
Hierarchical Inducing Point Gaussian Process for Inter-domain Observations

Luhuan Wu^{*1}

Andrew Miller^{*1}

Lauren Anderson²

Geoff Pleiss¹

David Blei¹

John Cunningham¹

¹ Columbia University

² The Observatories of the Carnegie Institution for Science

{lw2827, gmp2162, david.blei, jpc2181}@columbia.edu

{andrew.colin.miller, anders.astro}@gmail.com

Abstract

We examine the general problem of inter-domain Gaussian Processes (GPs): problems where the GP realization and the noisy observations of that realization lie on different domains. When the mapping between those domains is linear, such as integration or differentiation, inference is still closed form. However, many of the scaling and approximation techniques that our community has developed do not apply to this setting. In this work, we introduce the *hierarchical inducing point GP* (HIP-GP), a scalable inter-domain GP inference method that enables us to improve the approximation accuracy by increasing the number of inducing points to the millions. HIP-GP, which relies on inducing points with grid structure and a stationary kernel assumption, is suitable for low-dimensional problems. In developing HIP-GP, we introduce (1) a fast whitening strategy, and (2) a novel preconditioner for conjugate gradients which can be helpful in general GP settings.

1 INTRODUCTION

Gaussian processes (GPs) are a widely used statistical tool for inferring unobserved functions (Cressie, 1992, 1990; Rasmussen and Williams, 2006). The classic goal of GPs is to infer the unknown function given noisy observations. Here, we are interested in a more general

setting, inter-domain GPs, where the observed data is related to the latent function via some linear transformation, such as integration or differentiation, while an identity transformation recovers the standard GP problem. One motivating example is an astrophysics problem: mapping the three-dimensional spatial distribution of dust in the Milky way (Green et al., 2015; Leike and Enßlin, 2019; Kh et al., 2017). Interstellar dust is a latent function that can be inferred from star observations. However, because we are embedded in our own dust field, we can only observe some noisy *integral* of the dust function along the line of sight between Earth and a star. Beyond this example, inter-domain GPs arise often in the literature: integrated observations have been used in probabilistic construction of optimization algorithms (Wills and Schön, 2017), quadrature rules (Minka, 2000), and tomographic reconstructions (Jidling et al., 2018); while derivative observations have been used in dynamical systems (Solak et al., 2003), modeling monotonic functions (Riihimäki and Vehtari, 2010) and Bayesian optimization (Garnett et al., 2010; Siivola et al., 2018).

In practice, this type of inter-domain GP problem poses two intertwined obstacles that are beyond the reach of current techniques. First, large-scale exact modeling is usually intractable. The joint distribution of inter-domain observations and the underlying GP involves the transformed-domain and inter-domain kernel expressions, which rarely admits analytical solutions and requires approximations (Lázaro-Gredilla and Figueiras-Vidal, 2009; Hendriks et al., 2018). Common approximations are often handled by numerical integration, which is infeasible for big datasets since it requires integrating all pairwise correlations.

Moreover, inter-domain GPs suffer from the same scalability issues as regular GPs. For a dataset with N observations, the likelihood function depends on N^2 pairwise correlations. The leading strategy to scale

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s). * Authors contributed equally.

standard GP inference is to use $M \ll N$ *inducing points* to represent the global behavior of GP functions (Rasmussen and Williams, 2006). One popular inducing point method is stochastic variational Gaussian process (SVGP), which factorizes the objective over mini-batches of data and requires only $O(M^2)$ storage and $O(M^3)$ computation (Hensman et al., 2013). In the current practice of SVGP, M is limited to under 10,000 (Wilson and Nickisch, 2015; Izmailov et al., 2018). However, many inter-domain problems are spatial or temporal in nature, and the data do not lie in some small manifold in that space. In the interstellar dust problem, for example, we aim to make inference at every point in a dense 3D space. A small set of inducing points is incapable of resolving the resolution of interest, which is around 4 orders of magnitude smaller than the domain size. Furthermore, Bauer et al. (2016) shows that more inducing points are needed to reduce the overestimated observation noise parameter induced by SVGP. All of these facts necessitate the need to scale both N and M to larger quantities.

To this end, we develop the hierarchical inducing point GP (HIP-GP), a method to scale GP inference to millions of inducing points and observations for spatial-temporal inter-domain problems. In particular,

- We adapt the SVGP framework to inter-domain settings by decoupling observations and inducing points into different domains. This framework alleviates the difficulties of computing the full transformed kernel matrices, and enables the exploitations of the latent kernel structure.
- We then develop the HIP-GP algorithm to address the computational bottlenecks of standard SVGP objectives, employing two core strategies:
 - Fast matrix inversion with conjugate gradient method using the *hierarchical Toeplitz structure*. Upon this structure, we design a novel *preconditioner* and a new *whitening strategy* to further speed up computations;
 - A *structured variational approximation* of the posterior over inducing point values.

HIP-GP is suitable for low-dimensional inter-domain GP problems, and applies in settings where the kernel function is stationary and inducing points fall on a fixed, evenly-spaced grid. In addition, the technical innovations in developing HIP-GP are useful in a variety of more general settings.

2 BACKGROUND

2.1 Inter-domain GPs

Following the notations in van der Wilk et al. (2020), we consider a statistical model of the form

$$\rho \sim GP(0, k_\theta(\cdot, \cdot)) \quad (1)$$

$$\rho^* = \mathcal{L} \circ \rho \quad (2)$$

$$y_n | \mathbf{x}_n, \rho^* \sim \mathcal{N}(\rho^*(\mathbf{x}_n), \sigma_n^2) \quad (3)$$

for a dataset of N observations $\mathcal{D} \triangleq \{y_n, \mathbf{x}_n, \sigma_n^2\}_{n=1}^N$, where \mathcal{L} is a linear operator and $k_\theta(\cdot, \cdot)$ is the covariance function that encodes prior assumptions about the function ρ . Note that GPs are closed under linear operators, therefore ρ^* is also a GP (Rasmussen and Williams, 2006).

One common linear operator is the integral operator, $\mathcal{L} \circ \rho(\cdot) = \int \rho(\mathbf{x}) w(\mathbf{x}) d\mathbf{x}$, as used in Lázaro-Gredilla and Figueiras-Vidal (2009). We see that this \mathcal{L} maps the entire function $\rho(\cdot)$ to a single real value. Another example is the derivative of the d th input dimension $\mathcal{L} \circ \rho(\cdot) = \frac{\partial \rho}{\partial x_d}(\mathbf{x}_n)$. In this case, the operator only depends on the neighborhood around \mathbf{x}_n . Derivative observations are often useful for algorithmic purposes, e.g. in Riihimäki and Vehtari (2010). In application problems, they could be either collected, e.g. velocity measured by physical detectors, or identified from function observations (Solak et al., 2003). We also notice that setting \mathcal{L} to an identity map fits regular GPs into this framework.

The goal of inter-domain GPs is to infer the underlying function $\rho(\mathbf{x})$ — either to compute $p(\rho(\mathbf{x}_*) | \mathcal{D})$ for new test locations \mathbf{x}_* or to improve estimates of $p(\rho(\mathbf{x}_n) | \mathcal{D})$ for observed location \mathbf{x}_n given *all* observations.

2.2 Stochastic Variational Gaussian Process

The stochastic variational Gaussian process (SVGP) is an approximate method that scales GP inference to large N (Hensman et al., 2013). Denote the M inducing point *locations* $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)$, and the vector of inducing point *values* $\mathbf{u} \triangleq (\rho(\bar{\mathbf{x}}_1), \dots, \rho(\bar{\mathbf{x}}_M))$. SVGP defines a variational distribution over the inducing point values \mathbf{u} and the latent process values $\rho \triangleq (\rho(\mathbf{x}_1), \dots, \rho(\mathbf{x}_N))$ of the form

$$q(\mathbf{u}, \rho) = q_\lambda(\mathbf{u}) p(\rho | \mathbf{u}), \quad q_\lambda(\mathbf{u}) = \mathcal{N}(\mathbf{u} | \mathbf{m}, \mathbf{S}),$$

where $q_\lambda(\mathbf{u})$ is a multivariate Gaussian, $p(\rho | \mathbf{u})$ is determined by the GP prior and $\lambda \triangleq (\mathbf{m}, \mathbf{S})$ are variational parameters. This choice of variational family induces a convenient cancellation, resulting in a separa-

rable objective (Titsias, 2009)

$$\begin{aligned} \mathcal{L}(\lambda) & \quad (4) \\ &= \underbrace{\mathbb{E}_{q_\lambda(\mathbf{u})} [\mathbb{E}_{p(\rho|\mathbf{u})} [\ln p(\mathbf{y}|\rho)]]}_{\text{(a)}} - \underbrace{KL(q_\lambda(\mathbf{u}) || p(\mathbf{u}))}_{\text{(b)}}. \end{aligned}$$

We can write (a) as a sum over N observations

$$\text{(a)} = \sum_{n=1}^N \underbrace{\mathbb{E}_{q_\lambda(\mathbf{u})} [\mathbb{E}_{p(\rho_n|\mathbf{u})} [\ln p(y_n|\rho_n)]]}_{\triangleq a_n}. \quad (5)$$

The factorization of (a) enables the objectives to be estimated with mini-batches in a large dataset. However, notice that (b), the KL-divergence of two Gaussians, will involve a term $\ln |\mathbf{K}_{\mathbf{u},\mathbf{u}}|$ which requires $O(M^3)$ computation.

2.3 Matrix Solves with Conjugate Gradients

Conjugate gradients (CG) is an iterative algorithm for solving a linear system using only matrix-vector multiplies (MVM). CG computes $\mathbf{K}^{-1}\mathbf{p}$ for any $\mathbf{p} \in \mathbb{R}^M$ by computing $\mathbf{K}\mathbf{v}$ for a sequence of vectors $\mathbf{v} \in \mathbb{R}^M$ determined by the algorithm. For \mathbf{K} of size $M \times M$, CG computes the exact solution after M iterations, and typically converges after some smaller number of steps $S < M$ (Hestenes and Stiefel, 1952; Nocedal and Wright, 2006).

Preconditioned conjugate gradients (PCG) is an augmented version of CG that solves the system in a transformed space. A good preconditioner can dramatically speed up convergence (Shewchuk et al., 1994; Cutajar et al., 2016).

3 SCALING M : HIP-GP for INTER-DOMAIN PROBLEMS

We first formulate the SVGP framework for inter-domain observations, and identify its computational bottlenecks in Section 3.1. We then address these bottlenecks by the HIP-GP algorithm using the techniques developed in Section 3.2 - 3.3. In Section 3.4, we summarize our methods and discuss optimization procedures for HIP-GP.

3.1 Inter-domain SVGP Formulation

We show that the inter-domain observations can be easily incorporated into the SVGP framework. We place a set of inducing points $\mathbf{u} = \rho(\bar{\mathbf{x}})$ in the latent domain at input locations $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_M)$. Connections to the observations are made through the inter-

domain covariance, while the observations are characterized by the transformed-domain covariance. Formally, we have the inter-domain GP prior:

$$\begin{pmatrix} \rho_n^* \\ \mathbf{u} \end{pmatrix} \sim \mathcal{N} \left(0, \begin{pmatrix} k_{n,n}^{**} & \mathbf{k}_{n,\mathbf{u}}^* \\ \mathbf{k}_{\mathbf{u},n}^{**} & \mathbf{K}_{\mathbf{u},\mathbf{u}} \end{pmatrix} \right), \quad (6)$$

where the inter-domain covariance and the transformed-domain covariance are defined as

$$\mathbf{k}_{\mathbf{u},n}^* \triangleq \text{Cov}(\rho(\bar{\mathbf{x}}), \rho^*(\mathbf{x}_n)) = \text{Cov}(\mathbf{u}, \rho_n^*), \quad (7)$$

$$k_{n,n}^{**} \triangleq \text{Cov}(\rho^*(\mathbf{x}_n), \rho^*(\mathbf{x}_n)) = \text{Cov}(\rho_n^*, \rho_n^*), \quad (8)$$

and the latent domain covariance is

$$\mathbf{K}_{\mathbf{u},\mathbf{u}} \triangleq \text{Cov}(\rho(\bar{\mathbf{x}}), \rho(\bar{\mathbf{x}})) = \text{Cov}(\mathbf{u}, \mathbf{u}). \quad (9)$$

This form of the prior suggests formulating the inter-domain SVGP objective as follows

$$\begin{aligned} \mathcal{L}(\lambda) &= \sum_{n=1}^N \mathbb{E}_{q_\lambda(\mathbf{u})} [\mathbb{E}_{p(\rho_n^*|\mathbf{u})} [\ln p(y_n|\rho_n^*)]] \\ &\quad - KL(q_\lambda(\mathbf{u}) || p(\mathbf{u})), \end{aligned} \quad (10)$$

where

$$p(\rho_n^*|\mathbf{u}) = N(\rho_n^* | \mathbf{k}_{n,\mathbf{u}}^* \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{u}, k_{n,n}^{**} - \mathbf{k}_{n,\mathbf{u}}^* \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*).$$

Note that this framework can be extended to observations in multiple domains by including them with their corresponding inter-domain and transformed-domain covariances. Under this formulation, we avoid computing the $N \times N$ transformed-domain covariance matrix $\mathbf{K}_{N,N}^{**}$ that appears in the exact GP objective. Instead, only N terms of variance $k_{n,n}^{**}$ need to be evaluated. Importantly, the disentanglement of observed and latent domains enables us to exploit structure of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ for efficient computations. Such exploitation would be difficult without a variational approximation, especially in the case of mixed observations from multiple domains.

Whitened Parameterization Whitened parameterizations are used to improve inference in models with correlated priors because they offer a better-conditioned posterior (Murray and Adams, 2010; Hensman et al., 2015). Here we will show an additional computational benefit in the variational setting — the whitened posterior allows us to avoid computing $\ln |\mathbf{K}_{\mathbf{u},\mathbf{u}}|$ which appears in the KL term in Equation 10. To define the whitened parameterization, we describe the GP prior over \mathbf{u} as a deterministic function of standard normal parameters $\boldsymbol{\epsilon}$:

$$\boldsymbol{\epsilon} \sim \mathcal{N}(0, I), \quad \mathbf{u} = \mathbf{R}\boldsymbol{\epsilon}. \quad (11)$$

To preserve the covariance structure in the prior distribution of (ρ_n^*, \mathbf{u}) (Equation 6), the transformation

\mathbf{R} and the whitened correlation $\mathbf{k}_n \triangleq \text{Cov}(\boldsymbol{\epsilon}, \rho_n^*)$ need to satisfy the following two equalities:

$$\begin{aligned} \mathbf{K}_{\mathbf{u},\mathbf{u}} &= \text{Cov}(\mathbf{R}\boldsymbol{\epsilon}, \mathbf{R}\boldsymbol{\epsilon}) = \mathbf{R}\mathbf{R}^\top, \\ \mathbf{k}_{\mathbf{u},n}^* &= \text{Cov}(\mathbf{R}\boldsymbol{\epsilon}, \rho_n^*) = \mathbf{R}\mathbf{k}_n. \end{aligned} \quad (12)$$

The classical whitening strategy in GP inference is to use the Cholesky decomposition: $\mathbf{K}_{\mathbf{u},\mathbf{u}} = \mathbf{L}\mathbf{L}^\top$ where \mathbf{L} is a lower triangular matrix. In this case, $\mathbf{R} = \mathbf{L}$ and $\mathbf{k}_n = \mathbf{L}^{-1}\mathbf{k}_{\mathbf{u},n}^*$.

Now we can target the variational posterior over the whitened parameters $\boldsymbol{\epsilon}$: $q_\lambda(\boldsymbol{\epsilon}) = \mathcal{N}(\boldsymbol{\epsilon} | \mathbf{m}, \mathbf{S})$. The resulting *whitened variational objective* is

$$\mathcal{L}(\boldsymbol{\lambda}) = \sum_n \underbrace{\mathbb{E}_{q_\lambda(\boldsymbol{\epsilon})p(\rho_n^* | \boldsymbol{\epsilon})} [\ln p(y_n | \rho_n^*)]}_{(a'_n)} - \underbrace{KL(q_\lambda(\boldsymbol{\epsilon}) || p(\boldsymbol{\epsilon}))}_{(b')} \quad (13)$$

where

$$\begin{aligned} (a'_n) &= -\frac{1}{2} \ln \sigma_n^2 - \frac{1}{2\sigma_n^2} \left(y_n^2 + k_{n,n}^{**} - \mathbf{k}_n^\top \mathbf{k}_n \right. \\ &\quad \left. + \mathbf{k}_n^\top (\mathbf{S} + \mathbf{m}\mathbf{m}^\top) \mathbf{k}_n - 2y_n \mathbf{k}_n^\top \mathbf{m} \right), \\ (b') &= \frac{1}{2} (\text{tr}(\mathbf{S}) + \mathbf{m}^\top \mathbf{m} - \ln |\mathbf{S}| - M). \end{aligned} \quad (14)$$

Computational Bottlenecks The whitened objective above still factorizes over data points. However, there remain two computational bottlenecks. First, the correlation term \mathbf{k}_n in (a'_n) depends on the choice of the whitening strategy. The common Cholesky strategy requires $O(M^3)$ computation and $O(M^2)$ storage which is infeasible for large M . We address this bottleneck in Section 3.2. The second bottleneck lies in the variational covariance \mathbf{S} which is an $M \times M$ matrix, requiring $O(M^2)$ to store and $O(M^3)$ to compute the $\ln |\mathbf{S}|$ in (b') . We will address this problem by a structured variational approximation in Section 3.3.

3.2 Computational Accelerations

We now turn to the first bottleneck — how to design an efficient whitening strategy to compute the term \mathbf{k}_n . To do so, we rely on judicious placement of inducing points and assume a stationary covariance function, a general and commonly used class. We describe three key ingredients below.

Hierarchical Toeplitz Structure Consider a D -dimensional grid of evenly spaced points of size $M \triangleq M_1 \times \dots \times M_D$, characterized by one-dimensional grids of size M_i along dimension $i, i = 1 : D$, where D is the input dimension. Under a stationary kernel, we construct a covariance matrix for this set of points in x -major order (i.e. C-order). Such a matrix will have

hierarchical Toeplitz structure, which means the diagonals of the matrix are constant. Because of this data redundancy, a hierarchical Toeplitz matrix is characterized by its first row. Now we place the inducing points along a fixed, equally-spaced grid, resulting in a $M \times M$ hierarchical Toeplitz Gram matrix $\mathbf{K}_{\mathbf{u},\mathbf{u}}$. The efficient manipulation of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ is through its circulant embedding:

$$\mathbf{C} = \begin{pmatrix} \mathbf{K}_{\mathbf{u},\mathbf{u}} & \tilde{\mathbf{K}} \\ \tilde{\mathbf{K}}^\top & \mathbf{K}_{\mathbf{u},\mathbf{u}} \end{pmatrix} \quad (16)$$

where $\tilde{\mathbf{K}}$ is the appropriate reversal of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ to make \mathbf{C} circulant. \mathbf{C} admits a convenient diagonalization

$$\mathbf{C} = \mathbf{F}^\top \mathbf{D} \mathbf{F} = \mathbf{F}^\top \text{diag}(\mathbf{F}\mathbf{c}) \mathbf{F}, \quad (17)$$

where \mathbf{F} is the fast Fourier transform matrix, \mathbf{D} is a diagonal matrix of \mathbf{C} 's eigenvalues, and \mathbf{c} is the first row of \mathbf{C} . This diagonalization enables fast MVMs with \mathbf{C} and hence the embedded $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ via the FFT algorithm in $O(M \ln M)$ time, further making it sufficient for use within CG to efficiently solve a linear system.

The fast solves afforded by Toeplitz structure have been previously utilized for exact GP inference (Cunningham et al., 2008; Wilson et al., 2015). Here, we extend the applicability of Toeplitz structure to the variational inter-domain case by introducing a fast whitening procedure and an effective preconditioner for CG.

Fast Whitening Strategy Similar to the Cholesky decomposition, we aim to find a whitened matrix \mathbf{R} that serves as a root of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$, i.e. $\mathbf{R}\mathbf{R}^\top = \mathbf{K}_{\mathbf{u},\mathbf{u}}$. Directly solving $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{1/2}$ is not trivial. Alternatively, we access the root from the circulant embedding of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$. We consider the square root of the circulant matrix

$$\mathbf{C}^{1/2} = \mathbf{F}^\top \mathbf{D}^{1/2} \mathbf{F}, \quad (18)$$

and its block representation

$$\mathbf{C}^{1/2} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{D} \end{pmatrix}. \quad (19)$$

We make a key observation that the first row-block (\mathbf{A}, \mathbf{B}) can be viewed as a “rectangular root” of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$. That is, we define a *non-square* whitening matrix \mathbf{R} and the correlation vector \mathbf{k}_n as follows

$$\mathbf{R} \triangleq (\mathbf{A} \ \mathbf{B}), \quad \mathbf{k}_n \triangleq \mathbf{R}^\top \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*. \quad (20)$$

One can verify such \mathbf{R} and \mathbf{k}_n satisfy Equation 12, thus offering a valid whitening strategy. We note that since \mathbf{R} is non-square and $\mathbf{u} = \mathbf{R}\boldsymbol{\epsilon}$, this strategy doubles the number of variational parameters in each dimension of the whitened space.

Now we address how to efficiently compute \mathbf{k}_n defined in Equation 20. We first compute the intermediate quantity $\mathbf{k}'_n = \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*$ via CG in $O(M \ln M)$ time. We then compute $\mathbf{k}_n = \mathbf{R}^\top \mathbf{k}'_n$. Note that \mathbf{R}^\top is embedded in the matrix $\mathbf{C}^{1/2}$ which also admits the FFT diagonalization (Equation 18). Hence, MVM with \mathbf{R}^\top can be also done in $O(M \ln M)$ time.

Lastly, we show how to make CG's computation of $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*$ faster with a well-structured preconditioner.

Efficient Preconditioner The ideal preconditioner \mathbf{P} is a matrix that whitens the matrix to be inverted — the ideal \mathbf{P} is $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1}$. However, we cannot efficiently compute $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1}$. But due to the convenient diagonalization of the circulant embedding matrix \mathbf{C} , we can efficiently compute the inverse of \mathbf{C} :

$$\mathbf{C}^{-1} = \mathbf{F}^\top \mathbf{D}^{-1} \mathbf{F}. \quad (21)$$

Note that the upper left block of \mathbf{C}^{-1} *does not* correspond to $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1}$ as we explicitly write out

$$\mathbf{C}^{-1} = \begin{pmatrix} \left(\mathbf{K}_{\mathbf{u},\mathbf{u}} - \tilde{\mathbf{K}} \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \tilde{\mathbf{K}}^\top \right)^{-1} & \cdots \\ \cdots & \cdots \end{pmatrix}. \quad (22)$$

However, when the number of inducing points are large enough, $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ approaches a banded matrix, and so $\tilde{\mathbf{K}}$ is increasingly sparse. Therefore, the upper left block of \mathbf{C}^{-1} would be close to $\mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1}$, suggesting that it can serve as an effective *preconditioner* within PCG, and therefore an effective strategy for solving a linear system with the kernel matrix. We note that this banded property is often exploited in developing effective preconditioners (Chan and Ng, 1996; Saad, 2003). To justify this intuition, we analyze the PCG convergence speed under various settings of kernel functions and inducing point densities in appendix. We compare the performance of PCG and CG in systems of varying size in Section 5.1. We find that PCG converges faster than CG across all systems, taking only a fraction of the number of iterations that standard CG requires to converge. This speedup is crucial — PCG is a subroutine we use to compute the gradient term corresponding to each observation n .

Summary of Fast Computation for \mathbf{k}_n To summarize, we exploit additional computational benefits of the hierarchical Toeplitz matrix through its circulant embedding matrix, which enables fast matrix square-root and matrix inverse. We further utilize these fast operations to design novel whitening and preconditioning strategies. Thus, the whitened correlation term $\mathbf{k}_n = \mathbf{R}^\top \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*$ can be efficiently processed as follows:

1. embed $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ into a larger circulant matrix \mathbf{C} ;

2. solve $\mathbf{K}_{\mathbf{u},\mathbf{u}} \mathbf{k}'_n = \mathbf{k}_{\mathbf{u},n}^*$ for the intermediate term \mathbf{k}'_n with PCG, where we utilize the FFT diagonalization of \mathbf{C} and \mathbf{C}^{-1} ;
3. compute $\mathbf{k}_n = \mathbf{R}^\top \mathbf{k}'_n$, where we utilize the FFT diagonalization of $\mathbf{C}^{1/2}$.

The space and time complexity of this procedure are $O(M)$ and $O(M \ln M)$. This offers a speed-up over the Cholesky decomposition which has $O(M^2)$ space and $O(M^3)$ time complexity, respectively. In Section 5.2, we examine this acceleration by comparing the time of computing \mathbf{k}_n using Cholesky and using HIP-GP, as the system size M varying from 10^3 to 10^6 . HIP-GP's strategy outperforms Cholesky for small values of M , and scales to larger M where Cholesky is no longer feasible. We present HIP-GP's algorithmic details in appendix.

3.3 Structured Variational Approximation

Finally, we turn to the second bottleneck: how to represent and manipulate variational parameters of mean \mathbf{m} and covariance \mathbf{S} . We propose the *block independent* variational family

$$q(\mathbf{u}) = \prod_b^B \mathcal{N}(\mathbf{u}_b | \mathbf{m}_b, \mathbf{S}_b), \quad (23)$$

where \mathbf{u}_b denotes a subset of inducing points of size $M_b < M$ and \mathbf{S}_b is the $M_b \times M_b$ variational covariance for that subset. Note that when $M_b = 1$, it reduces to the *mean-field* variational family, and when $M_b = M$, it is the *full-rank* variational family. Calculations of the inverse and log-determinant of block independent \mathbf{S} scale $O(BM_b^3)$ — we must choose M_b to be small enough to be practical.

We note that independence in the posterior is a more reasonable approximation constraint *in the whitened parameterization* than the original space. The original GP prior, $p(\mathbf{u})$, is designed to have high correlation, and therefore data are unlikely to decorrelate inducing point values. In the whitened space, on the other hand, the prior is already uncorrelated. Hence the whitened posterior is not spatially correlated as much as the original posterior. This is in addition to the benefits of optimizing in the whitened space due to better conditioning.

Constructing Blocks The block independent approximation of Equation 23 requires assigning inducing points to B blocks. Intuitively, blocks should include nearby points, and so we focus on blocks of points that tile the space. To reconcile the Toeplitz ordering and the block orderings (they may not be the

same), we simply have to permute any M -length vector (e.g. $\mathbf{k}_{\mathbf{u},n}^*$ or \mathbf{m}) before multiplication with \mathbf{S} and then undo the permutation after multiplication. Fortunately, this permutation is linear in M .

3.4 Method Summary

The modeling difficulty of inter-domain GP problems arises from the numerical intractability of computing the full transformed-domain covariance $\mathbf{K}_{N,N}^{**}$ of size $N \times N$. We avoid this difficulty by decoupling the observations and the inducing points into different domains under the SVGP framework. Moreover, we leverage the kernel structure of the Gram matrix $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ in the latent domain for efficient computations.

The computational difficulty stems from the computations with the kernel matrix $\mathbf{K}_{\mathbf{u},\mathbf{u}}$ and the variational covariance \mathbf{S} . We avoid having to compute $\ln|\mathbf{K}_{\mathbf{u},\mathbf{u}}|$ by using a *whitened parameterization*; we develop a *fast whitening strategy* to compute the whitened correlation term $\mathbf{k}_n = \mathbf{R}^\top \mathbf{K}_{\mathbf{u},\mathbf{u}}^{-1} \mathbf{k}_{\mathbf{u},n}^*$ by exploiting the *hierarchical Toeplitz structure* with a *novel preconditioner*; and finally we explore a *structured representation* for \mathbf{S} .

Optimization We perform natural gradient descent on variational parameters using closed-form gradient updates. For gradient-based learning of kernel hyperparameters, automatically differentiating through the CG procedure is not numerically stable. Fortunately, we can efficiently compute the analytical gradient of CG solves utilizing the hierarchical Toeplitz structure, without increasing the computational complexity. See appendix for more details on gradient derivations.

4 RELATED WORK

Inter-domain GPs The idea of the inter-domain Gaussian processes has been discussed in (Lázaro-Gredilla and Figueiras-Vidal, 2009; van der Wilk et al., 2020). However, their primary interests are using inter-domain transformations to define inducing variables for specifying GP approximations, whereas our work explores the usage of SVGP framework to perform scalable modeling and inference with inter-domain observations.

Scalable Inducing Point Methods We note several recent approaches to scaling the number of inducing points in GP approximations. Shi et al. (2020) takes an orthogonal strategy to ours by approximating GP with inducing points in two independent directions, whereas HIP-GP requires inducing points to densely cover the input space. However, while improved over standard SVGP, their method still remains a cubic complexity. Izmailov et al. (2018) introduces the ten-

sor train decomposition into the variational approximation. Alternatively, Evans and Nair (2018) directly approximate the kernel with a finite number of eigenfunctions evaluated on a dense grid of inducing points. Both methods rely on *separable* covariance kernels to utilize the Kronecker product structure. This limits the class of usable kernels. The Matérn kernel, for example, is not separable across dimensions. To fill that gap, we instead focus on the class of *stationary kernels*.

Another line of inducing point work is based on *sparse kernel interpolations*. KISS-GP uses a local kernel interpolation of inducing points to reduce both the space and time complexity to $O(N + M^2)$ (Wilson and Nickisch, 2015). SV-DKL also uses local kernel interpolation, and exploits separable covariance structures and deep learning techniques to address the problem of multi-output classification (Wilson et al., 2016). But these kernel interpolation methods are not applicable to inter-domain observations under transformations. More specifically, for standard (non-inter-domain) problems, kernel interpolation methods approximate the $N \times N$ covariance matrix $\mathbf{K}_{\mathbf{N},\mathbf{N}}$ with $\mathbf{W}\mathbf{K}_{\mathbf{u},\mathbf{u}}\mathbf{W}^\top$, where \mathbf{W} is a sparse interpolation weight matrix. However, for problems with integral observations, we must compute the integrated kernel $\mathbf{K}_{\mathbf{N},\mathbf{N}}^{**} = [\int \int \text{Cov}(\rho(\mathbf{x}_i), \rho(\mathbf{x}_j)) d\mathbf{x}_i d\mathbf{x}_j]_{i,j=1}^N$. Approximating this integral with local interpolation is not straightforward, and computing every integrated cross-covariance term is costly. Alternatively, HIP-GP decouples observations and inducing points into different domains through the *inter-domain prior* (Equation 6). This decoupled prior enables mini-batch processing of $k_{n,n}^{**}$, eliminates the need to compute cross-covariance terms k_{n_i,n_j}^{**} , while still maintaining structure exploitation of $\mathbf{K}_{\mathbf{u},\mathbf{u}}$.

Fast Whitening Strategy As is mentioned before, the classical whitening strategy is the Cholesky decomposition that has $O(M^2)$ space and $O(M^3)$ time complexity. Pleiss et al. (2020) provides a more general purpose method for fast matrix roots and is in particular applicable to whitening GP. Their method is an MVM-based approach that leverages the contour integral quadrature and requires $O(M \log M + QM)$ time for Q quadrature points. Our whitening strategy specifically targets gridded inducing points and achieves more complexity savings ($O(M \log M)$ time).

5 EXPERIMENTS

Due to the mismatch in missions of different scalable GP methods, we focus most of our empirical study on HIP-GP and SVGP (with Cholesky whitening) which serve the most similar purposes. We also include in appendix a standard GP problem on a UCI benchmark




Figure 1: Convergence result of PCG v.s. CG. We compare PCG to standard CG for systems of size $M = 625, 2,500$ and $10,000$ over 25 independent runs. We report RMSE as a function of the fraction of *total CG iterations* (to converge). PCG converges faster than CG, and for larger M fewer iterations are required.

dataset (Dua and Graff, 2017) where we compare HIP-GP to exact GP (Wang et al., 2019), sparse Gaussian Process regression (SGPR) (Titsias, 2009) and SVGP.

5.1 Effect of the Preconditioner

We first examine the effect of the preconditioner developed in Section 3.2. We run CG and PCG with the preconditioner for systems of size $M = 625$ (25×25), $M = 2,500$ (50×50), and $M = 10,000$ (100×100) determined by a two-dimensional grid applied to the Matérn kernel. We run the algorithm to convergence (at tolerance $1e-10$) for 25 randomly initialized vectors of size M . We record the error at each iteration — the norm of the distance between the current solution and the converged solution.

We report the RMSE at each iteration in Figure 1. We rescale the x -axis to run from 0 to 1 for each system of size M . From this experiment we see two results: the Toeplitz preconditioner is extremely effective and the preconditioner seems to be more effective as the system size *becomes larger*. The fraction of CG iterations required for PCG to converge for $M = 10,000$ ($< 4.5\%$) is much smaller than the fraction of iterations required for $M = 625$ ($< 18\%$) to converge. Without this preconditioner, we would expect each HIP-GP iteration to take over twenty times longer to achieve similar precision.

5.2 Speedup over Cholesky Decomposition

We examine the speedup of HIP-GP’s whitening strategy over the Cholesky whitening strategy in standard SVGP, by comparing the time for solving the correlation term \mathbf{k}_n . We generate 200 random 1D observations, and evenly-spaced inducing grids of size M

M	10^3	10^4	10^5	10^6
HIP-GP	0.0045	0.0185	0.3475	1.4595
SVGP	0.0175	0.1745	n/a	n/a

Table 1: Whitening time comparison (second) of HIP-GP v.s. SVGP with Matérn(2.5) kernel.


ranging from 10^3 to 10^6 . We apply a set of kernels including the Matérn kernels with $\nu = 0.5, 1.5, 2.5$ and the squared exponential kernel. The marginal variance is fixed to 0.1 for all M . The length scale is set to L/M where L is the range of the data domain to utilize the inducing points efficiently. The PCG within the HIP-GP algorithm is run to convergence at tolerance $1e-10$. The Cholesky decomposition is only available up to $M = 10^4$ due to the memory limit. All experiments are run on a NVIDIA Tesla V100 GPU with 32GB memory.

We report the wall clock time of computations applied to Matérn (2.5) kernel in Table 1. The full report for all settings is presented in appendix. HIP-GP’s whitening strategy is consistently faster than the Cholesky whitening strategy across all experiments, and scales to larger M .


5.3 Synthetic Derivative Observations

To validate our inter-domain SVGP framework, we study a derivative GP problem. We follow the work in Solak et al. (2003), which introduces derivative observations in addition to regular function observations to reduce uncertainty in learning dynamic systems. We synthesize a 1D GP function from a random neural network with sinusoidal non-linearities, and obtain function derivatives using automatic differentiation. The total observations consist of 100 function observations and 20 derivative observations, with added noise level = 0.05 and 0.2 respectively, as depicted in Figure 2a.

We compare two inter-domain SVGP framework-based methods, HIP-GP and the standard SVGP, to the exact GP. We use the squared exponential kernel with signal variance 0.5 and length scale 0.1. For both HIP-GP and SVGP, we apply the full-rank variational family. The maximum number of PCG iterations within HIP-GP is set to 20. We evaluate the predictive performance on 100 test data. From Figure 2b and 2c, we see that the inter-domain SVGP framework successfully utilizes the derivative observations to improve the prediction quality with reduced uncertainty, and is comparable to the exact method.



(a) Synthetic function and derivative observations



(b) HIP-GP posterior prediction with / without derivative observations for 100 test data. The solid blue line is the mean prediction and the shaded blue area is the 1 posterior standard deviation band. The red points indicate derivative observation locations with true function values.

	HIP-GP	SVGP	Exact GP
RMSE	0.0192	0.0192	0.0192
Uncertainty	0.0198	0.0206	0.0198

(c) Predictive RMSE and uncertainty (i.e. average standard deviation) for 100 test data.

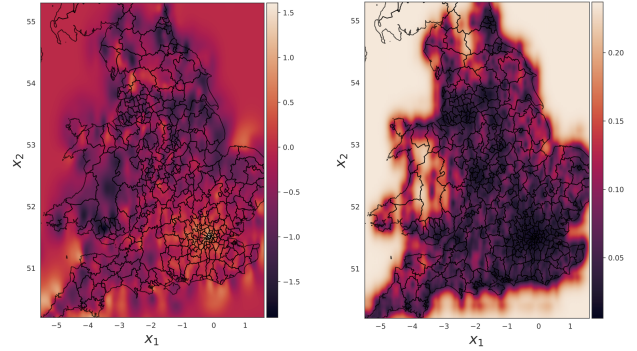
Figure 2: GP with derivative observations analysis.

5.4 Spatial Analysis: UK Housing Prices

Now we test HIP-GP on a standard GP problem, i.e., the transformation \mathcal{L} is an identity map. We apply HIP-GP to (log) prices of apartments as a function of latitude and longitude in England and Wales¹. The data include 180,947 prices from 2018, and we train on 160,947 observations and hold out 20,000 to report test error. We use the standard SVGP as baseline.

Scaling Inducing Points We run HIP-GP on an increasingly dense grid of inducing points M . In all experiments, we use the Matérn (2.5) kernel and apply the block-independent variational family with neighboring block size $M_b = 100$ (10×10) for HIP-GP and SVGP. The maximum number of PCG iterations within HIP-GP is set to 20 and 50 for training and evaluation. The predictive performance measured by RMSE and the training time are displayed in Figure 3c. From this result, we conclude that (i) increasing M improves prediction quality; (ii) the performance of HIP-GP is almost indistinguishable to that of SVGP given the same M . (iii) Again, HIP-GP runs faster than SVGP and scales to larger M . The best prediction of HIP-GP is depicted in Figure 3a and 3b.

¹HM land registry price paid data available [here](#).



(a) Posterior mean

(b) Posterior st. dev.

M	10,000	14,400	19,600	25,600	32,400	40,000
HIP-GP (RMSE)	0.411	0.409	0.400	0.397	0.393	0.389
SVGP (RMSE)	0.412	0.409	0.398	0.396	n/a	n/a
HIP-GP (time)	91.2	115.7	119.8	130.7	129.5	133.2
SVGP (time)	193.6	406.3	668.1	898.2	n/a	n/a

(c) Top row: predictive RMSE. Bottom row: average training time (second) per epoch.

Figure 3: UK Housing Analysis


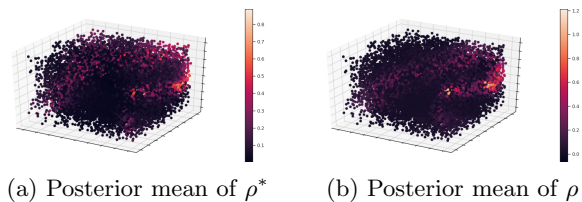


Figure 4: Stochastic optimization is robust to early stopping of PCG iterations.

PCG Iteration Early Stopping Additionally, we examine the effect of the maximum number of PCG iterations when computing \mathbf{k}_n on approximation quality. Figure 4 depicts test RMSE as a function of PCG iteration for $M = 14,400$ on the test dataset of size $N = 20,000$. The final approximation quality is robust to the number of PCG iterations used. The upshot is that HIP-GP needs only a small number of PCG iterations to be effective.

5.5 Inferring Interstellar Dust Map

Finally, we investigate an inter-domain GP problem with \mathcal{L} being the integral transformation: inferring the interstellar dust map from integral observations. The interstellar dust map ρ is a three-dimensional density function at each location in the Galaxy. The observations y , also known as the starlight extenctions, are noisy line integrals of the dust function (Rezaei Kh et al., 2017). We experiment with the Ananke dataset,



	MAE	MSE	loglike
HIP-GP ($M = 62,500$)	0.0101	0.0012	2.3517
SVGP ($M = 16,384$)	0.0153	0.0020	2.0690

(c) Predictive statistics for integrated observations

Figure 5: HIP-GP performance on 20,000 held-out data. Top: Posterior mean predictions in the integrated domain (ρ^*) and the latent domain (ρ). Bottom: We report the mean absolute error (MAE), the mean square error (MSE) and the test log likelihood in the intergrated domain.

which is comprised of 500,000 starlight extinctions within $4\text{kpc} \times 4\text{kpc} \times 2\text{kpc}$ region of a high resolution Milky Way like galaxy simulation — a cutting edge simulation in the field because of the gas and dust resolution (Wetzell et al., 2016; Hopkins et al., 2018; Sanderson et al., 2020). Our goal is to infer the underlying dust map ρ from the noisy extinctions y .

We compare HIP-GP with $M = 62,500$ ($50 \times 50 \times 25$) and SVGP with $M = 16,384$ ($32 \times 32 \times 16$) — the largest M feasible. For both methods, we apply the block-independent variational parameterization with neighboring block size $M_b = 8$ ($2 \times 2 \times 2$), and the Matérn(1.5) kernel. The maximum number of PCG iterations within HIP-GP is set to 200 and 500 for training and evaluation. We use Monte Carlo estimation to compute the inter-domain and transformed-domain covariance functions in Equations 7 and 8. We hold out 20,000 points for evaluation. The posterior mean predictions of the extinctions ρ^* and the latent dust map ρ are displayed in Figure 5a and 5b. The predictive test statistics are summarized in Table 5c. We see that with more inducing points, the predictive accuracy is enhanced. HIP-GP can scale to larger M which enables better prediction quality, while SVGP is limited to M around 16,000.

6 DISCUSSION

We formulate a general SVGP framework for inter-domain GP problems. Upon this framework, we further scale the standard SVGP inference by developing the HIP-GP algorithm, with three technical innovations (i) a fast whitened parameterization, (ii) a novel preconditioner for fast linear system solves with hierar-

chical Toeplitz structure, and (iii) a structured variational approximation. The core idea of HIP-GP lies in the structured exploitations of the kernel matrix and the variational posterior. Therefore, it can be potentially extended to various settings, e.g. the case where a GP is a part of a bigger probabilistic model, and the non-Gaussian likelihoods thanks to recent advance in non-conjugate GP inference (Salimbeni et al., 2018).

Future works involve more in-depth analysis of such CG-based approximate GP methods. On the applied side, we will to apply HIP-GP to the Gaia dataset (Gaia et al., 2018) which consists of nearly 2 billion stellar observations.

Acknowledgements

We thank the reviewers for their detailed feedback and suggestions. The interstellar dust map experiments in Section 5.5 were run on the Iron cluster at the Flatiron Institute, and we are grateful to the scientific computing team for their continual and dedicated technical assistance and support. The Flatiron Institute is supported by the Simons Foundation.

References

- Bauer, M., van der Wilk, M., and Rasmussen, C. E. (2016). Understanding probabilistic sparse gaussian process approximations. In *Advances in neural information processing systems*, pages 1533–1541.
- Chan, R. H. and Ng, M. K. (1996). Conjugate gradient methods for toeplitz systems. *SIAM review*, 38(3):427–482.
- Cressie, N. (1990). The origins of kriging. *Mathematical Geology*, 22(3):239–252.
- Cressie, N. (1992). Statistics for spatial data. *Terra Nova*, 4(5):613–617.
- Cunningham, J. P., Shenoy, K. V., and Sahani, M. (2008). Fast gaussian process methods for point process intensity estimation. In *International Conference on Machine Learning*, pages 192–199. ACM.
- Cutajar, K., Osborne, M., Cunningham, J., and Filippone, M. (2016). Preconditioning kernel matrices. In *International Conference on Machine Learning*, pages 2529–2538.
- Dua, D. and Graff, C. (2017). UCI machine learning repository.
- Evans, T. W. and Nair, P. B. (2018). Scalable gaussian processes with grid-structured eigenfunctions (GP-GRIEF). In *International Conference on Machine Learning*.
- Gaia, C., Brown, A., Vallenari, A., Prusti, T., de Bruijne, J., Babusiaux, C., Juhász, Á., Marschalló, G.,

- Marton, G., Molnár, L., et al. (2018). Gaia data release 2 summary of the contents and survey properties. *Astronomy & Astrophysics*, 616(1).
- Garnett, R., Osborne, M. A., and Roberts, S. J. (2010). Bayesian optimization for sensor set selection. In *Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks*, pages 209–219.
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., Rix, H.-W., Martin, N., Burgett, W., Draper, P. W., Flewelling, H., Hodapp, K., Kaiser, N., et al. (2015). A three-dimensional map of milky way dust. *The Astrophysical Journal*, 810(1):25.
- Hendriks, J. N., Jidling, C., Wills, A., and Schön, T. B. (2018). Evaluating the squared-exponential covariance function in gaussian processes with integral observations. *arXiv preprint arXiv:1812.07319*.
- Hensman, J., Fusi, N., and Lawrence, N. D. (2013). Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, pages 282–290. AUAI Press.
- Hensman, J., Matthews, A. G., Filippone, M., and Ghahramani, Z. (2015). MCMC for variationally sparse gaussian processes. In *Advances in Neural Information Processing Systems*, pages 1648–1656.
- Hestenes, M. R. and Stiefel, E. (1952). *Methods of conjugate gradients for solving linear systems*, volume 49. NBS Washington, DC.
- Hopkins, P. F., Wetzel, A., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C. C., Garrison-Kimmel, S., Hummels, C., et al. (2018). Fire-2 simulations: physics versus numerics in galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 480(1):800–863.
- Izmailov, P., Novikov, A., and Kropotov, D. (2018). Scalable gaussian processes with billions of inducing inputs via tensor train decomposition. In *Artificial Intelligence and Statistics*, pages 726–735.
- Jidling, C., Hendriks, J., Wahlström, N., Gregg, A., Schön, T. B., Wensrich, C., and Wills, A. (2018). Probabilistic modelling and reconstruction of strain. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 436:141–155.
- Kh, S. R., Bailer-Jones, C., Hanson, R., and Fouesneau, M. (2017). Inferring the three-dimensional distribution of dust in the galaxy with a non-parametric method-preparing for gaia. *Astronomy & Astrophysics*, 598:A125.
- Lázaro-Gredilla, M. and Figueiras-Vidal, A. (2009). Inter-domain gaussian processes for sparse inference using inducing features. In *Advances in Neural Information Processing Systems*, pages 1087–1095.
- Leike, R. and Enßlin, T. (2019). Charting nearby dust clouds using gaia data only. *arXiv preprint arXiv:1901.05971*.
- Minka, T. P. (2000). Deriving quadrature rules from gaussian processes. Technical report, Technical report, Statistics Department, Carnegie Mellon University.
- Murray, I. and Adams, R. P. (2010). Slice sampling covariance hyperparameters of latent gaussian models. In *Advances in Neural Information Processing Systems*, pages 1732–1740.
- Nocedal, J. and Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media.
- Pleiss, G., Jankowiak, M., Eriksson, D., Damle, A., and Gardner, J. R. (2020). Fast matrix square roots with applications to gaussian processes and bayesian optimization. *arXiv preprint arXiv:2006.11267*.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. The MIT Press.
- Rezaei Kh, S., Bailer-Jones, C., Hanson, R., and Fouesneau, M. (2017). Inferring the three-dimensional distribution of dust in the galaxy with a non-parametric method. preparing for gaia. *A&A*, 598:A125.
- Riihimäki, J. and Vehtari, A. (2010). Gaussian processes with monotonicity information. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 645–652. JMLR Workshop and Conference Proceedings.
- Saad, Y. (2003). *Iterative methods for sparse linear systems*. SIAM.
- Salimbeni, H., Eleftheriadis, S., and Hensman, J. (2018). Natural gradients in practice: Non-conjugate variational inference in gaussian process models. *arXiv preprint arXiv:1803.09151*.
- Sanderson, R. E., Wetzel, A., Loebman, S., Sharma, S., Hopkins, P. F., Garrison-Kimmel, S., Faucher-Giguère, C.-A., Kereš, D., and Quataert, E. (2020). Synthetic gaia surveys from the fire cosmological simulations of milky way-mass galaxies. *The Astrophysical Journal Supplement Series*, 246(1):6.
- Shewchuk, J. R. et al. (1994). An introduction to the conjugate gradient method without the agonizing pain.
- Shi, J., Titsias, M., and Mnih, A. (2020). Sparse orthogonal variational inference for gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 1932–1942. PMLR.
- Siivola, E., Vehtari, A., Vanhatalo, J., González, J., and Andersen, M. R. (2018). Correcting bound-

- ary over-exploration deficiencies in bayesian optimization with virtual derivative sign observations. In *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE.
- Solak, E., Murray-Smith, R., Leithead, W. E., Leith, D. J., and Rasmussen, C. E. (2003). Derivative observations in gaussian process models of dynamic systems. In *Advances in neural information processing systems*, pages 1057–1064.
- Titsias, M. (2009). Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574.
- van der Wilk, M., Dutordoir, V., John, S., Artemev, A., Adam, V., and Hensman, J. (2020). A framework for interdomain and multioutput gaussian processes. *arXiv preprint arXiv:2003.01115*.
- Wang, K. A., Pleiss, G., Gardner, J. R., Tyree, S., Weinberger, K. Q., and Wilson, A. G. (2019). Exact gaussian processes on a million data points. In *Advances in Neural Information Processing Systems*.
- Wetzel, A. R., Hopkins, P. F., Kim, J.-h., Faucher-Giguère, C.-A., Kereš, D., and Quataert, E. (2016). Reconciling dwarf galaxies with λ cdm cosmology: simulating a realistic population of satellites around a milky way–mass galaxy. *The Astrophysical Journal Letters*, 827(2):L23.
- Wills, A. G. and Schön, T. B. (2017). On the construction of probabilistic newton-type algorithms. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 6499–6504. IEEE.
- Wilson, A. and Nickisch, H. (2015). Kernel interpolation for scalable structured gaussian processes (kiss-gp). In *International Conference on Machine Learning*, pages 1775–1784.
- Wilson, A. G., Dann, C., and Nickisch, H. (2015). Thoughts on massively scalable gaussian processes. *arXiv preprint arXiv:1511.01870*.
- Wilson, A. G., Hu, Z., Salakhutdinov, R. R., and Xing, E. P. (2016). Stochastic variational deep kernel learning. In *Advances in Neural Information Processing Systems*, pages 2586–2594.