorithm is very different from ours as it is based on the fully-independent conditional approximation (see, e.g., [Quinero-Candela & Rasmussen, 2005](#)). More details can be found in [Appendix K](#).

6. Related Work

We have already described the main differences between our method and closely-related approaches throughout the paper, e.g., in [§§ 1, 4 and 5](#). We refer the reader to [Appendix K](#) for more details. Other works have considered the GPSSM in partially observable unstable settings ([Curi et al., 2020](#)) or combined variational inference with the Laplace approximation ([Lindinger et al., 2022](#)). With regards to scalable GPs, we note they have been the subject of much research effort in machine learning, with extensions to more general frameworks such as compositional models (see, e.g., [Wilson & Nickisch, 2015](#); [Salimbeni & Deisenroth, 2017](#); [Yu et al., 2019](#); [Cutajar et al., 2017](#); [Havasi et al., 2018](#); [Rossi et al., 2021](#)). Our approach can be seen as a generalization of the fully-Bayesian supervised learning method proposed by [Rossi et al. \(2021\)](#) to state-space models, where we have included the non-trivial component of transition dynamics and have proposed a collapsed optimal variational distribution for state trajectories.

Other models for GPs on sequential data have been proposed, see, for example, [Frigola \(2015\)](#) for an excellent overview. Of interest here is the state-space model view

Figure 1. Results on synthetic data: left: Observations shown as green dots, Ground truth as a solid black line, FFVD's mean prediction as a dashed red line with one standard deviation error bars. Right: Histograms of p-values for the hypothesis test that each marginal posterior over states $x_{t|g}$ (top) and inducing variables $u_{i|g}$ (bottom) is generated from a Normal distribution.

Figure 2. Training ($t \leq 150$) and test performance ($t = 150$) on the Furnace dataset. The black solid line is the underlying ground truth signal and the dashed red line is FFVD's mean prediction. The blue solid line indicates the training/test split. An underlying $x = 4$ -dimensional latent state was used.

of GPS that for time-series data with $d_x = 1$ and Markovian covariance functions can provide exact inference in linear time $O(T)$ (Solin, 2016). This has been extended to non-Gaussian likelihood models and made more efficient using several computational primitives (Nickisch et al., 2018). More recently, there has been some work on using the signature kernel (Toth & Oberhauser, 2020; Salvi et al., 2021) within GP models. In particular, Lemercier et al. (2021) generalize variational orthogonal features (Burt et al., 2020; Hensman et al., 2017) to the sequential case, constructing inducing variables associated with the signature kernel that yield a variational inference algorithm that does not require any matrix inversion.

7. Experiments

We evaluate our FFVD method on synthetic data and on six real-world system identification benchmarks (Ialongo et al., 2019; Doerr et al., 2018), comparing it with GPSSM (Frigola et al., 2014), PRSSM (Doerr et al., 2018), VCDT (Ialongo et al., 2019), and use a LSTM network (Hochreiter & Schmidhuber, 1997) as a baseline non-GP-based model. All experiment details can be found in Appendix I.

Figure 3. Traceplot of the training log-likelihood when using SGHMC (FFVD-M) and when collapsing the inducing variables FFVD-C-M. Collapsing generally improved convergence.

7.1. Synthetic Data

We generate data from a GPSSM with a squared exponential covariance function. Our goal here is to investigate the properties of our FFVD algorithm. Therefore, we fix the value of all parameters to their ground-truth values except for the latent states $x_{0:T}$ and the inducing variables u . The left panel of Fig. 1 illustrates that FFVD effectively learns the intricate transition function with two modes within regions containing latent states. As the number of latent states increases, FFVD more accurately approximates the true function. A lack of fit in regions without latent states is to be expected. Having a good fit, we are now interested in knowing whether the true posterior (as estimated by our algorithm) is close to a Gaussian distribution. The right panel in Fig. 1 indicates that more than 40% of the latent states $x_{0:T}$ and more than 50% of the inducing variables u do not provide enough support for the hypothesis that their marginal distributions follow a Normal distribution (see details in Appendix I). This brings into question the strong parametric assumptions over the variational posterior made by previous work (e.g., Ialongo et al., 2019; Doerr et al., 2018).

7.2. Real-World Data

Here we evaluate the different methods using six system identification benchmarks with a latent state dimension $d_x = 4$, as used by Ialongo et al. (2019); Doerr et al. (2018). See more details in Appendix I. The prediction performance is shown in Table 2 and Table 3, where we see that FFVD attains the lowest RMSE values in four out of the six benchmarks (in bold). Furthermore, for the Actuator and Dryer, FFVD-C-M ranks second among all the algorithms (underlined). The performance of GPSSM and PRSSM is usually worse than others, which is likely due to their strong mean-field and parametric assumptions, respectively. LSTM obtains good performance on three datasets, although its deterministic neural network structure is different from our random function setting. The performance of VCDT is

Table 2. Mean RMSE values (one standard deviations) on the real-world system identification benchmarks. Our method, when using SGHMC (FFVD-M); the collapsed version (FFVD-C-M); and when using pMCMC (FFVD-P).

Methods	Actuator		Ballbeam		Drive		Dryer		Flutter		Furnace	
LSTM	0:586	0:411	0:027	0:023	0:537	0:108	0:115	0:029	0:912	0:562	1:261	0:610
VGPSSM	0:580	0:274	0:073	0:011	0:722	0:087	0:241	0:023	1:482	0:218	1:115	0:358
PRSSM	0:497	0:381	0:059	0:013	0:813	0:101	0:017	0:042	1:371	0:156	1:243	0:407
VCDT	0:239	0:040	0:011	0:002	0:585	0:017	0:142	0:003	1:782	0:324	1:166	0:011
FFVD-M	0:358	0:242	0:019	0:018	0:673	0:207	0:205	0:313	0:280	0:193	0:571	0:185
FFVD-C-M	0:259	0:209	0:009	0:011	0:775	1:615	0:065	0:112	0:663	0:189	0:548	0:051
FFVD-P	0:388	0:087	0:199	0:045	0:342	0:057	0:317	0:050	0:562	0:088	0:669	0:174

Table 3. Negative Log-Likelihood values on the real-world system identification benchmarks.

Methods	Actuator		Ballbeam		Drive		Dryer		Flutter		Furnace	
VCDT	0:36	0:02	0:65	0:01	1:23	0:01	0:02	0:01	6:13	0:48	7:49	0:07
FFVD-M	0:03	0:10	0:09	0:05	1:66	0:01	0:08	0:09	0:48	0:33	0:43	0:05

Table 4. Mean running time values on Furnace

Methods	Iterations	Furnace
VCDT	10	2696
	100	26971
	500	161325
	1000	337478
FFVD-M	10	319
	100	2633
	500	12142
	1000	23868

the closest to our FFVD methods, as it models a coupled

In addition to this quantitative evaluation, we can see a qualitative illustration of using our algorithm for predicting the training and test (future) observations in Fig. 2. This is an example of good generalization, although it is (of course) not consistent across all problems, given the limited training data. Finally, we analyze the convergence of our algorithm in Fig. 3, where we see that FFVD-C-M uses less iterations (8000 iterations) than FFVD-M (40000 iterations) to achieve similar performance. This confirms the benefits of collapsing the inducing variables and sampling only on the lower-dimensional space of state trajectories. We can see these analyses for all benchmarks in Appendices J.1 and J.2.

We also have executed experiments to illustrate the advantages of our approach when considering running time regarding to different number of iterations in Table 4. We can clearly see that our approach is 10 times faster than VCDT. These time results were done using a Macbook Pro 2021

with 16GB in memory, M1 chip, and 8 cores. It is noted that caution must be taken with interpreting these results, as they depend on implementation specifics, computer architectures along with other practical details.

8. Limitations

Our approach is an approximation. Indeed, we did not realize our paper came across as claiming our posterior was the true posterior or exact in any particular limiting case. Such an approach for defining the conditional posterior in terms of the conditional prior is customary in variational methods for GP models since the seminal work of Titsias (2009). Sampling from the true full (non-sparse) posterior is not scalable due to its cubic cost on the number of observations.

9. Conclusion

We have presented FFVD, a new algorithm for inference in GPSSMs based on variational inference under the inducing-variable formalism. Unlike previous approaches, FFVD does not make any independence or parametric assumptions on the joint variational posterior over state trajectories and inducing variables $(x_{0:T}; u)$ and, instead, represents the posterior via samples from the optimal variational distribution. Furthermore, we have shown that our formulation allows us to collapse the inducing variables and carry out inference in the lower dimensional space of state trajectories. Given samples from trajectories, we can make predictions using the optimal conditional $(u_j | x_{0:T})$, for which we have derived a closed-form expression. Future work will investigate scalable multi-output GPs in order to deal with very high-dimensional state representations in a flexible and efficient way.

References

- Andrieu, C., Doucet, A., and Holenstein, R. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3): 269–342, 2010.
- Bauer, M., van der Wilk, M., and Rasmussen, C. E. Understanding probabilistic sparse gaussian process approximations. *Advances in neural information processing systems* 29, 2016.
- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., and Wilson, M. A. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience* 18(18):7411–7425, 1998.
- Burt, D. R., Rasmussen, C. E., and van der Wilk, M. Variational orthogonal features. *arXiv preprint arXiv:2006.13170* 2020.
- Chen, T., Fox, E., and Guestrin, C. Stochastic gradient Hamiltonian Monte Carlo. *ICML*, pp. 1683–1691. PMLR, 2014.
- Curi, S., Melchior, S., Berkenkamp, F., and Krause, A. Structured variational inference in partially observable unstable Gaussian process state space models. In *Bayen, A. M., Jadbabaie, A., Pappas, G., Parrilo, P. A., Recht, B., Tomlin, C., and Zeilinger, M. (eds.), Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pp. 147–157. PMLR, 10–11 Jun 2020.
- Cutajar, K., Bonilla, E. V., Michiardi, P., and Filippone, M. Random feature expansions for deep Gaussian processes. In *Precup, D. and Teh, Y. W. (eds.), Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 884–893. PMLR, 06–11 Aug 2017.
- D’agostino, R. and Pearson, E. S. Tests for departure from normality. *Biometrika* 60(3):613–622, 1973.
- Doerr, A., Daniel, C., Schiegg, M., Duy, N.-T., Schaal, S., Toussaint, M., and Sebastian, T. Probabilistic recurrent state-space models. *ICML*, pp. 1280–1289, 2018.
- Eleftheriadis, S., Nicholson, T., Deisenroth, M., and Hensman, J. Identification of Gaussian process state space models. *Advances in neural information processing systems* 30, 2017.
- Frigola, R. Bayesian time series learning with Gaussian processes. PhD thesis, University of Cambridge, 2015.
- Frigola, R., Lindsten, F., Schön, T. B., and Rasmussen, C. E. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. *NIPS*, pp. 3156–3164, 2013.
- Frigola, R., Chen, Y., and Rasmussen, C. E. Variational Gaussian process state-space models. *NIPS*, pp. 3680–3688, 2014.
- Havasi, M., Hernández-Lobato, J. M., and Murillo-Fuentes, J. J. Inference in deep Gaussian processes using stochastic gradient Hamiltonian Monte Carlo. *NeurIPS* 31, 2018.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. *Uncertainty in Artificial Intelligence* AUAI Press, 2013.
- Hensman, J., Matthews, A. G., Filippone, M., and Ghahramani, Z. MCMC for variationally sparse Gaussian processes. *NIPS*, pp. 1648–1656, 2015.
- Hensman, J., Durrande, N., Solin, A., et al. Variational Fourier features for Gaussian processes. *Mach. Learn. Res.* 18(1):5537–5588, 2017.
- Hernandez, E., Martin, F., and Valero, F. State-space modeling for atmospheric pollution. *Journal of Applied Meteorology and Climatology* 30(6):793–811, 1991.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation* 9(8):1735–1780, 1997.
- Ialongo, A. D., Van Der Wilk, M., Hensman, J., and Rasmussen, C. E. Overcoming mean-eld approximations in recurrent Gaussian process models. *ICML*, pp. 2931–2940, 2019.
- Kalman, R. E. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* 82(1):35–45, 1960. doi: 10.1115/1.3662552.
- Katzfuss, M., Stroud, J. R., and Wikle, C. K. Understanding the Ensemble Kalman Filter. *The American Statistician*, 70(4):350–357, 2016. doi: 10.1080/00031305.2016.1141709.
- Kitagawa, G. Non-Gaussian state-space modeling of non-stationary time series. *Journal of the American statistical association* 82(400):1032–1041, 1987.
- Lemercier, M., Salvi, C., Cass, T., Bonilla, E. V., Damoulas, T., and Lyons, T. J. SigGPDE: Scaling sparse Gaussian processes on sequential data. *International Conference on Machine Learning*, pp. 6233–6242. PMLR, 2021.
- Lindinger, J., Rakitsch, B., and Lippert, C. Laplace approximated Gaussian process state-space models. *38th Conference on Uncertainty in Artificial Intelligence* 2022.

- Marmin, S. and Filippone, M. Deep gaussian processes for calibration of computer models (with discussion). *Bayesian Analysis* 7(4):1301–1350, 2022.
- Murphy, K. P. Probabilistic Machine Learning: Advanced Topics. MIT Press, 2023. URL <http://probml.github.io/book2>.
- Nickisch, H., Solin, A., and Grigorevskiy, A. State space Gaussian processes with non-Gaussian likelihood. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3789–3798. PMLR, 10–15 Jul 2018.
- Ogata, K. et al. *Modern control engineering*, volume 5. Prentice hall Upper Saddle River, NJ, 2010.
- Quinero-Candela, J. and Rasmussen, C. E. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research* 6:1939–1959, 2005.
- Rossi, S., Heinonen, M., Bonilla, E., Shen, Z., and Filippone, M. Sparse Gaussian processes revisited: Bayesian approaches to inducing-variable approximations. *AI&STAT* pp. 1837–1845, 2021.
- Salimbeni, H. and Deisenroth, M. Doubly stochastic variational inference for deep Gaussian processes. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Salvi, C., Cass, T., Foster, J., Lyons, T., and Yang, W. The signature kernel is the solution of a goursat PDE. *SIAM Journal on Mathematics of Data Science* 3(3):873–899, 2021.
- Sauer, A., Cooper, A., and Gramacy, R. B. Vecchia-approximated deep Gaussian processes for computer experiments. *Journal of Computational and Graphical Statistics* 0(0):1–14, 2022.
- Solin, A. Stochastic differential equation methods for spatio-temporal Gaussian process regression. PhD thesis, Aalto University, 2016.
- Titsias, M. Variational learning of inducing variables in sparse Gaussian processes. *AI&STAT* pp. 567–574, 2009.
- Toth, C. and Oberhauser, H. Bayesian learning from sequential data using Gaussian processes with signature covariances. In *International Conference on Machine Learning* pp. 9548–9560. PMLR, 2020.
- Tsay, R. S. *Analysis of financial time series*. John Wiley & Sons, 2005.
- Williams, C. K. and Rasmussen, C. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.
- Wilson, A. and Nickisch, H. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1775–1784, Lille, France, 07–09 Jul 2015. PMLR.
- Yu, H., Chen, Y., Low, B. K. H., Jaillet, P., and Dai, Z. Implicit Posterior Variational Inference for Deep Gaussian Processes. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, pp. 14475–14486. Curran Associates, Inc., 2019.

A. Basic Results on Gaussian Distributions

A.1. Conditional Gaussian Distribution

Assuming the joint Gaussian distribution $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \Sigma)$ over the random vector \mathbf{x} such that:

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_a \\ \mathbf{x}_b \end{pmatrix}; \quad \boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_a \\ \boldsymbol{\mu}_b \end{pmatrix}; \quad \Sigma = \begin{pmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{pmatrix}; \quad (30)$$

where $\Sigma_{ba} = \Sigma_{ab}^T$, we have that the conditional distributions are given by:

$$p(\mathbf{x}_b | \mathbf{x}_a) = \mathcal{N}(\mathbf{x}_b; \boldsymbol{\mu}_b + \Sigma_{ba} \Sigma_{aa}^{-1}(\mathbf{x}_a - \boldsymbol{\mu}_a); \Sigma_{bb} - \Sigma_{ba} \Sigma_{aa}^{-1} \Sigma_{ab}); \quad (31)$$

A.2. Expectation over log of Normal Distribution

With an approximate marginal posterior $q(\mathbf{f}_?)$ and a Normal distribution $p(\mathbf{y} | \mathbf{f}_?)$ of the form:

$$q(\mathbf{f}_?) = \mathcal{N}(\mathbf{f}_?; \boldsymbol{\mu}_?, \Sigma_?); \quad (32)$$

$$p(\mathbf{y} | \mathbf{f}_?) = \mathcal{N}(\mathbf{y}; \mathbf{f}_?, \Sigma_y); \quad (33)$$

we can compute

$$E_{q(\mathbf{f}_?)} \log p(\mathbf{y} | \mathbf{f}_?) = \log \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_?, \Sigma_y) + \text{Tr}(\Sigma_y^{-1} \boldsymbol{\Sigma}_?): \quad (34)$$

B. Posterior Marginal in Sparse GP Models

In variational sparse GP models we usually have the joint Gaussian $p(\mathbf{f}_?, \mathbf{u})$ and a posterior distribution over \mathbf{u} with

$$p(\mathbf{f}_?) = \mathcal{N}(\mathbf{m}_?; \mathbf{K}_?); \quad (35)$$

$$p(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{m}_Z; \mathbf{K}_Z); \quad (36)$$

$$\text{Cov}(\mathbf{f}_?; \mathbf{u}) = \mathbf{K}_{?:Z}; \quad (37)$$

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_u; \Sigma_u); \quad (38)$$

$$q(\mathbf{f}_?; \mathbf{u}) := q(\mathbf{u})p(\mathbf{f}_? | \mathbf{u}); \quad (39)$$

and we wish to compute:

$$q(\mathbf{f}_?) = \int_Z q(\mathbf{f}_?; \mathbf{u}) d\mathbf{u}; \quad (40)$$

where we have omitted conditional dependencies $\mathbf{x}_?$ and Z for simplicity in the notation. For example this is necessary to compute the expectations over the conditional log likelihood term during inference or to make predictions on a new test point. Using Eq. (31) we obtain:

$$p(\mathbf{f}_? | \mathbf{u}) = \mathcal{N}(\mathbf{f}_?; \mathbf{m}_? + \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1}(\mathbf{u} - \mathbf{m}_Z); \mathbf{K}_? - \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1} \mathbf{K}_{Z:}); \quad (41)$$

In order to integrate out over \mathbf{u} we know that the result is a Gaussian with mean and covariances obtained from the following linear transformation of:

$$\mathbf{f}_? = \mathbf{m}_? + \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1}(\mathbf{u} - \mathbf{m}_Z) + \boldsymbol{\epsilon}; \quad \mathbf{N}(0; \mathbf{K}_? - \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1} \mathbf{K}_{Z:}); \quad (42)$$

Thus, we obtain:

$$q(\mathbf{f}_?) = \mathcal{N}(\mathbf{f}_?; \mathbf{m}_? + \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1}(\boldsymbol{\mu}_u - \mathbf{m}_Z); \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1} \Sigma_u \mathbf{K}_{Z:} + \mathbf{K}_? - \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1} \mathbf{K}_{Z:}); \quad (43)$$

$$= \mathcal{N}(\mathbf{f}_?; \mathbf{m}_? + \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1}(\boldsymbol{\mu}_u - \mathbf{m}_Z); \mathbf{K}_? + \mathbf{K}_{?:Z} \mathbf{K}_Z^{-1}(\Sigma_u - \mathbf{K}_Z) \mathbf{K}_Z^{-1} \mathbf{K}_{Z:}); \quad (44)$$

C. Graphical Model

Fig. 4 illustrates GPSSMs graphical model.

Figure 4. Generative process of the GPSSM (left panel) and the sparse GPSSM (right panel).

D. Details of Prior Reparameterization: Whitening Prior over Inducing Variables

Here we show a more detailed derivation of the prior reparameterization and, consequently, the new form of the required conditional distribution. We know that our prior over the inducing variables is $p(u) = \mathcal{N}(u; m_Z; K_Z)$ and that we use the new whitened prior $u = L_Z^{-1}(u - m_Z)$ with $L_Z L_Z^T = K_Z$ and $p(v) = \mathcal{N}(v; 0; I_M)$. We also have that the marginal prior over $f_?$ is $p(f_?) = \mathcal{N}(f_?; m_?; K_?)$. It follows that $p(f; v)$ is a Gaussian with cross-covariance:

$$\text{Cov}(f_?; v) = E[(f_? - m_?)(L_Z^{-1}(u - m_Z))^T] \quad (45)$$

$$= E\left[\underbrace{(f_? - m_?)(u - m_Z)^T}_{\text{Cov}(f_?; u)}\right](L_Z^{-1})^T \quad (46)$$

$$= K_{?:Z} (L_Z^{-1})^T \quad (47)$$

Hence, using Eq. (31), we have that:

$$p(f_? j; x_?; v; Z) = \mathcal{N}(f_?; m_? + K_{?:Z} (L_Z^{-1})^T v; K_? - K_{?:Z} (L_Z^{-1})^T L_Z^{-1} K_{Z,t}); \quad (48)$$

$$p(f_? j; x_?; v; Z) = \mathcal{N}(f_?; m_? + K_{?:Z} (L_Z^{-1})^T v; K_? - K_{?:Z} K_Z^{-1} K_{Z,?}); \quad (49)$$

Alternatively, we can simply obtain this result by replacing $u = m_Z + L_Z v$ in the conditional distribution. Here we note that this is consistent with the GP definition in that if we were to integrate out the variables u from the joint model $p(f_?; v)$ we would obtain the exact GP prior $p(f_?)$.

D.1. Posterior Marginal

With the whitened joint model $p(f_?; v)$, Gaussian approximate posterior $q(v)$ and a variationally sparse GP model we have:

$$p(f_? j v) = \mathcal{N}(f_?; m_? + K_{?:Z} (L_Z^{-1})^T v; K_? - K_{?:Z} K_Z^{-1} K_{Z,?}); \quad (50)$$

$$q(v) = \mathcal{N}(v; \hat{v}; \hat{\Sigma}); \quad (51)$$

$$q(f_?; v) := q(v)p(f_? j v); \quad (52)$$

It is easy to show that the posterior marginal $q(f_?) = \int_{\mathbb{R}^M} q(f_?; v) dv$ is:

$$q(f_?) = \mathcal{N}(f_?; m_? + K_{?:Z} (L_Z^{-1})^T \hat{v}; K_? + K_{?:Z} (L_Z^{-1})^T (\hat{\Sigma} - I_M) L_Z^{-1} K_{Z,?}); \quad (53)$$

E. Free-Form Posterior Estimation By Bounding the Log Marginal Likelihood Directly

Computing the log marginal:

$$\log p(y_{0:T-1}; x_{0:T-1}; u j Z) = \log p(u j Z) p(x_0) p(y_0 j x_0) \prod_{t=1}^{T-1} p(y_t j x_t) + L(x_{0:T-1}; u; Z); \quad (54)$$

where

$$L(x_{0:T-1}; u; Z) := \log \int_{f_{1:T-1}} p(x_t | f_t) p(f_{1:T-1}; x_{0:T-1}; u; Z) df_{1:T-1}; \quad (55)$$

which we note is a log of an expectation, which we can bound using Jensen's inequality:

$$L(x_{0:T-1}; u; Z) \leq \mathbb{E}_{p(f_{1:T-1}; x_{0:T-1}; u; Z)} \log \int_{f_t} p(x_t | f_t) \quad (56)$$

$$= \mathbb{E}_{p(f_t | x_{t-1}; u; Z)} \log p(x_t | f_t); \quad (57)$$

where $p(f_t | x_{t-1}; u; Z)$ is the GP predictive distribution over f_t at x_{t-1} given the (pseudo observations) inducing variables u at inducing input Z :

$$p(f_t | x_{t-1}; u; Z) = N(f_t; A_{t-1}u; B_{t-1}) \quad \text{with} \quad (58)$$

$$A_{t-1} = K_{t-1;Z} K_Z^{-1} \quad (59)$$

$$B_{t-1} = K_{t-1} - K_{t-1;Z} K_Z^{-1} K_{Z;t-1}; \quad (60)$$

Hence, the expectations in Eq. (57) can be computed in closed-form:

$$\mathbb{E}_{p(f_t | x_{t-1}; u; Z)} \log p(x_t | f_t) = \log N(x_t; A_{t-1}u; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}); \quad (61)$$

Thus, we have that:

$$\log p(y_{0:T-1}; x_{0:T-1}; u; Z) = \log [p(u; Z) p(x_0) p(y_0 | x_0)] + \sum_{t=1}^{T-1} \log p(y_t | x_t) + \log N(x_t; A_{t-1}u; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}); \quad (62)$$

Then setting the latent variables $z := f_{0:T-1}$, we can obtain the approximate log unnormalized posterior:

$$\log q(z) := \log p(u; Z) + \log p(x_0) + \log p(y_0 | x_0) + \sum_{t=1}^{T-1} \log p(y_t | x_t) + \log N(x_t; A_{t-1}u; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}); \quad (63)$$

where $q(z) / q(z) = p(z | y_{0:T-1})$. Thus, we can draw samples from our approximate posterior using MCMC and the energy function $U(z) = -\log p(z; y_{0:T-1}) = -\log p(z | y_{0:T-1}) + C = -\log q(z) + C$.

E.1. Whitened Version

In the whitened version, $v = L_Z^{-1}(u - m_Z)$ or $u = m_Z + L_Z v$ and the effective prior is

$$p(v) = N(v; 0; I); \quad (64)$$

Letting $\mu = (u; Z)$; $\nu = (v; Z)$, the determinant of the Jacobian can be calculated as:

$$\begin{aligned} \left| \frac{\partial \mu}{\partial \nu} \right| &= \begin{vmatrix} \frac{\partial (m_Z + L_Z v)}{\partial v} & \frac{\partial (m_Z + L_Z v)}{\partial Z} & \frac{\partial (m_Z + L_Z v)}{\partial Z} \\ \frac{\partial v}{\partial v} & \frac{\partial v}{\partial Z} & \frac{\partial v}{\partial Z} \\ \frac{\partial Z}{\partial v} & \frac{\partial Z}{\partial Z} & \frac{\partial Z}{\partial Z} \end{vmatrix} \\ &= \begin{vmatrix} L_Z & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{vmatrix} = \begin{vmatrix} L_Z & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{vmatrix} = |L_Z| = |K_Z|^{1/2} \end{vmatrix} \quad (65) \end{aligned}$$

Thus, we can re-write the approximate log unnormalized posterior as:

$$\begin{aligned}
 \log q(v; x_{0:T-1}) &= \log jK_Z j^{\frac{1}{2}} - \log jK_Z j^{\frac{1}{2}} - \frac{1}{2} v^T v + \log p(x_0) + \log p(y_0 | x_0) + \\
 &\quad \prod_{t=1}^{T-1} \log p(y_t | x_t) + \log N(x_t; K_{t-1;Z} (L_Z^{-1})^T v; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}) : \\
 &= \log p(v) + \log p(x_0) + \log p(y_0 | x_0) + \\
 &\quad \prod_{t=1}^{T-1} \log p(y_t | x_t) + \log N(x_t; K_{t-1;Z} (L_Z^{-1})^T v; Q) - \frac{1}{2} \text{Tr}(Q^{-1} B_{t-1}) : \quad (66)
 \end{aligned}$$

F. Collapsed Method

We can integrate out u in Eq. (21) to obtain the marginal distribution $q(x_{1:T})$.

$$q(x_{1:T}) = \int_u q(u; x_{1:T}) du = \int_u p(u | Z) p(x_0) \prod_{t=1}^{T-1} p(y_t | x_t) \exp\left(-\frac{1}{2} \text{Tr}(Q^{-1} B_{t-1})\right) N(x_t; x_t; Q) du \quad (67)$$

$$= p(x_0) \prod_{t=1}^{T-1} p(y_t | x_t) \exp\left(-\frac{1}{2} \text{Tr}(Q^{-1} B_{t-1})\right) \int_u p(u) \prod_{t=1}^{T-1} N(x_t; x_t; Q) du : \quad (68)$$

In the case of using our prior re-parameterization, i.e., whitened inducing variables, for multidimensional states and, therefore, multi-output GPs, we can write the integral above for dimension d as:

$$\begin{aligned}
 \int_{v^{(d)}} p(v^{(d)}) \prod_{t=1}^{T-1} N(x_t^{(d)}; m_{t-1}^{(d)} + A_{t-1}^{(d)} v^{(d)}; Q_d) dv^{(d)} \\
 = \prod_{t=1}^{T-1} N(x_t^{(d)}; m_{t-1}^{(d)}; Q_d) \exp\left(-\frac{1}{2} (g^{(d)})^T (H^{(d)})^{-1} g^{(d)}\right) \det(H^{(d)})^{-\frac{1}{2}} \quad (69)
 \end{aligned}$$

where

$$g^{(d)} := I + \prod_{t=1}^{T-1} (A_{t-1}^{(d)})^T Q_d^{-1} (A_{t-1}^{(d)}) \quad x_t^{(d)} - m_{t-1}^{(d)} > Q_d^{-1} A_{t-1}^{(d)} ; \quad (70)$$

$$H^{(d)} = I + \prod_{t=1}^{T-1} (A_{t-1}^{(d)})^T Q_d^{-1} (A_{t-1}^{(d)}) \quad (71)$$

$$K_Z^{(d)} = L_Z^{(d)} (L_Z^{(d)})^T : \quad (72)$$

Thus, our optimal log variational posterior marginal over state trajectories is

$$\begin{aligned}
 \log q(x_{0:T}) &= \log p(x_0) + \prod_{t=1}^{T-1} [\log p(y_t | x_t) + \log N(x_t; m_{t-1}; Q)] \\
 &\quad - \frac{1}{2} \sum_{d=1}^D \prod_{t=1}^{T-1} \text{Tr}(Q_d^{-1} B_{t-1}^{(d)}) + \log \det(H^{(d)})^{-1} (x^{(d)})^T H^{(d)} (x^{(d)}) ; \quad (73)
 \end{aligned}$$

where $x^{(d)} := \prod_{t=1}^{T-1} (A_{t-1}^{(d)})^T Q_d^{-1} (x_t^{(d)} - m_{t-1}^{(d)})$:

Algorithm 1 Particle MCMC for inferring latent states $x_{0:T}$

Require: Number of particles S , observations $y_{1:T}$, samples of $x_{1:T}^0$ in the previous iteration, likelihood parameters $\theta = f; C; d; R; g$,

Initialize state x_0 , weights $W_t^{(i)} = 1$ for $i = 1; \dots; S$.

Fix the last particle as setting $x_{1:T}^{(S)} = x_{1:T}^0$

for $t = 1; \dots; T$ do

Generate $x_t^{(i)}$ from Eq. (22) for $i = 1; \dots; S-1$.

Calculate weight $W_t^i = p(y_t | x_t^{(i)})$ for $i = 1; \dots; S$

Normalize the weights $\bar{W}_t^{(i)} = W_t^{(i)} / \sum_i W_t^{(i)}$

if $t < T$ then

For $i^0 = 1; \dots; S-1$, re-sample the index j_{i^0} from the categorical distribution, with the event probabilities being $(\bar{W}_t^{(1)}; \dots; \bar{W}_t^{(S)})$.

For $i^0 = 1; \dots; S-1$, let $x_{1:t}^{(i^0)} := x_{1:t}^{(j_{i^0})}$ and set $W_t^{i^0} = 1/S$

else

Re-sample the index j from the categorical distribution, with the event probabilities being $(\bar{W}_t^{(1)}; \dots; \bar{W}_t^{(S)})$.

Return $x_{0:T} := x_{1:T}^{(j)}$

end if

end for

G. Closed-Form Optimal Variational Conditional Distribution over Inducing Variables

Given a trajectory, the optimal variational conditional over the whitened inducing variables given a state trajectory is given by:

$$q(v | x_{0:T}) = \prod_{d=1}^D N(v^{(d)}; g^{(d)}; H^{(d)}); \quad (74)$$

where g and H are given in Eqs. (70) and (71), respectively. A similar expression can be obtained for the original (non-whitened) inducing variables. Of importance here is that our formulation has allowed us to obtain an optimal conditional variational distribution, given the state trajectories, over the inducing variables in closed form. We see then that, in effect, our approximate posterior $q_{\text{FFVD}}(x_{0:T}; u) = q(x_{0:T})q(u | x_{0:T})$. In order to make predictions, we run SGHMC to obtain samples $x_{0:T}^{(s)}$ from the marginal $q(x_{0:T})$ and then use the closed-form expression for the conditional to obtain $f(u^{(s)} | x_{0:T}^{(s)})$. Running SGHMC over the much lower-dimensional space of state trajectories (instead of the joint space of trajectories and inducing variables) should generally converge faster. Our experiments with both versions of the algorithm confirm this.

G.1. Optimal Closed-Form Conditional vs Assumed Parametric Factorizations

We contrast our variational distribution with that of the variational Gaussian process state space model (VGPSSM; Frigola et al., 2014) and the variationally coupled dynamic trajectory model (VCDT; Lalongo et al., 2019). While, implicitly, our variational distribution is $q_{\text{FFVD}}(x_{0:T}; u) := q(x_{0:T})q(u | x_{0:T})$ the VGPSSM and VCDT assume $q_{\text{VGPSSM}}(x_{0:T}; u) := q(x_{0:T})q(u)$ and $q_{\text{VCDT}}(x_{0:T}; u) := q(x_{0:T} | u)q(u)$, respectively. Although the former (VGPSSM) obtained an “optimal” posterior (given the factorization assumption), it is a mean-field approximation and ignores the dependencies between state trajectories and inducing variables. The latter (VCDT), does not make a mean-field assumption but its factorization forces parametric constraints over the individual distributions. In fact, Lalongo et al. (2019) assume Gaussian posteriors for both the marginal $q(u)$ and each of the time-dependent conditionals $q(x_{0:T} | u)$. Our proposal is the only variational distribution that is theoretically optimal while yielding a closed-form conditional.

H. Particle Markov chain Monte Carlo (MCMC)

We can use the sequential structure of the latent states for efficient inference. More specifically, using the transition in Eq. (22). The details of this are given in Algorithm 1.

I. Experiment Details

Regarding the variational Gaussian process state space model (VGPSSM), probabilistic recurrent state space model (PRSSM), and VCDT, we download the authors' implementations from their websites. For a fair comparison, we try to follow the same treatment for hyper-parameters across all methods. We note, however, that the Matlab implementation of PRSSM does not provide inference for GP hyper-parameters, conditional likelihood parameters, inducing inputs Z , process variance σ or observation variance σ_o .

I.1. Settings for synthetic data

In generating the synthetic data, we set the kernel's signal variance $\sigma_s = 0$ and lengthscale $l = 0.5$ in the sparse GPSSM. We reduce the impact of observational error and set the observation variance as small as 0.01, i.e., each observation would be generated as $N(x_t; 0:0.01)$. We also set the process variance $\sigma = 0.01$. The number of inducing points is set to 20, and these 20 inducing points are evenly spread in the interval $[2, 2]$. We set the length of the training trajectory as 120, the number of iterations as 50 000 and the number of posterior samples as 500.

I.1.1. TESTING THE MARGINALS FOR GAUSSIANTY

We conduct a hypothesis test for each individual trajectory state and inducing variable. Given the posterior samples for x_t (or u_m), which we denote them as $x_t^{(s)}$ ($u_m^{(s)}$), we use the implementation of `scipy.stats.normaltest` from Python's `scipy` package to test whether we have sufficient evidence to reject the hypothesis that (u_m) follows a Gaussian distribution `scipy.stats.normaltest` is based on the work of [D'agostino & Pearson \(1973\)](#).

I.2. Settings for Real-world Data

We adopt similar settings as in [lalongo et al. \(2019\)](#) to initialize the hyper-parameters for all the models, which first uses a factorized nonlinear model to optimize $x_{0:T}; u; Z; C; d; R; Q$. We use the identity function as the mean function and the squared exponential function with automatic relevance determination (ARD) for GPSSM. When dealing with more than 1 latent dimension, which requires a multi-output GP, we use a different set of kernel hyper-parameters for each output. We set the number of inducing points $M = 100$ and the number of dimensions for latent states to $d_x = 4$. We set standard diagonal Gaussian priors for the initial latent state and (where applicable) for likelihood parameter C , the logarithm of observation standard deviation $\log(R) = -2$ and process variance $\log(Q)$.

For optimizing these hyper-parameters, we use Adam, with default settings for the optimizer parameters except for a decayed learning rate. We run our FVD algorithm for 50 000 iterations and VCDT for 200 000. We set the learning rate as 0.01 and the decay parameter as 0.05 during SGHMC sampling. We use $S = 10^5$ posterior samples for VCDT (as in the results in the original paper) and $S = 100$ for ours. Following the evaluation in [lalongo et al. \(2019\)](#), we use the RMSE between the models' predictions and the ground truth in the future time steps (not seen in training) as the comparison criterion for all the methods.

Regarding the benchmarks, we have 1024 observations for Actuator, 1 000 observations for Ballbeam, 500 observations for Drive, 1 000 observations for Dryer, 1 024 for Flutter, and 296 observations for Furnace. Training lengths are: $T_{\text{train}} = 500; 500; 250; 500; 500; 150$, respectively and test lengths are the rest of the sequences. All these observations are 1-dimensional but, as in previous work, we consider 4-dimensional latent states. Each benchmark also contains a 1-dimensional control input vector $u \in \mathbb{R}^{1 \times 1}$.

I.3. Details of Performance Metrics

We use the test RMSE and the negative mean log likelihood (NMLL).

Given samples from the predictive distribution $y_t^{(s)}$ we can compute the RMSE as:

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T y_t - \text{Ef } y_t^{(s)} g^2}; \quad (75)$$

where $\text{Ef } y_t^{(s)} g$ is estimated as the empirical mean.

if we the corresponding method gives us $\hat{f}_{\mathbf{x}_t^{(s)}}g$, then we can estimate the NMLL as:

$$\text{NMLL} = \frac{1}{S} \frac{1}{T} \prod_{s=1}^S \prod_{t=1}^T \log \mathcal{N}(\mathbf{y}_t; \mathbf{C}\mathbf{x}_t^{(s)} + \mathbf{d}, \mathbf{R}). \quad (76)$$

If the method only provides us with $\hat{f}_{\hat{\mathbf{y}}_t^{(s)}}g$ then a reasonable estimate is:

$$\text{NMLL} = \frac{1}{T} \prod_{t=1}^T \log \mathcal{N}(\mathbf{y}_t; \mathbb{E} \hat{f}_{\hat{\mathbf{y}}_t^{(s)}}g, \mathbb{V} \hat{f}_{\hat{\mathbf{y}}_t^{(s)}}g), \quad (77)$$

where $\mathbb{E} \hat{f}_{\hat{\mathbf{y}}_t^{(s)}}g, \mathbb{V} \hat{f}_{\hat{\mathbf{y}}_t^{(s)}}g$ are the empirical mean and variance of the predicted samples.

For FFVD, we simply use:

$$\text{NMLL}_{\text{FFVD}} = \frac{1}{S} \frac{1}{T} \prod_{s=1}^S \prod_{t=1}^T \log \mathcal{N}(\mathbf{y}_t; \mathbf{C} \hat{\mathbf{x}}_t^{(s)} + \mathbf{d}, \mathbf{C}\hat{\Sigma}_t^{(s)}\mathbf{C}^T + \mathbf{R}), \quad (78)$$

where $\hat{\mathbf{x}}_t^{(s)}$ and $\hat{\Sigma}_t^{(s)}$ are given by Eq. (22).

I.3.1. DATA NORMALIZATION

The normalization (i.e., standardization) of observations is given by the transformation:

$$\tilde{y}_t = \frac{y_t - m_y}{\sigma_y}, \quad (79)$$

where m_y and σ_y are the mean and standard deviation of the training data. Let us denote the $\hat{\text{RMSE}}$ and $\hat{\text{NMLL}}$ as the error metrics in the normalized space. It is easy to show we can obtain the metrics in the original space as:

$$\text{RMSE} = \sigma_y \hat{\text{RMSE}} \quad (80)$$

$$\text{NMLL} = \log \sigma_y + \hat{\text{NMLL}}. \quad (81)$$

I.4. Reproducibility

We will make our code publicly available upon acceptance of the paper.

J. Additional Results

J.1. Training and Test Performance on All Benchmarks

Fig. 5 shows the training and testing regimes and predictions for the six system identification benchmarks considered.

J.2. Traceplot of Log-Likelihood on Training data for FFVD-M and FFVD-C-M

Figure 6 shows the convergence trace plots for all benchmarks.

K. Extended Related Work

Here we give more details about the differences between our method and closely related approaches for inference in the GPSSM. The VGPSSM of Frigola et al. (2014) is the first developing variational approaches to inference in GPSSMs, and it uses mean-field assumptions to factorize the variational distribution of inducing variables \mathbf{u} and latent state trajectories $\mathbf{x}_{0:T}$ as $q_{\text{VGPSSM}}(\mathbf{u}, \mathbf{x}_{0:T}) = q(\mathbf{u})q(\mathbf{x}_{0:T})$. This assumption clearly overlooks the complex dependencies between \mathbf{u} and $\mathbf{x}_{0:T}$. Doerr et al. (2018) model these dependencies with their PRSSM algorithm, but make the unrealistic assumption of the posterior dynamics over the latent state trajectories being the same as the prior. Ialongo et al. (2019) show that this assumption can have critical consequences on the posterior and predictive distributions and propose a more general

