



Hilbert space methods for approximating multi-output latent variable Gaussian processes

Soham Mukherjee^{1,3} · Manfred Claassen^{2,3} · Paul-Christian Bürkner¹

Received: 3 October 2025 / Accepted: 12 March 2026

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2026

Abstract

Gaussian processes are a powerful class of non-linear models, but have limited applicability for larger datasets due to their high computational complexity. In such cases, approximate methods are required, for example, the recently developed class of Hilbert space Gaussian processes. They have been shown to significantly reduce computation time while retaining most of the favorable properties of exact Gaussian processes. However, Hilbert space approximations have so far only been developed for uni-dimensional outputs and manifest (known) inputs. Thus, we generalize Hilbert space methods to multi-output and latent input settings. Through extensive simulations, we show that the developed approximate Gaussian processes are indeed not only faster, but also provide similar or even better uncertainty calibration and accuracy of latent variable estimates compared to exact Gaussian processes. While not necessarily faster than alternative Gaussian process approximations, our new models provide better calibration and estimation accuracy, thus striking an excellent balance between trustworthiness and speed. We additionally illustrate our methods on a real-world case study from single cell biology.

Keywords Approximate GPs · Bayesian inference · Hilbert space · Latent variables · MCMC · Single-cell RNA

1 Introduction

Gaussian processes (GPs) are a popular class of non-linear models well-known for their flexibility in modeling a wide range of data scenarios (Williams and Rasmussen 1995). They are also well-known for their steep computational complexity for model fitting, scaling cubically with the sample size. Thus, exact GPs cannot be used to model large datasets within feasible time. These scalability issues are even more pronounced for extensions of standard GP models. Here, we consider a combination of GPs for latent variable estimation (Lawrence 2003, 2005) and GPs with multi-dimensional outputs (Cressie 1993; Teh et al. 2005). Multi-output GPs are tasked to jointly model several, potentially correlated outcome variables, each with their own

GP. In turn, estimating latent variables via GPs not only increases the number of estimable parameters drastically, but also requires multiple (potentially many) outputs to achieve sufficient estimation accuracy. Multi-output GPs for latent variable estimation have important applications, for example, in single-cell biology where estimating latent cell orderings from multi-dimensional RNA gene sequencing data is crucial to understand the dynamics of the biological process (Haque et al. 2017). Yet, when using exact GPs for this purpose, only small datasets can be analyzed, which strongly limits their practical applicability.

Several methods have been proposed as scalable approximations for multi-output GPs (Álvarez and Lawrence 2011; Hensman et al. 2013a) and latent variable GPs (Titsias 2009; Titsias and Lawrence 2010). These methods primarily depend on reduced rank representations that approximate the GP covariance matrix using inducing points (Seeger et al. 2003; Snelson and Ghahramani 2005; Quinonero-Candela and Rasmussen 2005) followed by approximate model inference via mean-field variational inference (VI) (Titsias 2009). These GPs scale linearly with sample size and quadratically only with the number of inducing points. Thus, they are computationally efficient as long as the number of inducing points is not too large. However, the lack of support for custom

✉ Soham Mukherjee
soham.mukherjee@tu-dortmund.de

¹ Department of Statistics, Technische Universität Dortmund, Vogelpothsweg 87, 44227 Dortmund, Germany

² Department of Internal Medicine, Universitätsklinikum Tübingen, Tübingen, Germany

³ Department of Computer Science, University of Tübingen, Tübingen, Germany

prior specifications of GP hyperparameters across multiple outputs. Although GP parameters can be fixed to a certain value, without varying custom priors, applicability for such models can be limiting for complex data scenarios. Moreover, an in-depth study of the statistical properties of latent variable estimates using these methods are, to the best of our knowledge, yet to be discussed.

An alternative approach is to approximate the covariance function through its spectral decomposition computed from a finite set of representative basis functions. The latter method falls under the category of Hilbert space approximate GPs (HSGPs; Solin and Särkkä 2020; Riutort-Mayol et al. 2022). By exploiting the spectral representation of a stationary covariance function, the computational complexity of HSGPs scales linearly with both sample size and the number of basis functions. What is more, they come with powerful diagnostics that indicate whether the chosen number of basis functions was sufficient for an accurate GP approximation given the data at hand (Riutort-Mayol et al. 2022). HSGPs are typically estimated via Markov-chain Monte Carlo (MCMC), which is slower than VI but produces posterior approximations of much higher quality in terms of accuracy and posterior uncertainty estimation (Yao et al. 2018). The increased, diagnosable accuracy of HSGPs combined with their efficiency has a lot of potential for latent variable estimation. However, HSGPs have so far only been developed for single outputs and manifest (known) inputs. We develop extensions for the HSGP framework to address both latent variable inputs and multi-dimensional outputs and compare their benefits over exact GPs as well as other GP approximation methods.

1.1 Overview of Contributions

- We generalize HSGPs for latent variable inputs and multi-dimensional outputs.
- We validate our new HSGPs against the corresponding exact GPs (where feasible) as well as against variational GPs based on inducing points approximation through a wide range of simulation scenarios.
- Our extensive simulation studies investigate the statistical properties of HSGPs in terms of latent variable estimation accuracy, uncertainty calibration, and model convergence.
- We illustrate the applicability of our HSGPs on real-world single-cell RNA sequencing data, whose analysis is infeasible with similarly specified exact GPs.

2 Related work

Among the extensions of GPs, previous works discuss exact multi-output GPs (Teh et al. 2005) as well as their scalable

counterparts (Álvarez and Lawrence 2011; Hensman et al. 2013a) for high-dimensional data. Recent works in multi-output GPs (Moreno-Muñoz et al. 2018; Joukov and Kulic 2022) also account for output-dimension specific information using different GP hyperparameters for each output dimension. In case of latent variable GPs (Lawrence 2003, 2005) current approximations use a combination of inducing points (Snelson and Ghahramani 2005; Quinonero-Candela and Rasmussen 2005) and mean-field VI (Titsias 2009; Titsias and Lawrence 2010) for scalable solutions. A combination of the above approximation methods for multi-output and latent variable GPs are presented in Lalchand et al. (2022). In this paper, we consider HSGPs (Solin and Särkkä 2020; Riutort-Mayol et al. 2022) and extend it for multi-dimensional outputs and latent variable inputs. Recently, Mukherjee et al. (2025) developed exact multi-output latent variable GPs with derivative information. We only take inspirations from the relevant parts (multi-output and latent inputs) of that work for our model development and experimental design.

In terms of real-world applications, we highlight the field of single-cell biology. There, multi-output structures naturally arise in a plethora of research problems (Haque et al. 2017) and latent GPs are common to model cellular ordering (Hensman et al. 2013b; Campbell and Yau 2015). Later works involve estimating the cellular ordering (Reid and Wernisch 2016; Campbell and Yau 2018) along with inference on different cellular pathway dynamics. A direct combination of the above multi-output latent GPs for single-cell RNA sequencing data is presented in Ahmed et al. (2019), which used the combination of inducing points and VI methods for estimating cellular ordering. As a real-world case study, we showcase our models for a single-cell data of the cell cycle (Mahdessian et al. 2021).

3 Methods

We initially discuss the structure of exact GPs and HSGPs for univariate output y and covariate (known) with paired samples (y_i, x_i) for $i = 1, 2, \dots, N$ where N is the sample size. Then, we generalize the Hilbert space framework for multi-dimensional outputs and latent inputs before discussing methods of inference and alternate GP approximation methods.

3.1 Gaussian processes

A Gaussian process (GP) is a stochastic process specified by a mean function $\mu = \mu(x)$ and a covariance function $k = k(x, x')$ where x, x' are two arbitrary input values. We write $f \sim \mathcal{GP}(\mu, k)$ to indicate that $f = f(x)$ is distributed as a GP such that any finite subset of f jointly follows a multivariate Gaussian distribution. For a univariate output y ,

a standard Gaussian process regression is given by

$$y_i = f(x_i) + \varepsilon_i \tag{1}$$

where ε_i is the i^{th} sample of the independent additive noise ε with error variance σ^2 such that $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The entire model can thus be written as

$$y_i | f \sim \mathcal{N}(f(x_i), \sigma^2). \tag{2}$$

For $i \neq j$, $\text{Cov}(y_i, y_j) = k(x_i, x_j)$ such that $i, j = 1, 2, \dots, N$. When $i = j$, $\text{Cov}(y_i, y_j) = \text{Var}(y_i) = k(x_i, x_j) + \sigma^2$. In terms of k , the Matern class of stationary covariance functions is a highly common choice (Williams and Rasmussen 1995; Rasmussen and Williams 2006). Among them, we specifically consider Squared Exponential (SE), Matern 3/2 and 5/2, which are defined as:

$$\begin{aligned} k_{se} &= \alpha^2 \exp\left(-\frac{r^2}{2\rho^2}\right), \\ k_{m32} &= \alpha^2 \left(1 + \frac{\sqrt{3r^2}}{\rho}\right) \exp\left(-\frac{\sqrt{3r^2}}{\rho}\right), \\ k_{m52} &= \alpha^2 \left(1 + \frac{\sqrt{5r^2}}{\rho} + \frac{5r^2}{3\rho^2}\right) \exp\left(-\frac{\sqrt{5r^2}}{\rho}\right), \end{aligned} \tag{3}$$

where $r = |x - x'|$ is the distance between two arbitrary inputs, $\alpha > 0$ is the marginal standard deviation (SD) and $\rho > 0$ is the length scale parameter. The above covariance functions vary in terms of their implied functional smoothness, with the SE covariance function being the most and the Matern 3/2 being the least smooth. We use a constant mean function μ (as an estimated parameter independent of inputs) for all of our exact GPs and HSGPs discussed in this paper.

3.2 Hilbert space approximations

In the HSGP framework (Solin and Särkkä 2020; Riutort-Mayol et al. 2022), stationary covariance functions are approximated using their spectral densities in the frequency domain. Briefly, Bochner’s theorem (Gihman and Skorokhod 2004) states that a covariance function of a stationary process can be expressed as the Fourier transform of a positive finite measure. If that measure has a density, then we call it a spectral density of the covariance function. Further, using Wiener-Khinchine theorem (Chatfield 2003), we can show that a covariance function and its spectral density are Fourier duals. Combining the results from both these theorems, any arbitrary function satisfying the criteria of being a covariance function (symmetric and positive-definite) can be expressed in terms of its spectral density. The spectral densities corresponding to the covariance functions from Eq.(3) are given

by

$$\begin{aligned} S_{se}(\omega) &= \alpha^2 \rho(\sqrt{2\pi}) \exp\left(-\frac{1}{2}\rho^2 \omega^T \omega\right), \\ S_{m32}(\omega) &= \alpha^2 \left(\frac{2\Gamma(2)3^{3/2}}{\rho^3/2}\right) \exp\left(\frac{3}{\rho^2} + \omega^T \omega\right)^{-2}, \\ S_{m52}(\omega) &= \alpha^2 \left(\frac{2\Gamma(3)5^{5/2}}{3\rho^5/4}\right) \exp\left(\frac{5}{\rho^2} + \omega^T \omega\right)^{-3}, \end{aligned} \tag{4}$$

where $\omega \in \mathbb{R}^N$ is the input in the frequency domain, and $\rho > 0$ and $\alpha > 0$ are the covariance function hyperparameters. The spectral densities can be derived for the general Matern class of covariance functions (Riutort-Mayol et al. 2022), but for this paper, we specifically focus on the above special cases.

We follow the procedure of Riutort-Mayol et al. (2022) and start by defining the compact cover $\Omega \in [-L, L] \subset \mathbb{R}$ containing each element x from the input space for any positive real number L . Following this, we can write any stationary covariance function with $x, x' \in \Omega$ as

$$k_\theta(x, x') = \sum_{j=1}^{\infty} S_\theta(\sqrt{\lambda_j}) \phi_j(x) \phi_j(x'), \tag{5}$$

where S_θ is the spectral density of the covariance function k_θ with θ being the set of hyperparameters. The sets of eigenvalues $\{\lambda_j\}_{j=1}^{\infty}$ and eigenfunctions $\{\phi_j\}_{j=1}^{\infty}$ of the Laplacian operator in the domain Ω satisfy the following problem in Ω under Dirichlet boundary conditions (Solin and Särkkä 2020):

$$\begin{aligned} -\nabla^2 \phi_j(x) &= \lambda_j \phi_j(x), \quad x \in \Omega \\ \phi_j(x) &= 0, \quad x \notin \Omega \end{aligned} \tag{6}$$

where the eigenvalues $\lambda_j > 0$ are real positive due to the Laplacian being a positive-definite Hermitian operator. Here, the ϕ_j are sinusoidal functions and the solution to the eigenvalue problem in Eq.(6) are independent of the choice of covariance functions. They are given by

$$\begin{aligned} \lambda_j &= \left(\frac{j\pi}{2L}\right)^2, \\ \phi_j(x) &= \sqrt{\frac{1}{L}} \sin(\sqrt{\lambda_j}(x + L)). \end{aligned} \tag{7}$$

We approximate the covariance function using the linear combination of its basis functions. Considering the first M number of basis terms, we obtain from Eq.(5),

$$k_\theta(x, x') \approx \sum_{j=1}^M S_\theta(\sqrt{\lambda_j}) \phi_j(x) \phi_j(x'). \tag{8}$$

The full covariance matrix K for inputs $x_i, i \in 1, \dots, N$ is thus approximated by its finite basis function as

$$K \approx \Phi \Delta \Phi^T \tag{9}$$

where $\Delta = \text{diag}\{S_\theta(\sqrt{\lambda_1}), \dots, S_\theta(\sqrt{\lambda_M})\}$ is the diagonal matrix containing spectral densities with eigenvalues and $\Phi \in \mathbb{R}^{N \times M}$ is the matrix of eigenfunctions

$$\Phi = \Phi(\mathbf{x}) = \begin{bmatrix} \phi_1(x_1) & \dots & \phi_M(x_1) \\ \vdots & \ddots & \vdots \\ \phi_1(x_N) & \dots & \phi_M(x_N) \end{bmatrix}. \tag{10}$$

for a vector of input points $\mathbf{x} = (x_1, \dots, x_N)$. We can thus re-write the GP functions from Section 3.1 as

$$f \sim \mathcal{N}(\mu, \Phi \Delta \Phi^T). \tag{11}$$

This is equivalent to the linear representation

$$f(\mathbf{x}) \approx \mu + \sum_{j=1}^M (S_\theta(\sqrt{\lambda_j}))^{1/2} \Phi(\mathbf{x}) \beta_j \tag{12}$$

where $\beta_j \sim \mathcal{N}(0, 1)$ is standard normal. The eigenvalues λ_j are monotonically increasing in j and S_θ rapidly decreases to zero for a bounded covariance functions. As a result, usually only a small number of basis functions M are required for an adequate approximation (Rasmussen and Williams 2006; Riutort-Mayol et al. 2022).

3.3 Extending HSGPs

We extend the HSGP framework to model multi-dimensional outputs and latent variable inputs. We follow a similar framework pertaining to the multi-output latent variable models considered in Mukherjee et al. (2025) (but do not consider the derivative methods here).

3.3.1 Multi-output HSGPs

We specify multi-output GPs with response variables (y_1, \dots, y_D) over $D > 1$ output dimensions (Rasmussen and Williams 2006), using y_{di} to denote the response for dimension d and observation i . As the usual approach, we first set up D independent, univariate Gaussian processes f_d each with their own set of hyperparameters, say θ_d (Teh et al. 2005; Mukherjee et al. 2025). Thus, extending HSGPs to multi-output GPs, we first modify Eq.(12) such that

$$f_d(\mathbf{x}) \approx \mu_d + \sum_{j=1}^M (S_{\theta_d}(\sqrt{\lambda_j}))^{1/2} \Phi(\mathbf{x}) \beta_{jd} \tag{13}$$

where $\beta_{jd} \sim \mathcal{N}(0, 1)$. The univariate GPs are then related to one another by folding them with a (D -dimensional) across-dimension correlation matrix C (Teh et al. 2005; Bonilla et al. 2007). Specifically, for each observation i , we obtain a vector of across-dimension correlated GP values as

$$\begin{pmatrix} f_1^*(x_i) \\ \dots \\ f_D^*(x_i) \end{pmatrix} = A \times \begin{pmatrix} f_1(x_i) \\ \dots \\ f_D(x_i) \end{pmatrix}, \tag{14}$$

where A is the Cholesky factor of C such that $C = AA^T$ with A being lower-triangular. The Cholesky factor A is matrix multiplied (\times) to the vector of across-dimension GP values. This way, multi-output GPs combine two dependency structures, one within dimensions (and across observations) as expressed by the univariate GPs through corresponding covariance functions and one across output dimensions (but within observations) as expressed by C (or A). We specify C through a Cholesky prior $C \sim \text{LKJ}(\eta = 1)$, which is uniform across all valid correlation matrices of a given dimension (Lewandowski et al. 2009). For efficiency reasons, we equivalently estimate A instead of C via MCMC along with all other model parameters. More details on model setup are discussed later in 4.2. Adding independent Gaussian noise to our multi-output GP model, we extend Eq.(2) for all d and i :

$$y_{di} | f_d^* \sim \mathcal{N}(f_d^*(x_i), \sigma_d^2) \tag{15}$$

An alternative way of specifying C is through a set of input-dependent covariance functions (Gneiting et al. 2010; Bourotte et al. 2016). In our study, we however only consider an input-independent C , as modeling the across-dimension correlations is not the main focus of this paper. In Section 6, we provide a more detailed discussion on this topic.

3.3.2 Latent variable HSGPs

Within latent-variable GPs, the inputs x are considered as unobserved and are treated similar to other estimable parameters. To that end, we specify an observed quantity \tilde{x} that guides the estimation of the latent inputs x . From a Bayesian perspective, \tilde{x} acts as a prior-like object for the latent x to be then further refined by the GPs learning from y . Specifically, we consider the scenario where \tilde{x} is a noisy measurement of x . If we assume that the measurements \tilde{x} are Gaussian with known measurement SD s , we can write for each observation i :

$$\tilde{x}_i \sim \mathcal{N}(x_i, s^2). \tag{16}$$

The vector of latent inputs $\mathbf{x} = (x_1, \dots, x_N)$ is then passed to the approximation step in Eq. (12) (and subsequently Eq.

(13) for the multi-output case). We call the resulting model (multi-output) latent variable HSGPs.

Latent variable GPs, exact or approximate, are more difficult to fit than their manifest counterparts. The primary reasons are the substantial increase in the number of estimable parameters as well as identification issues arising due to both x and ρ now being treated as unknown. We alleviate this identifiability issue using MCMC methods discussed in Sect. 3.4. We further present a detailed simulation study investigating the challenges of latent variable GPs in Sect. 4.

3.4 Bayesian inference

We fit HSGPs using full Bayesian inference via MCMC sampling. Our primary parameters of interest are the latent inputs x and GP hyperparameters θ given data (outputs) y . We re-purpose θ to include not only the covariance function hyperparameters ρ and α as specified earlier but also the error SD σ . For a specific output dimension d , we assume independent prior distributions on θ_d as

$$\theta_d \sim p(\theta_d) = p(\rho_d) p(\alpha_d) p(\sigma_d). \tag{17}$$

The joint probability density factorizes as

$$p(y, x, \theta | \tilde{x}) = p(x | \tilde{x}) \prod_d^D p(y_d | x, \theta_d) p(\theta_d). \tag{18}$$

where $p(y_d | x, \theta_d)$ denotes the GP-based likelihood for a single output dimension and $p(x | \tilde{x})$ denotes the prior for the latent x implied by the measurement model in Eq.(16). The details of prior specifications used in our experiments are further discussed in Section 4. Using Bayes' rule, we obtain the joint posterior over x and θ as

$$p(x, \theta | y, \tilde{x}) = \frac{p(y, x, \theta | \tilde{x})}{\int \int p(y, x, \theta | \tilde{x}) dx d\theta}. \tag{19}$$

The inference framework specified above is same for both exact and HSGPs. Posterior samples of x and θ for all output dimensions are obtained via MCMC sampling, specifically adaptive Hamiltonian Monte Carlo (Neal 2011; Hoffman and Gelman 2014). We implemented the exact GP and the HSGP models in Stan using the RStan interface (Development Team 2024).

3.5 Approximate latent variable GPs using VI

Approximating latent variable GPs using VI methods (VIGPs) usually include a combination of approximating both the covariance function and the posterior. Approximating the covariance function is carried out using inducing points to obtain a reduced rank representation of the Gram matrix

(Rasmussen and Williams 2006). There are many inducing point methods (Quinero-Candela and Rasmussen 2005), however, based on the suggestions of Bauer et al. (2016), we specifically consider the implementation discussed in Titsias (2009). In this framework, the estimation of latent variables x and GP hyperparameters θ after marginalizing over f is carried out via the marginal likelihood

$$p(y | x, \theta) = \mathcal{N}(y | \mu, K_{N \times N} + \sigma^2 I), \tag{20}$$

where $K_{N \times N}$ is the full rank GP covariance matrix for a sample size N . This full rank covariance matrix is then replaced by $Q_{N \times N} = K_{N \times M} K_{M \times M}^{-1} K_{M \times N}$ where M is the number of inducing points. This reduces the computational complexity of inverting a full rank matrix to only computing its Nyström approximation. While fundamentally different, the basis points in our method explained in Section 3.2 and inducing points here serve a similar purpose. For that reason, we choose to denote both the number of basis points (in HSGPs) and inducing points (in VIGPs) by M for simplicity. Effectively, the marginal likelihood is then reformulated as

$$p(y | x, \theta^*) = \mathcal{N}(y | \mu, Q_{N \times N} + \sigma^2 I). \tag{21}$$

Following the specifications discussed in Titsias and Lawrence (2010), the joint probability distribution for the d^{th} output dimension is given by

$$p(y_d, f_d, M_d) = p(y_d | f_d) p(f_d | M_d) p(M_d), \tag{22}$$

where M_d is the dimension-specific set of inducing points. The inducing points are designated as variational parameters and optimized through variational inference (Titsias 2009; Titsias and Lawrence 2010). The posterior $p(f_d, M_d | y_d) = p(f_d | M_d, y_d) p(M_d | y_d)$ is subsequently approximated by a sparse variational distribution of the form

$$q_p(f_d, M_d) = p(f_d | M_d) q_M(M_d), \tag{23}$$

where $q_M()$ is the variational distribution over the inducing points M_d . Model inference is carried out by the usual method of computing the variational lower bound (see Titsias and Lawrence 2010 for further details).

While the latent variable GPs were earlier designed for classification tasks (Lawrence 2003, 2005), VIGPs have later been used for GP regressions as well, thus making it an ideal candidate for reference and comparison to our proposed HSGPs. In this paper, we follow the implementations for VIGPs in Titsias and Lawrence (2010); Ahmed et al. (2019) using the GP-LVM module of Bingham et al. (2019) where the user only needs to set the number of inducing points and the internal working takes care of their optimization for sparse latent input estimation.

4 Simulation study

For real-world data, we lack the ground truth values of latent variables. In contrast, in a simulated setting, we have full control over the data generating process including the true latent values. Thus, it is crucial to investigate the statistical properties of latent variable models using simulated data. To this end, we investigate exact GPs, HSGPs and VIGPs under various simulation scenarios exploring their behavior in terms of model convergence, model calibration, and estimation accuracy. The data generating conditions are inspired by Mukherjee et al. (2025).

4.1 Data generating scenarios

We consider six different simulation scenarios. These scenarios are broadly categorized into data generated from a GP and data generated from a periodic function (non-GP data scenario). Under the former category, we consider three data generating processes based on multi-output GPs with SE, Matern 3/2, and Matern 5/2 covariance functions, respectively. We sample length-scale $\rho \sim \text{Normal}^+(1, 0.05^2)$, GP marginal SD $\alpha \sim \text{Normal}^+(3, 0.25^2)$, and error SD $\sigma \sim \text{Normal}^+(1, 0.25^2)$. Additionally, we consider a fourth GP data scenario, where we allow the length-scale of the SE covariance function to be highly variable across each output dimensions, $\rho \sim \text{Normal}^+(1, 0.25^2)$, thus increasing the challenges in model fitting. Each of the hyperparameters are sampled D times, one value per output dimension. Under the GP data scenarios, the fitted models align with the true data generating process (or are approximations to them). However, simulating from GPs results in a lot of variation in the true functional forms across output dimensions and simulation trials.

Compared to that, the periodic data generating process have more consistency since we ensure a similar level of non-linearity across each outputs. Concretely, we use the following periodic process:

$$f_{id} = \alpha_d \sin\left(\frac{x_i}{\rho_d}\right), \quad (24)$$

$$y_{id} \sim \mathcal{N}(f_{id}, \sigma_d^2).$$

In our simulations, we consider two version of this process, one with low oscillations (high true ρ_d values) and one with higher oscillations (low true ρ_d values). These two periodic data scenarios mimic the challenges of modeling starkly different functional smoothness with the higher oscillations being the more challenging one. The hyperparameters of the periodic process closely resemble that of the GP data scenarios and are thus named the same for simplicity. The sampling distributions are also similarly specified as mentioned pre-

viously. For the periodic data with higher oscillation, we specify $\rho \sim \text{Normal}^+(0.5, 0.05^2)$ as compared to the lower oscillation case where $\rho \sim \text{Normal}^+(1, 0.25^2)$.

In all scenarios, we sample the between-dimension (output) correlation matrix $C \sim \text{LKJ}(\eta = 1)$ (Lewandowski et al. 2009). This implies C to be uniformly distributed within the set of all valid correlation matrices of dimension D . For the sample size N , we consider three cases $N = 20, 50$ and 200 . Due to the computational complexity of exact GPs, we only consider them for $N = 20$ case. Further, a special case for the GP data scenarios is also demonstrated for the HSGPs and VIGP with $N = 1000$ in Section 4.6. Unlike the smaller N cases, here the fitted models aren't aligned with the true data generating process, that is, we specify vague priors for our HSGP hyperparameters for $N = 1000$ case. We generate the ground truth latent inputs as $x_i \sim \text{Uniform}(0, 10)$, where $i = 1, \dots, N$. Further, we assume a prior measurement SD of the noisy \tilde{x} as $s = 0.3$ (see Section 3.3). For the number of output dimensions D , we consider the three cases $D = 5, 10$ and 20 . We perform 50 trials for each of the choices of N and D per simulation scenario.

4.2 Model specifications

We fit exact GPs, HSGPs, and VIGPs for all simulation scenarios, with the exception of omitting exact GPs for sample sizes higher than $N = 20$ due to the scalability issues discussed earlier (see Section 1). For the HSGPs, we select the boundary conditions L such that it contains the input space completely. This is obtained by multiplying a scalar adjustment $c = 1.25$ (Riutort-Mayol et al. 2022) to the range of the input prior \tilde{x} . Thus, we ensure that we have a compact cover $[-L, L]$ that prevents input values x to be near the boundaries. By doing so, we ensure stable MCMC model convergence for the latent variable inputs (see Section 4.3). We select the minimum number of basis functions M_{\min} loosely based on the suggestions of Riutort-Mayol et al. (2022). The choices for M_{\min} depends on the choice of covariance function to accurately approximate the degree of non-linearity. Thus, for the SE covariance function,

$$M_{\min} = 1.75 \frac{cS}{\mu_\rho}, \quad (25)$$

for the Matern 3/2,

$$M_{\min} = 3.42 \frac{cS}{\mu_\rho}, \quad (26)$$

and for the Matern 5/2,

$$M_{\min} = 2.65 \frac{cS}{\mu_\rho}. \quad (27)$$

Table 1 Primary parameters of interest

Name	Notation	Number of parameters
Latent inputs	x	N
Length scale	ρ	D
GP marginal SD	α	D
Error SD	σ	D

Note: N stands for number of observations or sample size. D denotes the number of output dimensions

In the above equations, c is the scaler adjustment to the boundary conditions, S is the range of inputs and μ_ρ is the mean of the length-scale prior. The SE covariance function requires much lower M_{\min} compared to, say, the Matern 3/2 which is on the other end of the smoothness spectrum among the Matern class.

Prior specifications for the exact GPs and HSGPs are kept same for all model parameters. Our prior choices are aligned with the data generating process since that is required for testing model calibration (see Section 4.4). Thus, we specify priors for length-scale $\rho \sim \text{Normal}^+(1, 0.05^2)$, GP marginal SD $\alpha \sim \text{Normal}^+(3, 0.25^2)$ and error SD $\sigma \sim \text{Normal}^+(1, 0.25^2)$. For the periodic data with higher oscillation, we specify a lower prior mean of $\rho \sim \text{Normal}^+(0.5, 0.05^2)$. Lastly, in case of the SE simulation scenario with highly varying length-scale, we specify a prior of $\rho \sim \text{Normal}^+(1, 0.25^2)$. The number of estimated primary parameters of interest for each HSGPs are provided in Table 1.

Following the suggestions of the number of inducing points discussed in Titsias (2009), we use $M = 10$ inducing points for the VIGPs. In this framework, a small constant is added to the diagonal terms of the matrix. This helps maintain numerical stability while computing the GP covariance matrix. In case of VIGPs, we tested different values of diagonal constants starting from $\delta = 10^{-12}$ (usually suggested default). Based on our tests, for 10 inducing variables (points), the VIGPs need at least $\delta = 10^{-3}$ for a numerically positive definite GP covariance matrices and reliably finish an entire simulation scenario. Higher number of inducing points would require an even larger δ . Since $\delta = 10^{-3}$ is already quite large, we keep $M = 10$ to not further increase the bias caused by large values of δ . As in the other methods, we set the latent prior measurement SD to $s = 0.3$. As for the GP hyperparameters, since the VIGP implementation in Pyro does not support custom priors we keep them at their uniform defaults for the simulation scenarios with $N = 20, 50$ and 200 . For the $N = 1000$ (see Section 4.6), we introduce a VIGP case where GP hyperparameters are fixed to the data generating conditions, essentially creating a favorable condition for VIGPs.

Both exact GPs and HSGPs are implemented in Development Team (2024) and fitted with a single MCMC chain of

2000 iterations of which the first 1000 are discarded as warm-up. As shown in Mukherjee et al. (2025), model convergence is similar for multiple chains when fitting latent variable GPs. Thus, we run only a single chain per model to reduce overall computation times. The simulation studies are conducted on 50 vCPUs (Intel(R) Xeon(R) Gold 6230R CPU @ 2.10 GHz) with 720 GB work memory. The average runtime per dataset for the HSGPs with $N = 200$ and $D = 20$, the most computationally expensive case, was approximately 28 minutes for the SE model, 42 minutes for the Matern 3/2 model, and 34 minutes for the Matern 5/2 model. In the same scenarios, the VIGPs took approximately 2 minutes regardless of the covariance function. The exact GPs (just for $N = 20$ and $D = 20$) took 48 minutes for the SE model, 1.6 hours for the Matern 3/2 and 2.4 hours for the Matern 5/2 model. The HSGPs were significantly faster than the exact GPs but slower than VIGPs. The slower speed of HSGPs compared to VIGPs is counter-balanced by highly consistent model calibrations, superior estimation accuracy for latent variables, and overall much more stable model performance, as shown below.

4.3 Model convergence

We investigate MCMC convergence of our fitted exact GPs and HSGPs for all of the six simulation study scenarios discussed before. We use standard MCMC sampling diagnostics including state-of-the-art versions of the scale reduction factor \hat{R} , bulk effective sample size (Bulk-ESS) and tail effective sample size (Tail-ESS) (Vehtari et al. 2021). A combined check of these measures provide a comprehensive picture of the parameter-specific model convergence. In general, \hat{R} should be very close to 1 and should ideally not exceed 1.01 (Vehtari et al. 2021). We additionally consider a more relaxed threshold of 1.1 in our simulation studies, since the ground truth is available as another layer of evaluation. Bulk-ESS indicates the reliability of measures of central tendency such as the posterior mean or median. Tail-ESS indicates the reliability of the 5% and 95% quantile estimates, which are then used to construct credible intervals. Both Bulk-ESS and Tail-ESS should have values greater than 100 times the number of MCMC chains (higher is better). We compute all the convergence measures with the posterior package (Bürkner et al. 2023). Inference of the VIGPs is done via numerical optimization, which we considered to have converged if no convergence warnings were issued by the optimizer.

In Fig. 1, we show \hat{R} , Bulk-ESS, and Tail-ESS for the latent x and GP hyperparameters of both exact GPs and HSGPs with SE covariance function. While the exact GPs reach and exceed the relaxed \hat{R} threshold of 1.1 for some simulated datasets, the HSGPs consistently satisfy the much stricter 1.01 threshold of model convergence. HSGPs subsequently also have much higher Bulk and Tail-ESS. Overall,

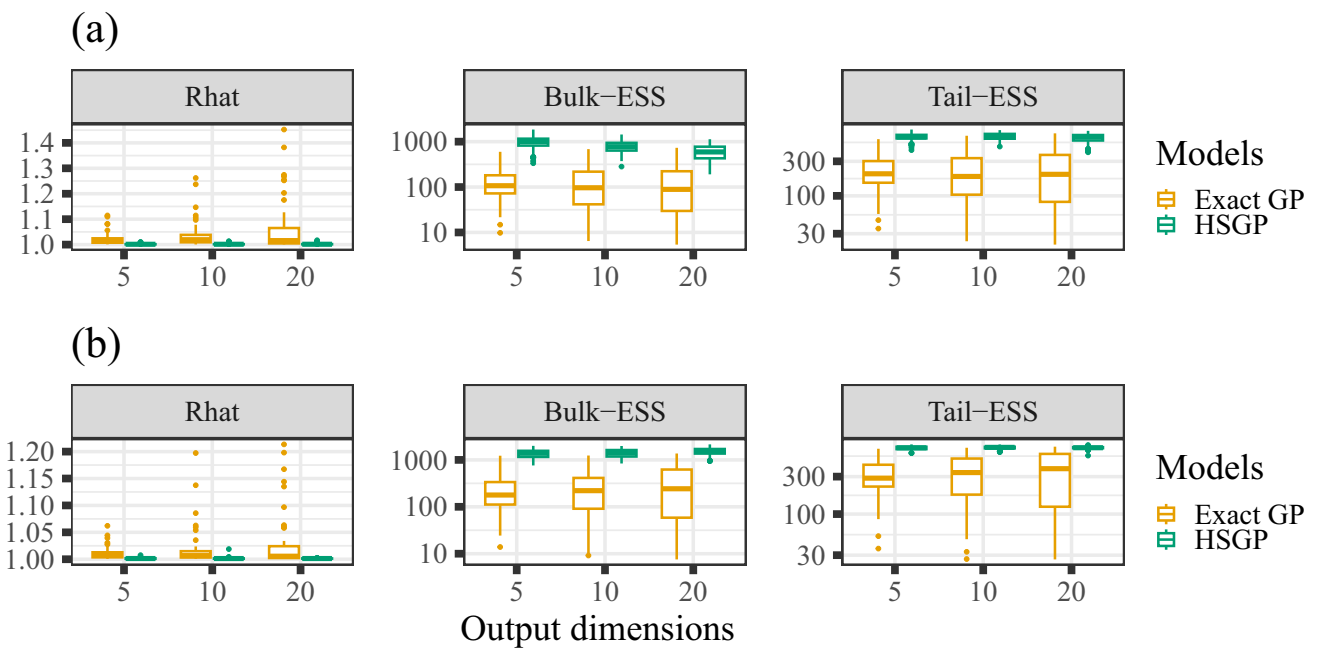


Fig. 1 Squared exponential scenario: Convergence check for (a) latent inputs and (b) GP hyperparameters of the exact GPs (for $N = 20$) and HSGPs (combined for $N = 20, 50$ and 200 cases). The y-axes for Bulk and Tail ESS plots are \log_{10} transformed

based on the diagnostics, HSGPs show much more consistent and stable convergence as compared to exact GPs. We provide convergence diagnostic figures for the other simulation scenarios in the supplementary materials Section C (Figures S3-S7) where we see similar results. We additionally provide further analysis for MCMC convergence of latent inputs and GP hyperparameters compared between exact GPs and HSGPs in supplementary materials Section D (Figures S8-S11).

4.4 Testing model calibration

We use simulation based calibration (SBC) (Modrák et al. 2023; Talts et al. 2020) to test model calibration for estimating latent inputs x . Through SBC, we check if the model is able to produce posterior distributions that are consistent with the data generating process, thus correctly accounting for the implied uncertainty.

To that end, briefly, SBC aims to test the goodness of posterior approximations by exploiting self-consistency properties of Bayesian models. This results in a rank statistic that is uniformly distributed under the assumption of a well calibrated model (Modrák et al. 2023), which can subsequently be checked both graphically and numerically. As a graphical test, we plot the empirical cumulative distribution function (ECDF) of the ranks along with their 95% confidence regions under the assumption of uniformity (Säilynoja et al. 2022). ECDFs lying outside of their confidence region indicate miscalibrations. These graphical tests also yield a test statistic

(the $\log \gamma$ score; see Säilynoja et al. 2022), which we further analyze numerically. More details on SBC is provided in supplementary materials Section A.

In Fig. 2, we show the rank ECDFs obtained from the exact GPs, HSGPs, and VIGPs for the SE simulation scenario ($N = 20$ and $D = 20$). The results indicate that the latent variable estimates obtained from the HSGPs are consistently well calibrated. In contrast, for exact GPs, we see miscalibrations for some latent variables shown by their corresponding rank sample ECDFs lying outside the threshold region. In case of VIGPs, the instances of miscalibrations are even more severe. The rank ECDFs for Matern 3/2 and 5/2 scenarios are reported in supplementary materials Section E (Figures S12 and S13) where the results are same as the SE scenario.

To streamline calibration checking across simulation scenarios, we analyze the $\log \gamma$ scores using a multi-level model (Bürkner 2017) described in supplementary materials Section A. In Fig. 3, we show the predicted $\log \gamma$ scores of the different GP models under the SE simulation scenario. This is equivalent to summarizing the different cases of rank ECDFs across all simulation scenarios.

For all the choices of N and D , the latent variable estimates from HSGPs are consistently well calibrated. By comparison, the exact GPs fail the calibration test for $D = 20$ as their $\log \gamma$ scores remain consistently less than the critical value corresponding to the 95% confidence threshold. This can be primarily attributed to the computational issues of exact GPs for latent variable inputs as seen through MCMC con-

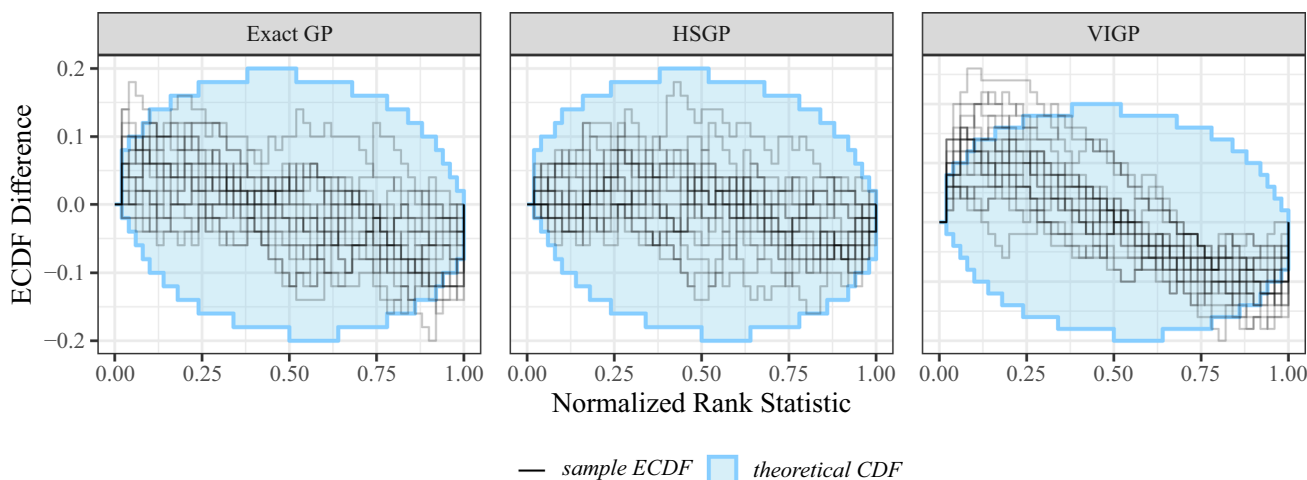


Fig. 2 Squared exponential scenario: ECDF-difference calibration plots of the latent x estimated by exact GP, HSGP, and VIGP. Only the HSGP is consistently well calibrated

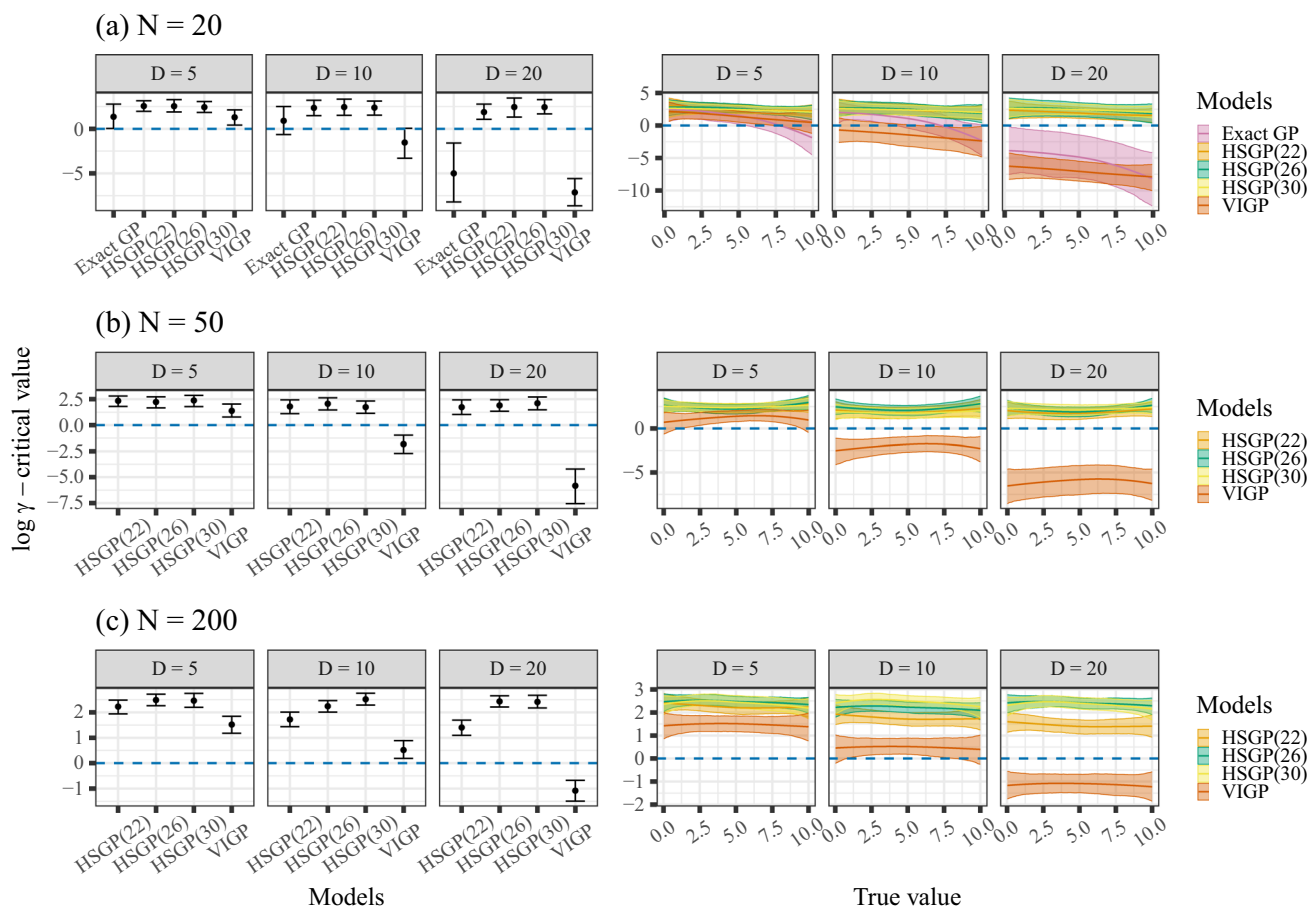


Fig. 3 Squared exponential scenario: $\log \gamma$ scores offset by the 95% confidence threshold for all the fitted models. The behavior of scores across true latent x values are shown in the right-hand panel. The blue

dashed line denotes the threshold to reject uniformity, that is, models with values less than 0 are miscalibrated. The HSGP(M) shows the HSGPs with their corresponding number of basis functions

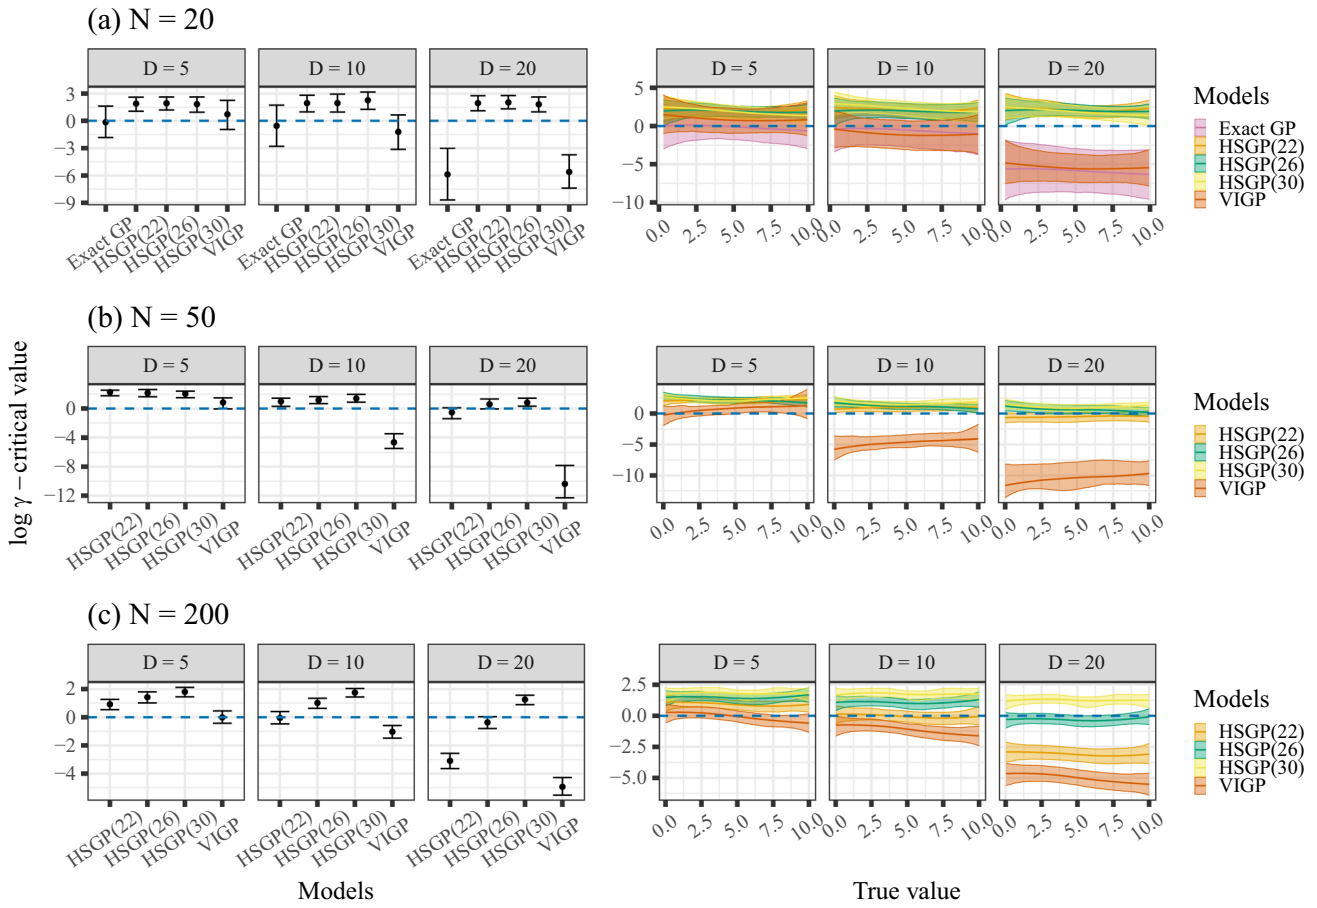


Fig. 4 Squared exponential scenario (highly varying ρ): $\log \gamma$ scores offset by the 95% confidence threshold for all the fitted models. The behavior of scores across true latent x values are shown in the right-hand

panel. The blue dashed line denotes the threshold to reject uniformity, that is, models with values less than 0 are miscalibrated. The HSGP(M) shows the HSGPs with their corresponding number of basis functions

vergence diagnostics in Sect. 4.3 along with supplementary materials Section C and D. The VIGPs shows miscalibrations especially for higher output dimension cases throughout all the choices of N . Additionally, exact GPs exhibit worse calibrations (decreasing $\log \gamma$ scores) when estimating larger values of the latent inputs. This issue directly causes problems in estimating latent values closer to the boundary. VIGPs show similar behavior for lower sample sizes $N = 20$ although not as strong as exact GPs. The HSGPs are able to overcome and rectify this behavior. For the special scenario of SE data with highly variable length-scales as shown in Fig. 4, VIGPs and HSGPs with small number of basis function struggle to pass the calibration tests, especially for $N = 200$. The miscalibrations are however alleviated with a higher number of basis functions in HSGPs. The $\log \gamma$ results for the Matern 3/2 and 5/2 simulation scenarios are presented in supplementary materials Section E (Figures S14 and S15).

4.5 Latent variable estimation

For all of the simulation scenarios discussed in Sect. 4.1, we evaluate the estimation capabilities of the different fitted models specified in Sect. 4.2. We show model performance in terms of absolute bias $|\text{Bias}(x_{post}, x_{true})|$ of the latent posterior estimates x_{post} with respect to their true values x_{true} as well as $\text{SD}(x_{post})$ as a measure of posterior sharpness. Since mean-field VI is known to underestimate the posterior SD (Kucukelbir et al. 2017), we avoid using a metric like RMSE that favors (too) narrow posteriors, thus providing misleading interpretations if used as a sole metric.

We summarize our findings using the multi-level model (Bürkner 2017) described in supplementary materials Section A for both posterior bias and SD. The SE simulation scenario presented in Fig. 5 shows that, in most cases, HSGPs have lower bias compared to exact GPs and VIGPs. The VIGPs consistently show the highest bias. Overall, for the SE data scenario, we see on average about 11% (for higher sample sizes) to about 18% (for lower sample sizes) decrease

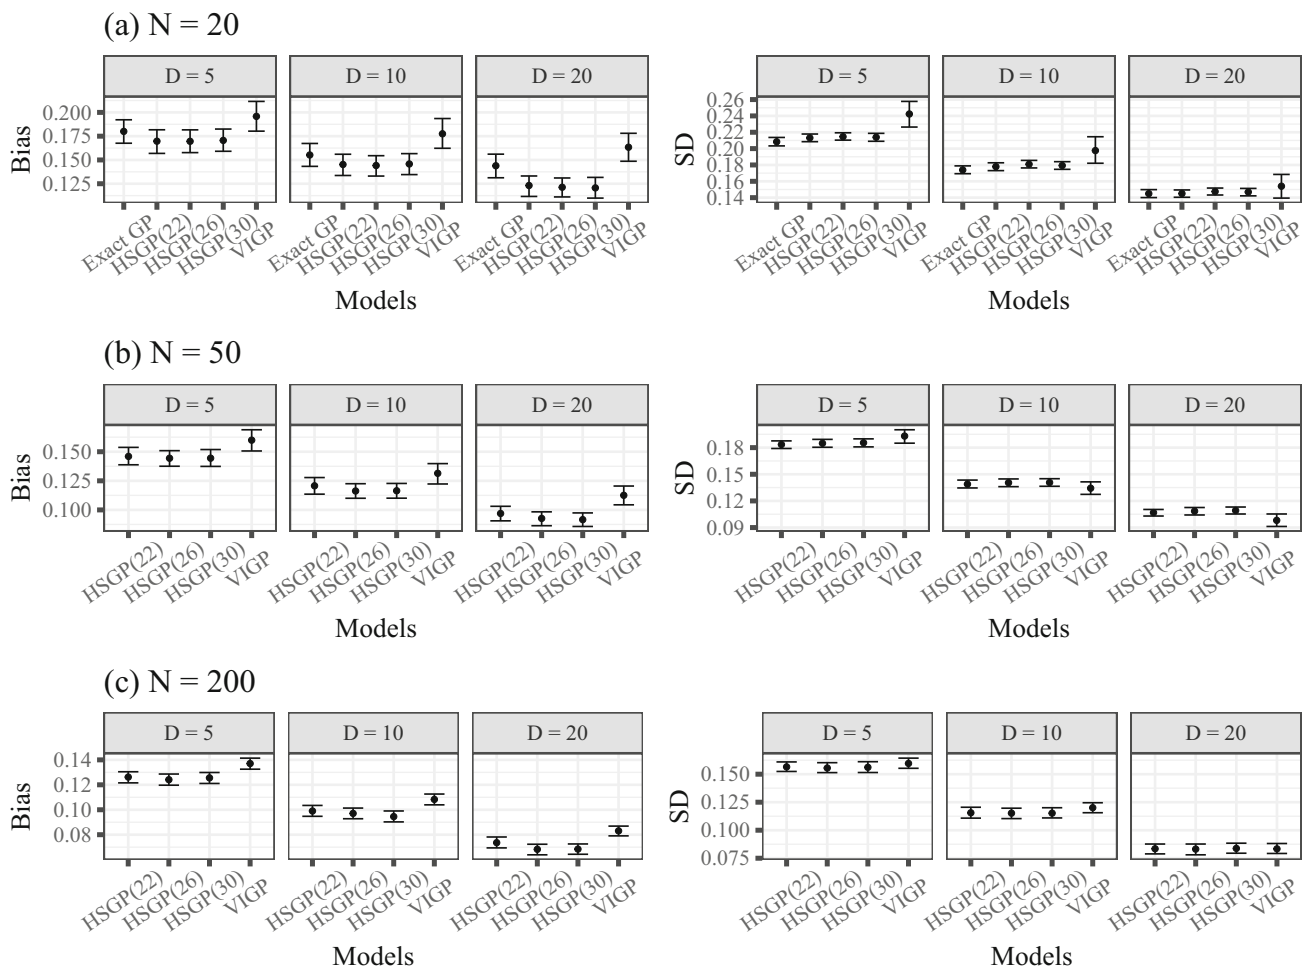


Fig. 5 Squared exponential scenario: posterior bias and SD on recovery of latent inputs for all fitted models. The HSGP(M) shows the HSGPs with their corresponding number of basis functions

in posterior bias between VIGPs and HSGPs (with $M = 22$ basis functions). Results for the posterior SD are more mixed, with VIGPs showing higher SDs for small sample sizes and about the same SDs for larger sample sizes. Specifically, HSGPs overcome the computational challenges (in MCMC convergence and posterior uncertainty calibration) of exact GPs to provide better latent variable estimation accuracy.

In Figure 6 we show the more challenging simulation scenario. Here, the data generating process had a higher amount of variability in length-scales across the output dimensions. We observe a stronger difference in posterior bias between VIGPs and HSGPs (in favor of HSGPs). We hypothesize that HSGPs are more strongly favored here as they allow for varying hyperparameters across dimensions, while the VIGPs (at least their Pyro implementation) only support single hyperparameters across all output dimensions. For example, a single length scale parameter ρ is optimized for all D output dimensions. While this is a limitation of the current VIGP implementation, we compare a favorable case where the GP

hyperparameters are fixed to the data generating conditions in the special simulation scenario with $N = 1000$ in Section 4.6.

Among the HSGPs, higher number of basis functions ($M = 26$ and 30 compared to $M = 22$) result in lower posterior bias for both the simulation scenarios shown in Figs. 5 and 6. Combining this result with their corresponding model calibration tests, we recommend using $M = 30$ to achieve the highest accuracy for the latent variable estimates. For faster but reliable results, HSGPs with $M = 22$ can be used in simpler cases like the SE data scenario. Furthermore, in the remaining scenarios (see Figures S16-S19 in supplementary materials Section F), HSGPs perform consistently better than their competitors. In all the scenarios, each of the fitted models (exact GPs only for $N = 20$) show a consistent decrease in posterior bias and SD as the number of output dimensions D and sample size N increases. We report the estimation results for GP hyperparameters in the supplementary materials Section B (followed by additional figures in Section H).

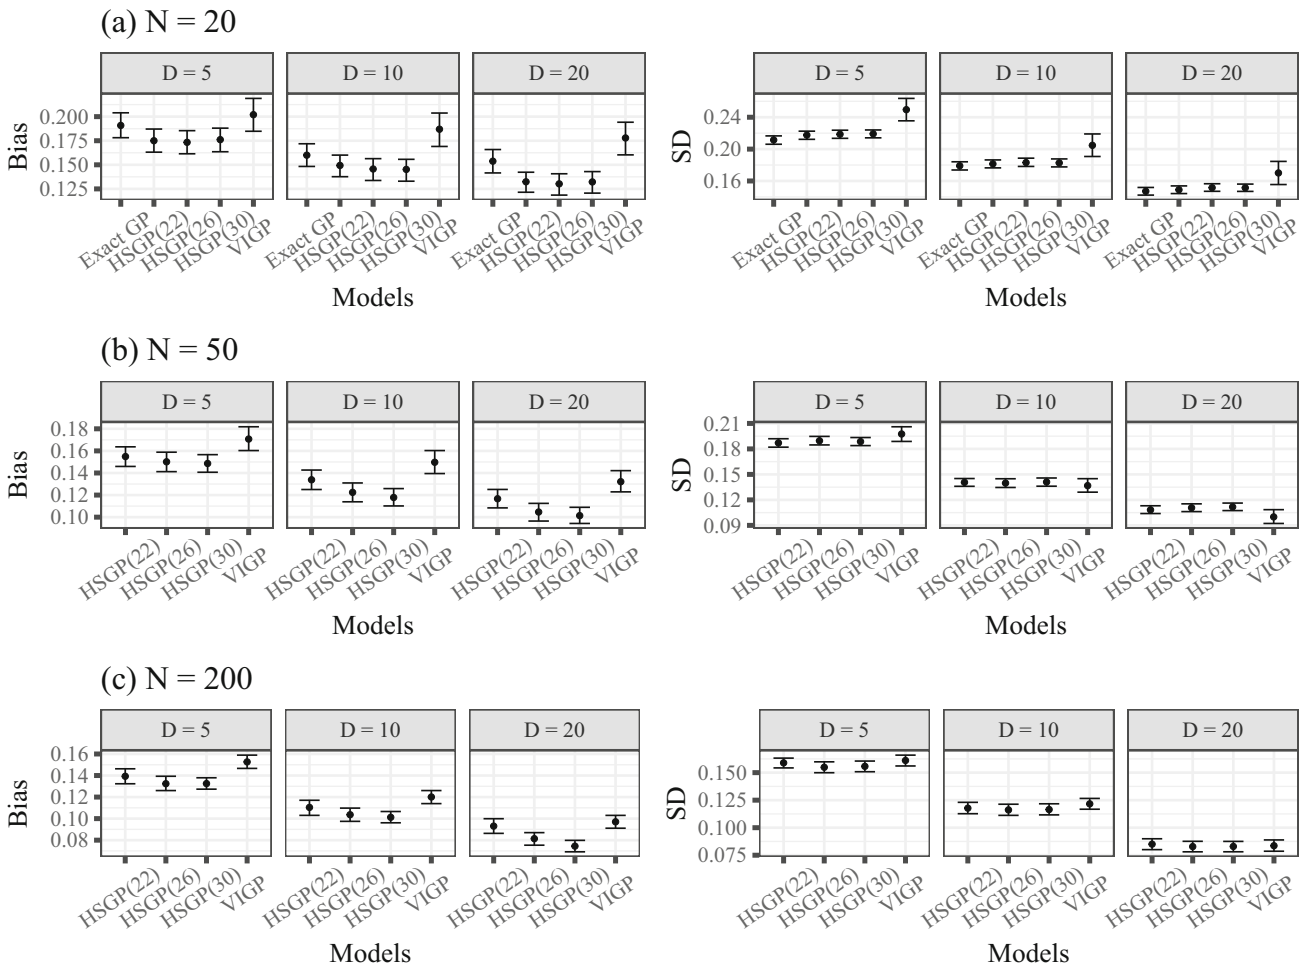


Fig. 6 Squared exponential scenario (highly varying ρ): posterior bias and SD on recovery of latent inputs for all fitted models. The HSGP(M) shows the HSGPs with their corresponding number of basis functions

4.6 Checking robustness of HSGPs with wide priors

In the above discussed GP simulation scenarios, the hyperparameter priors for exact GPs and HSGPs are aligned with the data generating conditions as a requirement for SBC checks (see Sect. 4.4). Therefore, in those simulation studies, the exact GPs and HSGPs receive a favorable treatment as compared to VI-GPs. Hence, to check the robustness of latent variable HSGPs, we perform additional simulations where we apply hyperparameter priors to the HSGPs that are less informed (i.e., wider priors) along with a larger sample size $N = 1000$. We leave exact GPs out of this simulation study due to their computational limitations.

Keeping the data generating conditions same as discussed in Section 4.1, we specify priors for length-scale $\rho \sim \text{Normal}^+(1.5, 0.5^2)$, GP marginal SD $\alpha \sim \text{Normal}^+(3.5, 1^2)$ and error SD $\sigma \sim \text{Normal}^+(1.5, 1^2)$. That is, we shift the prior mean away from the data generating conditions along with a much wider prior variance. A comparison of the data

Table 2 Comparison of GP hyperparameter priors and data generating conditions

Hyperparameter	Dgp prior	Model prior
ρ	$\text{Normal}^+(1, 0.05^2)$	$\text{Normal}^+(1.5, 0.5^2)$
α	$\text{Normal}^+(3, 0.25^2)$	$\text{Normal}^+(3.5, 1^2)$
σ	$\text{Normal}^+(1, 0.25^2)$	$\text{Normal}^+(1.5, 1^2)$

Note: The Dgp priors show the true sampling distribution for simulated data and Model priors show the prior distribution for the HSGP parameters

generating conditions and the GP hyperparameter priors are presented in Table 2. Simulations are carried out for SE, Matern 3/2 and 5/2 covariance functions. Based on the specified length-scale prior and chosen covariance function, the minimum required number of basis functions for HSGPs change according to the empirical relation shown in Eq. (25)-(27).

We consider two VIGP variants: VIGP (default) and VIGP (prior). The former is the same VIGP model considered in other simulation scenarios with the default priors used by Pyro. In contrast, for VIGP (prior), we fix the GP hyperparameters to the mean of the true sampling distributions, thereby acting as a highly informed reference.

We summarize the posterior bias and SD using the same multilevel model (described in supplementary materials Section A) used for other simulation scenarios. For the SE scenario in Fig. 7, the HSGPs achieve lower bias compared to VIGPs with a higher number of basis functions M . This is especially true for higher output dimensions $D = 20$ where we need much higher M than the minimum M suggested by the empirical relation in Eq.(25). On the other hand, posterior SDs for HSGPs are consistently lower than VIGPs. Among the VIGP variants, both VIGP (default) and VIGP (prior) performs similarly in terms of latent variable estimation accuracy. That is, fixing the GP hyperparameters for VIGPs to informed values does not result in any improvement in model performance. The results for Matern 3/2 and 5/2 covariance functions (see Figures S20 and S21 in supplementary materials section G) are highly similar, with Matern 3/2 showing even more pronounced differences between HSGPs and VIGPs.

The average runtime per dataset of sample size $N = 1000$ on the computing cluster (details mentioned previously) was about 6.6 hours on average for SE, 11.5 hours for Matern 3/2 and 10.2 hours for Matern 5/2 models. For these specific simulation conditions, the runtime for HSGPs are reasonable compared to exact GPs that are computationally infeasible. Although the HSGP computation times are affected due to wider prior specifications for GP hyperparameters, the model performance in terms of latent variable estimation accuracy remains superior to VIGPs as in other simulation scenarios. In terms of speed, the VIGPs were similarly fast as in other simulation scenarios.

5 Real-world case study

We illustrate the application of the HSGPs using a real-world scRNA sequencing data. To that end, we re-analyze cell-cycle data from Mahdessian et al. (2021). Our aim is to estimate the underlying cell ordering known as "pseudotime". The dataset comprises of spliced single-cell messenger RNA (or mRNA) expression profiles of cells along the cell cycle. Each cell is additionally labeled with its own approximated time in the cell cycle process, i.e. the experimental time called "cell hours". We choose this dataset because it covers a cyclic process going through three distinct biological phases (G1, G2M and S-ph) depicting substantial variation in gene expressions. Additionally, the provided cell hours is a continuous temporal sequence that can be utilized as a prior guide to estimate

pseudotime. For the purpose of this real-world case study, we use the data from all the cells along with a selection of 12 influential genes (see supplementary materials Section I for the list) based on the recommendations of Mahdessian et al. (2021). In other words, each sample point corresponds to a single cell and each output dimension corresponds to a single gene, with the value being the gene expressions per cell. Thus, the analyzed data comprises of $N = 960$ observations and $D = 12$ output dimensions.

We use cell hours in the context of this specific data as \tilde{x} for our latent pseudotime x . Both \tilde{x} and x are real numbers with values ranging between 0 and 1. For our prior measurement SD which is unknown, we choose $s = 0.03$, keeping it proportional to our choices in the simulation studies in Section 4.1 and 4.2.

We fit a HSGP and VIGP with SE covariance function. We specify the priors for length-scale $\rho \sim \text{Normal}^+(0.4, 0.1^2)$, GP marginal SD $\alpha \sim \text{Normal}^+(14, 3.5^2)$, and error SD $\sigma \sim \text{Normal}^+(7, 3.5^2)$. The prior specifications are based on the similar real-world case study in Mukherjee et al. (2025). Based on the ρ prior and the (0, 1) range of the input space, we use the minimum number of basis function $M = 6$ (see Eq.(25)). For the VIGP, we again resort to default priors since custom priors are not supported. Additionally, based on the $N = 1000$ simulation study results in Section 4.6, the default implementation performs equally well compared to fixed hyperparameter settings. For the number of inducing points, we follow Ahmed et al. (2019) where 10 inducing points are shown to be sufficient for pseudotime estimation. Moreover, keeping the number of inducing points to 10 similar to our simulation studies allows us to avoid compromising the numerical stability of VIGPs.

As in any real-world latent variable problem, we lack knowledge of the ground truth for the latent variable. Thus, we show the deviation of the posterior pseudotime x estimates from the prior cell hours \tilde{x} . The results from our simulation studies provide evidence that this shift is indeed in the direction of the ground truth. The pseudotime deviations from cell hours for both HSGP and VIGP are presented in Fig. 8. The first two sub-figures show the estimated pseudotime deviations from the HSGP and VIGP, respectively, plotted against the prior cell hours. The third sub-figure show the difference between the estimated pseudotimes from these two models. For some cells, prior-posterior differences are on average about 10% of the total cell hours time scale for both HSGP and VIGP. Comparing their two pseudotime estimates, we see differences of about 5% (with some outlying cells).

In Fig. 9, we show the posterior estimates of the HSGP hyperparameters along with their posterior 95% CIs. A key observation here is the strong differences in the hyperparameter estimates across the genes (output dimensions). This confirms the importance of assuming gene-specific

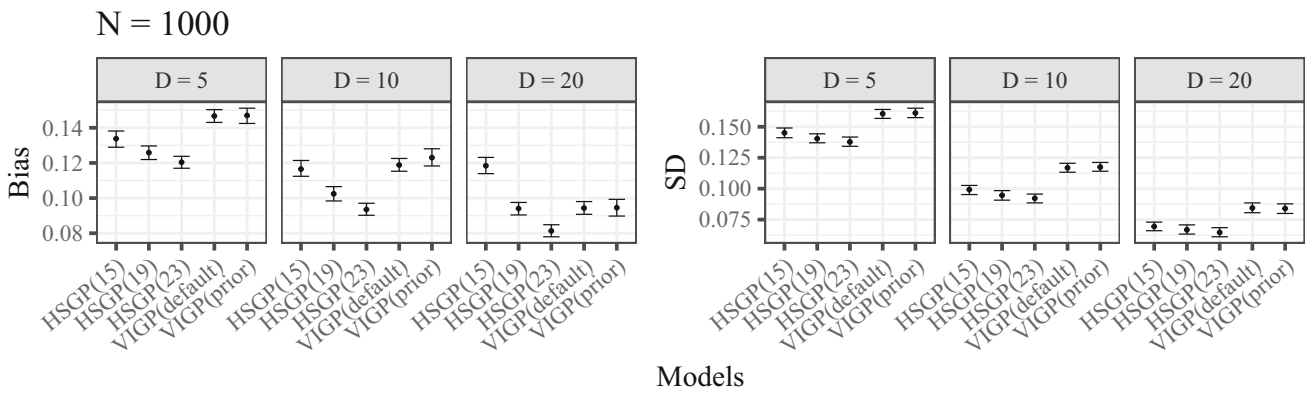


Fig. 7 Squared exponential scenario (special case): posterior bias and SD on recovery of latent inputs for all fitted models. The HSGP(M) shows the HSGPs with their corresponding number of basis functions.

VIGP (default) shows the default Pyro implementation and VIGP (prior) has fixed GP hyperparameters informed by the true sampling distributions

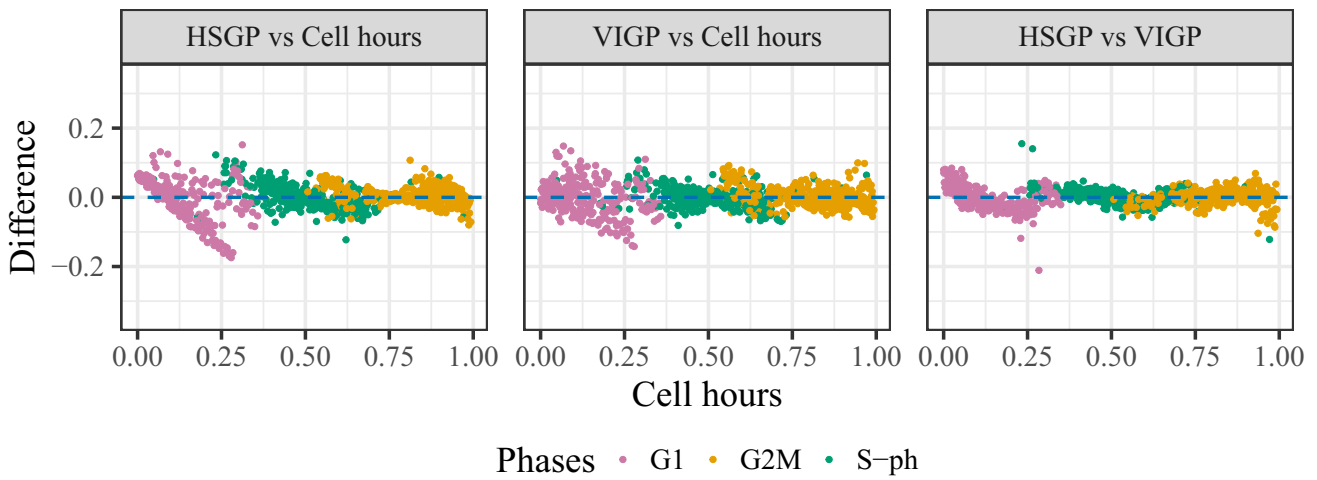


Fig. 8 (a) Difference of HSGP pseudotime estimates from the prior cell hours; (b) Difference of VIGP pseudotime estimates from the prior cell hours; (c) Difference of HSGP and VIGP pseudotime estimates. Both

models assume an SE covariance function. The different colors denote the three distinct biological phases (G1, G2M and S-ph).




Fig. 9 Hyperparameter estimates for the HSGP with SE covariance function for cell cycle data. Points indicate posterior means and the point ranges indicate 95% CIs for each hyperparameter per output dimension

hyperparameters while modeling these biological processes. Modeling varying GP hyperparameters, while supported in our HSGPs, is not possible yet for the (pyro implemented) VIGPs, so we do not include their estimates in Fig. 9.

The runtime for the HSGP was approximately 16 minutes on an Apple M4 chip (2 cores in parallel) with 24 GB working memory. Same model fitted using the cluster computing resources detailed in Section 4.2 took about 58 minutes. In comparison, fitting the VIGP took about 3 minutes using the same cluster computing resources.

6 Discussion

We generalize Hilbert space approximations for Gaussian processes (HSGPs; Solin and Särkkä 2020; Riutort-Mayol et al. 2022) to develop a scalable class of latent variable models for high-dimensional multi-output data. Our HSGPs account for both varying hyperparameters across output dimensions and correlations between output dimensions. They reduce the steep computational complexity of $O(N^3D + ND^2)$ for exact multi-output GPs to $O(NMD + ND^2)$ for a dataset of N samples and D output dimensions, where M is the number of basis functions of the HSGP. Since M is typically much smaller than N , using HSGPs reduce computation times drastically. Through extensive simulations, we investigated the statistical properties of latent GPs in terms of model convergence, uncertainty calibration, and estimation accuracy. In all of the above criteria, the HSGPs performed better than both exact GPs and approximate GPs based on inducing points and variational inference (VIGPs) (Titsias 2009; Ahmed et al. 2019). While VIGPs are faster yet, they provide limited modeling capabilities and are outperformed by HSGPs in terms of model calibration and estimation accuracy. Moreover, even with less informed hyperparameter priors (wide priors), the HSGPs achieve higher accuracy than VIGPs. Although HSGPs are comparably slower, they are reasonably fast even for relatively large datasets such as the one investigated in our real-world case study. This stands in sharp contrast to exact GPs, which would have taken at least several hours, if not more, to fit the same data.

6.1 Limitations and future research

For GPs with Matern class of covariance functions, the length-scale ρ is directly related to the scale of the input space. In latent variable GPs, this implies strong dependencies between the prior on ρ , the range of latent x , and the prior measurement SD s (see Sects. 3 and 4.2). As a result, specifying all three independently may lead to convergence issues and overall bad model behavior. Future research should investigate these dependencies more closely to pro-

vide better guidelines for prior specification in latent variable GPs.

With regard to the correlation structure C across output dimensions (see Sect. 3.3.1), we make the choice of modeling it on the level of the GP functions f rather than on the level of observations y . It is not always apparent which of these two options is preferable for a given application. For example, in spatial statistics, correlations between discrete locations are also either modeled on the mean (similar to the GP functions in our case) level, e.g., via conditional autoregressive (CAR) structures, or on the observation level, e.g., via spatial autoregressive (SAR) structures (Wall 2004). In this paper, we apply our methods to RNA gene expressions for cell cycle data, where the genes are known to be correlated at the biological process level (Riba et al. 2022). Thus, specifying across-dimension correlations on the GP functions appear sensible here. For other applications, it may be preferable to specify the correlations at the observational level, or even jointly at both levels.

Another consideration concerning C is whether to model it as input dependent or input independent. In this paper, we chose it to be input independent, thus estimating the full correlation matrix directly. As an alternative, an input dependent formulation could be achieved through another GP that models the correlations across output dimensions within each observation (similar to a cross-covariance structure shown in Gneiting et al. (2010)). In this regard, it would further be interesting to study Hilbert space approximations also for this across-output dimension GP, especially when the number of dimensions is large enough for exact GPs to become computationally impractical. We leave these extensions of our framework to future research.

Lastly, an important generalization of latent variable GPs concern the inclusion of derivative information, which can further increase estimation accuracy (Mukherjee et al. 2025). However, such methods have so far only been developed for exact GPs, which hamper their real-world applicability. Extending HSGPs to include derivative information will likely provide a scalable solution. However, such an extension requires spectral densities for composite covariance functions, which to our knowledge is yet to be derived.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11222-026-10869-x>.

Acknowledgements We acknowledge the valuable insights provided by all group members. P.B further acknowledges support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) via the Collaborative Research Center 391 (Spatio-Temporal Statistics for the Transition of Energy and Transport) – 520388526. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting S.M.

Author Contributions S.M wrote the main manuscript, prepared figures and conducted the formal analysis. S.M and P.B was involved in conceiving the idea. S.M., P.B and M.C contributed to the case study analysis. All authors reviewed the manuscript.

Data Availability Case study data was based on [Mahdessian, D., Cesnik, A.J., Gnann, C. et al. Spatiotemporal dissection of the cell cycle with single-cell proteogenomics. *Nature* 590, 649–654 (2021). <https://doi.org/10.1038/s41586-021-03232-9>] and the processed data that is used for this manuscript is made publicly available on Github (<https://github.com/Soham6298/Latent-variable-HSGPs>).

Code Availability The code for the model development, simulation studies as well as the results can be found here: <https://github.com/Soham6298/Latent-variable-HSGPs>

Declarations

Competing interests The authors declare no competing interests.

References

- Ahmed, S., Rattray, M., Boukouvalas, A.: GrandPrix: scaling up the Bayesian GPLVM for single-cell data. *Bioinformatics* **35**(1), 47–54 (2019)
- Álvarez, M.A., Lawrence, N.D.: Computationally efficient convolved multiple output gaussian processes. *J. Mach. Learn. Res.* **12**(41), 1459–1500 (2011)
- Bauer, M., van der Wilk, M., Rasmussen, C.E.: Understanding probabilistic sparse gaussian process approximations. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29. Curran Associates Inc (2016)
- Bingham, E., Chen, J.P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., Goodman, N.D.: Pyro: deep universal probabilistic programming. *J. Mach. Learn. Res.* **20**(1), 973–978 (2019)
- Bonilla, E.V., Chai, K., Williams, C.: Multi-task Gaussian Process Prediction. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc (2007)
- Bourotte, M., Allard, D., Porcu, E.: A flexible class of non-separable cross-covariance functions for multivariate space–time data. *Spatial Statistics*, 18:125–146. *Spatial Statistics Avignon: Emerging Patterns* (2016)
- Bürkner, P.-C.: brms: An R Package for Bayesian Multilevel Models Using Stan. *J. Stat. Softw.* **80**, 1–28 (2017)
- Bürkner, P.-C., Gabry, J., Kay, M., Vehtari, A.: posterior: Tools for working with posterior distributions. R package version 1.5.0 (2023)
- Campbell, K., Yau, C.: Bayesian Gaussian Process Latent Variable Models for pseudotime inference in single-cell RNA-seq data. Pages: 026872 Section: New Results (2015)
- Campbell, K.R., Yau, C.: A descriptive marker gene approach to single-cell pseudotime inference. *Bioinformatics* **35**(1), 28–35 (2018)
- Chatfield, C.: *The Analysis of Time Series*, chapter 6, pages 92–104. Chapman and Hall/CRC (2003)
- Cressie, N.A.C.: *Statistics for Spatial Data*. Wiley, New York (1993)
- Gihman, I.I., Skorokhod, A.V.: *Linear Theory of Random Processes*, pages 195–307. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
- Gneiting, T., Kleiber, W., Schlather, M.: Matérn cross-covariance functions for multivariate random fields. *J. Am. Stat. Assoc.* **105**(491), 1167–1177 (2010)
- Haque, A., Engel, J., Teichmann, S.A., Lönnberg, T.: A practical guide to single-cell RNA-sequencing for biomedical research and clinical applications. *Genome Medicine* **9**(1), 75 (2017)
- Hensman, J., Fusi, N., Lawrence, N.D.: Gaussian Processes for Big Data. [arXiv:1309.6835](https://arxiv.org/abs/1309.6835) [cs, stat] (2013a)
- Hensman, J., Lawrence, N.D., Rattray, M.: Hierarchical Bayesian modelling of gene expression time series across irregularly sampled replicates and clusters. *BMC Bioinformatics* **14**(1), 252 (2013)
- Hoffman, D.M., Gelman, A.: The no-u-turn sampler. *Journal of Machine Learning Research* (2014)
- Joukov, V., Kulic, D.: Fast Approximate Multioutput Gaussian Processes. *IEEE Intelligent Systems*, 37(4):56–69. Conference Name: IEEE Intelligent Systems (2022)
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., Blei, D.M.: Automatic differentiation variational inference. *J. Mach. Learn. Res.* **18**(14), 1–45 (2017)
- Lalchand, V., Ravuri, A., Lawrence, N.D.: Generalised GPLVM with Stochastic Variational Inference. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 7841–7864. PMLR. ISSN: 2640-3498 (2022)
- Lawrence, N.: Gaussian Process Latent Variable Models for Visualisation of High Dimensional Data. In *Advances in Neural Information Processing Systems*, volume 16. MIT Press (2003)
- Lawrence, N.: Probabilistic non-linear principal component analysis with gaussian process latent variable models. *J. Mach. Learn. Res.* **6**(60), 1783–1816 (2005)
- Lewandowski, D., Kurowicka, D., Joe, H.: Generating random correlation matrices based on vines and extended onion method. *J. Multivar. Anal.* **100**(9), 1989–2001 (2009)
- Mahdessian, D., Cesnik, A.J., Gnann, C., Danielsson, F., Stenström, L., Arif, M., Zhang, C., Le, T., Johansson, F., Schutten, R., Bäckström, A., Axelsson, U., Thul, P., Cho, N.H., Carja, O., Uhlén, M., Mardinoglu, A., Stadler, C., Lindskog, C., Ayoglu, B., Leonetti, M. D., Pontén, F., Sullivan, D.P., Lundberg, E.: Spatiotemporal dissection of the cell cycle with single-cell proteogenomics. *Nature*, 590(7847):649–654. (2021) Number: 7847 Publisher: Nature Publishing Group
- Modrák, M., Moon, A.H., Kim, S., Bürkner, P., Huurre, N., Faltejsková, K., Gelman, A., Vehtari, A.: Simulation-Based Calibration Checking for Bayesian Computation: The Choice of Test Quantities Shapes Sensitivity. [arXiv preprint](https://arxiv.org/abs/2023.08.01) (2023)
- Moreno-Muñoz, P., Artés, A., Álvarez, M.: Heterogeneous Multi-output Gaussian Process Prediction. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc (2018)
- Mukherjee, S., Claassen, M., Bürkner, P.-C.: Dgp-lvm: Derivative gaussian process latent variable models. [arXiv preprint](https://arxiv.org/abs/2025.01.01) (2025)
- Neal, R.: “MCMC Using Hamiltonian Dynamics.” In *Handbook of Markov Chain Monte Carlo*, edited by Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng, 116–62. Chapman; Hall/CRC (2011)
- Quinero-Candela, J., Rasmussen, C.E.: A unifying view of sparse approximate gaussian process regression. *J. Mach. Learn. Res.* **6**, 1939–1959 (2005)
- Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press (2006)
- Reid, J.E., Wernisch, L.: Pseudotime estimation: deconfounding single cell time series. *Bioinformatics* **32**(19), 2973–2980 (2016)
- Riba, A., Oravec, A., Durik, M., Jiménez, S., Alunni, V., Cerciat, M., Jung, M., Keime, C., Keyes, W.M., Molina, N.: Cell cycle gene regulation dynamics revealed by rna velocity and deep-learning. *Nat. Commun.* **13**(1), 2865 (2022)
- Riutort-Mayol, G., Bürkner, P.-C., Andersen, M.R., Solin, A., Vehtari, A.: Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming. *Statistics and Computing*, 33(1) (2022)

- Seeger, M.W., Williams, C.K.I., Lawrence, N.D.: Fast Forward Selection to Speed Up Sparse Gaussian Process Regression. In International Workshop on Artificial Intelligence and Statistics, pages 254–261. PMLR. ISSN: 2640-3498 (2003)
- Snelson, E., Ghahramani, Z.: Sparse Gaussian Processes using Pseudo-inputs. In Advances in Neural Information Processing Systems, volume 18. MIT Press (2005)
- Solin, A., Särkkä, S.: Hilbert space methods for reduced-rank Gaussian process regression. *Stat. Comput.* **30**(2), 419–446 (2020)
- Stan Development Team. Stan Modeling Language Users Guide and Reference Manual, 2.32.0 (2024)
- Säilynoja, T., Bürkner, P.-C., Vehtari, A.: Graphical test for discrete uniformity and its applications in goodness-of-fit evaluation and multiple sample comparison. *Stat. Comput.* **32**(2), 32 (2022)
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., Gelman, A.: Validating Bayesian Inference Algorithms with Simulation-Based Calibration. arXiv preprint (2020)
- Teh, Y.W., Seeger, M., Jordan, M.I.: Semiparametric Latent Factor Models. *Proceedings of Machine Learning Research* (2005)
- Titsias, M.: Variational Learning of Inducing Variables in Sparse Gaussian Processes. In Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, pages 567–574. PMLR. ISSN: 1938-7228 (2009)
- Titsias, M., Lawrence, N.D.: Bayesian Gaussian Process Latent Variable Model. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pages 844–851. *JMLR Workshop and Conference Proceedings*. ISSN: 1938-7228 (2010)
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., Bürkner, P.-C.: Rank-Normalization, Folding, and Localization: An Improved \hat{R} for Assessing Convergence of MCMC (with Discussion). *Bayesian Anal.* **16**(2), 667–718 (2021)
- Wall, M.M.: A close look at the spatial structure implied by the car and sar models. *Journal of statistical planning and inference* **121**(2), 311–324 (2004)
- Williams, C., Rasmussen, C.: Gaussian processes for regression. In: Touretzky, D., Mozer, M., Hasselmo, M. (eds.) *Advances in Neural Information Processing Systems*, vol. 8. MIT Press (1995)
- Yao, Y., Vehtari, A., Simpson, D., Gelman, A.: Yes, but did it work?: Evaluating variational inference. In International Conference on Machine Learning, pages 5581–5590. PMLR (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.