
Efficiently Sampling Functions from Gaussian Process Posteriors

James T. Wilson^{*1} Viacheslav Borovitskiy^{*23} Alexander Terenin^{*1}
Peter Mostowsky^{*2} Marc Peter Deisenroth⁴

Abstract

Gaussian processes are the gold standard for many real-world modeling problems, especially in cases where a model’s success hinges upon its ability to faithfully represent predictive uncertainty. These problems typically exist as parts of larger frameworks, wherein quantities of interest are ultimately defined by integrating over posterior distributions. These quantities are frequently intractable, motivating the use of Monte Carlo methods. Despite substantial progress in scaling up Gaussian processes to large training sets, methods for accurately generating draws from their posterior distributions still scale cubically in the number of test locations. We identify a decomposition of Gaussian processes that naturally lends itself to scalable sampling by separating out the prior from the data. Building off of this factorization, we propose an easy-to-use and general-purpose approach for fast posterior sampling, which seamlessly pairs with sparse approximations to afford scalability both during training and at test time. In a series of experiments designed to test competing sampling schemes’ statistical properties and practical ramifications, we demonstrate how *decoupled sample paths* accurately represent Gaussian process posteriors at a fraction of the usual cost.

1. Introduction

Gaussian processes (GPs) are a powerful framework for reasoning about unknown functions f given partial knowledge of their behaviors. In decision-making scenarios, well-calibrated predictive uncertainty is crucial for balancing important tradeoffs, such as exploration versus exploita-

tion and long-term versus short-term rewards. Bayesian methods naturally strike this balance (Ghavamzadeh et al., 2015; Shahriari et al., 2015). While many quantities of interest defined with respect to Bayesian posteriors cannot be computed analytically (such as expectations of nonlinear functionals), they may be readily estimated via Monte Carlo methods. Depending on this sample-based estimator’s relative cost and statistical behavior, its performance may vary from state-of-the-art to method-of-last-resort.

Unlike methods for scalable training (Hensman et al., 2013; Wang et al., 2019), techniques for efficiently sampling from GP posteriors have received relatively little attention in the machine learning literature. On the one hand, naïve approaches to sampling are statistically well-behaved, but scale poorly owing to their need to solve for increasingly large linear systems at test time. On the other hand, fast approximation strategies using Fourier features (Rahimi and Recht, 2008) avoid costly matrix operations, but are prone to misrepresenting predictive posteriors (Wang et al., 2018; Mutny and Krause, 2018; Calandriello et al., 2019). Investigating their respective behaviors, we find that many of these strategies are complementary, with one often excelling where others falter. Motivated by this comparison of strengths and weaknesses, we leverage a lesser known decomposition of GP posteriors that allows us to incorporate the best of both worlds.

Our approach centers on the observation that we may implicitly condition Gaussian random variables by combining them with an explicit corrective term. Translating this intuition to GPs, we may decompose the posterior as the sum of a prior and an update. By doing so, we are able to separately represent each of these terms using a basis well-suited for sampling. This notion of “conditioning by kriging” was first presented by Matheron in the early 1970s, with various applications to geostatistics (Journal and Huijbregts, 1978; de Fouquet, 1994; Chiles and Delfiner, 2009). The concept was later rediscovered in astrophysics (Hoffman and Ribak, 1991; Van de Weygaert and Bertschinger, 1996), where it has been used to help simulate the universe as we know it.

We unite these ideas with techniques from the growing literature on approximate GPs to obtain an easy-to-use and general-purpose approach for accurately sampling from GP posteriors in linear time.

^{*}Equal contribution ¹Imperial College London ²St. Petersburg State University ³St. Petersburg Department of Steklov Mathematical Institute of Russian Academy of Sciences ⁴University College London. Correspondence to: j.wilson17@imperial.ac.uk, viacheslav.borovitskiy@gmail.com, a.terenin17@imperial.ac.uk, and pmostowsky@gmail.com.

2. Review of Gaussian processes

As notation, let $f : \mathcal{X} \rightarrow \mathbb{R}$ be an unknown function with domain $\mathcal{X} \subseteq \mathbb{R}^d$ whose behavior is indicated by a training set consisting of n Gaussian observations $y_i = f(\mathbf{x}_i) + \varepsilon_i$ subject to measurement noise $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$.

A Gaussian process is a random function $f : \mathcal{X} \rightarrow \mathbb{R}$ such that, for any finite set of locations $\mathbf{X}_* \subseteq \mathcal{X}$, the random vector $\mathbf{f}_* = f(\mathbf{X}_*)$ follows a Gaussian distribution. In particular, if $f \sim \mathcal{GP}(\mu, k)$, then $\mathbf{f}_* \sim \mathcal{N}(\boldsymbol{\mu}_*, \mathbf{K}_{*,*})$ is multivariate normal with covariance $\mathbf{K}_{*,*} = k(\mathbf{X}_*, \mathbf{X}_*)$ specified by a kernel k . Henceforth, we assume a zero-mean prior $\mu(\cdot) = 0$ and continuous, stationary covariance function $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$.

Given n observations \mathbf{y} , the GP posterior at \mathbf{X}_* is defined as $\mathbf{f}_* | \mathbf{y} \sim \mathcal{N}(\mathbf{m}_{*,*|n}, \mathbf{K}_{*,*|n})$, where we have defined

$$\begin{aligned} \mathbf{m}_{*,*|n} &= \mathbf{K}_{*,n}(\mathbf{K}_{n,n} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \\ \mathbf{K}_{*,*|n} &= \mathbf{K}_{*,*} - \mathbf{K}_{*,n}(\mathbf{K}_{n,n} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}_{n,*}. \end{aligned} \quad (1)$$

When using (1) to help guide reinforcement learning agents (Kuss and Rasmussen, 2004), black-box optimizers (Snoek et al., 2012), and other complex algorithms, we often rely on samples to estimate quantities of interest. The standard way of generating these samples is via a location-scale transform of Gaussian random variables $\boldsymbol{\zeta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, namely

$$\mathbf{f}_* | \mathbf{y} = \mathbf{m}_{*,*|n} + \mathbf{K}_{*,*|n}^{1/2} \boldsymbol{\zeta}, \quad (2)$$

where $(\cdot)^{1/2}$ denotes a matrix square root, such as a Cholesky factor. Since this scheme is exact up to numerical error, we take it to be the gold standard against which the sample quality of alternatives will be judged. Unfortunately, this sampling strategy is also one of the least scalable, since the cost of computing $\mathbf{K}_{*,*|n}^{1/2}$ is already $\mathcal{O}(n^3)$.

The first column of Figure 1 visualizes sampling from a GP posterior given varying amounts of training data n . Since matrices on the right-hand-side of (1) grow as training sets increases in size, this method of sampling can be seen to accumulate little to no error as n increases. However, this growth requires us to invert increasingly large matrices both during training and at test time, causing standard GP inference and sampling methods to scale poorly in n .

2.1. Function-space approximations to GPs

The preceding interpretation of GPs, as distributions over functions with Gaussian marginals, is commonly known as the *function-space* view (Rasmussen and Williams, 2006). From this perspective, a natural way of approximating GPs is to represent f in terms of its behavior $\mathbf{u} = f(\mathbf{Z})$ at a carefully chosen set of *inducing locations* $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$. In line with this function-space intuition of reasoning about f

via a small set of locations, this family of approximations is commonly referred to as *sparse Gaussian processes*.

Rather than directly conditioning on observations \mathbf{y} , sparse GPs begin by defining *inducing distributions* $q(\mathbf{u})$ that explain for the data. Over the years, distinct iterations of sparse GPs have proposed different inducing paradigms (Snelson and Ghahramani, 2006; Titsias, 2009; Hensman et al., 2017). In this work, we remain agnostic regarding the choice of $q(\mathbf{u})$ and simply assume access to draws $\mathbf{u} \sim q(\mathbf{u})$.

Given $q(\mathbf{u})$, we approximate posterior distributions as

$$p(\mathbf{f}_* | \mathbf{y}) \approx \int_{\mathbb{R}^m} p(\mathbf{f}_* | \mathbf{u}) q(\mathbf{u}) d\mathbf{u}. \quad (3)$$

If $\mathbf{u} \sim \mathcal{N}(\boldsymbol{\mu}_u, \boldsymbol{\Sigma}_u)$, we compute this integral analytically to obtain a Gaussian distribution with mean and covariance

$$\begin{aligned} \mathbf{m}_{*,*|m} &= \mathbf{K}_{*,m} \mathbf{K}_{m,m}^{-1} \boldsymbol{\mu}_m \\ \mathbf{K}_{*,*|m} &= \mathbf{K}_{*,*} + \mathbf{K}_{*,m} \mathbf{K}_{m,m}^{-1} (\boldsymbol{\Sigma}_u - \mathbf{K}_{m,m}) \mathbf{K}_{m,m}^{-1} \mathbf{K}_{m,*}. \end{aligned} \quad (4)$$

By virtue of explaining for n observations using m inducing variables, sparse GPs can be trained with $\mathcal{O}(\tilde{n}m^2)$ time complexity, where the choice of batch size $1 \leq \tilde{n} \leq n$ depends on the particular algorithm. Since high-quality approximations can be constructed using $m \ll n$ (Burt et al., 2019), sparse GPs drastically improve upon their exact counterparts' $\mathcal{O}(n^3)$ scaling.

While posterior moments (4) may be computed at reduced cost, this benefit does not carry over when sampling. The standard procedure for sampling from sparse GPs is the same as in (2) and incurs $\mathcal{O}(n^3)$ cost.¹ When used to drive Monte-Carlo-based algorithms, sparse GPs can therefore be fast during training but slow during deployment. The middle column of Figure 1 depicts samples from a sparse GP posterior with $m = 8$ inducing locations.

2.2. Weight-space approximations to GPs

In the function-space view of GPs, we reason about f in terms of the values it may assume at locations $\mathbf{x} \in \mathcal{X}$. We now turn to the *weight-space* view, where we will reason about f as a weighted sum of basis functions. Per the kernel trick (Schölkopf and Smola, 2001), k can be viewed as the inner product in a reproducing kernel Hilbert space (RKHS) \mathcal{H} equipped with a feature map $\varphi : \mathcal{X} \rightarrow \mathcal{H}$. If \mathcal{H} is separable, we may approximate this inner product as

$$k(\mathbf{x}, \mathbf{x}') = \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle_{\mathcal{H}} \approx \boldsymbol{\phi}(\mathbf{x})^\top \boldsymbol{\phi}(\mathbf{x}'), \quad (5)$$

where $\boldsymbol{\phi} : \mathcal{X} \rightarrow \mathbb{R}^\ell$ is a finite-dimensional feature map (Rasmussen and Williams, 2006). For stationary covariance functions, Bochner's theorem implies that a suitable

¹Select inducing point methods allow fast sampling from degenerate posteriors, see Quiñero-Candela and Rasmussen (2005).




Figure 1: Comparison of GP posteriors and sample paths given $n = 4$ (top) and $n = 1000$ (bottom) observations at shaded locations. Error values shown in bottom right-hand corner of each figure denote 2-Wasserstein distances (see Section 4) between empirical (closed-form) posteriors and the true posterior (dashed black). *Left*: Mean and two standard deviations of exact GP posterior (green) along with samples at $*$ = 1024 test locations. *Middle*: Sparse GP with inducing variables \mathbf{u} at $m = 8$ locations $\mathbf{z} \in \mathcal{X}$ denoted by ‘ \circ ’. *Right*: Random Fourier feature-based GP with $\ell = 2000$ basis functions; for $n = 1000$, variance starvation has started to set in and predictions away from the data show visible signs of deterioration.

ℓ -dimensional feature map can be constructed via a set of *random Fourier features* (RFF) (Rahimi and Recht, 2008).

In this case, we have $\phi_i(\mathbf{x}) = \sqrt{2/\ell} \cos(\boldsymbol{\theta}_i^\top \mathbf{x} + \tau_i)$, where $\boldsymbol{\theta}_i$ are sampled proportional to the kernel’s spectral density and $\tau_j \sim U(0, 2\pi)$. By defining the *Bayesian linear model*

$$f(\cdot) = \sum_{i=1}^{\ell} w_i \phi_i(\cdot) \quad w_i \sim \mathcal{N}(0, 1), \quad (6)$$

we obtain an ℓ -dimensional GP approximation. As in previous sections, f is now a random function with Gaussian marginals. However, this stochasticity is now entirely controlled by the distribution of weights \mathbf{w} .

For Gaussian likelihoods, the posterior weight distribution $\mathbf{w} \mid \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{w}|n}, \boldsymbol{\Sigma}_{\mathbf{w}|n})$ is Gaussian with moments

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{w}|n} &= (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma^2 \mathbf{I})^{-1} \boldsymbol{\Phi}^\top \mathbf{y} \\ \boldsymbol{\Sigma}_{\mathbf{w}|n} &= (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma^2 \mathbf{I})^{-1} \sigma^2, \end{aligned} \quad (7)$$

where $\boldsymbol{\Phi} = \phi(\mathbf{X})$ is an $n \times \ell$ feature matrix. In both cases, we may solve for the right-hand side at $\mathcal{O}(\min\{\ell, n\}^3)$ cost by applying the Woodbury matrix identity.

Approximating the posterior $f \mid \mathbf{y}$ as weighted sums of basis functions in (6) is particularly advantageous for purposes of sampling. As before, we may generate draws from (7) by first computing $\boldsymbol{\Sigma}_{\mathbf{w}|n}^{1/2}$ at $\mathcal{O}(\ell^3)$ cost.² Unlike before, we now sample weight vectors rather than function values and each draw now defines an actual *function* evaluable at arbitrary locations $\mathbf{x} \in \mathcal{X}$. These methods have recently attracted

²Alternatively, we may generate draws at $\mathcal{O}(n^3)$ cost by instead utilizing an eigen-decomposition (Seeger, 2008).

attention in Bayesian optimization (Hernández-Lobato et al., 2014; Shahriari et al., 2015), where the ability to fine-tune test locations \mathbf{X}_* by differentiating through samples is particularly valuable (Wilson et al., 2018).

Unfortunately, these efficiency gains are counterbalanced by loss in expressivity. GP approximations equipped with covariance functions arising from finite-dimensional feature maps are well-known to exhibit undesirable pathologies at test time, see Rasmussen and Quinero-Candela (2005). In the case of Fourier-feature-based approximations, this tendency manifests as *variance starvation*, whereby their extrapolatory predictions become increasingly ill-behaved as n increases (Wang et al., 2018; Mutny and Krause, 2018; Calandriello et al., 2019). Intuitively, this occurs because the Fourier basis is only efficient at representing stationary GPs. The posterior, however, is generally nonstationary. This tendency is evident in the right column of Figure 1: samples from the posterior clearly deteriorate in quality as we transition from low to high-data regimes.

Motivation. Prior to presenting our primary contributions, we briefly pause to restate key trends discussed above and shown in Figure 1. Sampling from sparse GPs accommodates large amounts of training data $n = |\mathbf{X}|$, but scales poorly with the number of test locations $* = |\mathbf{X}_*|$. Conversely, sampling from Fourier-feature-based weight-space approximations scales gracefully with $*$, but results in high approximation error as n increases. Function- and weight-space approaches to sampling from GP posteriors therefore exhibit opposing strengths and weaknesses.

Hence, the question: *can we obtain the best of both worlds?*

3. Sampling with Matheron’s rule

Our approach to designing an improved sampling scheme, which doubles as a rough outline for this section, is as follows: (i) analyze the shortcomings of existing methods, (ii) identify a decomposition of GPs that isolates these issues, (iii) represent the different terms using bases that address their corresponding issues. We begin by reviewing *Matheron’s rule* for Gaussian random variables (Journel and Huijbregts, 1978; Chiles and Delfiner, 2009; Doucet, 2010), which is central to our analysis.

Theorem 1 (Matheron’s Rule). *Let \mathbf{a} and \mathbf{b} be jointly Gaussian random variables. Then the random variable \mathbf{a} conditional on $\mathbf{b} = \beta$ is equal in distribution to*

$$(\mathbf{a} \mid \mathbf{b} = \beta) \stackrel{d}{=} \mathbf{a} + \text{Cov}(\mathbf{a}, \mathbf{b}) \text{Cov}(\mathbf{b}, \mathbf{b})^{-1}(\beta - \mathbf{b}). \quad (8)$$

Proof. Follows immediately by computing the mean and covariance of both sides. \square

Intuitively, Matheron’s rule tells us that conditional random variable $\mathbf{a} \mid \mathbf{b}$ can be broken up into a term representing the prior $p(\mathbf{a}, \mathbf{b})$ and a term that communicates the error in the prior upon observing that $\mathbf{b} = \beta$. Hence, we may sample $\mathbf{a} \mid \mathbf{b}$ by drawing \mathbf{a} and \mathbf{b} together from the prior and, subsequently, updating \mathbf{a} to account for residuals $\beta - \mathbf{b}$ as in (8). The corresponding statement for GPs is as follows.

Corollary 2. *For a Gaussian process $f \sim \mathcal{GP}(0, k)$ with marginal $\mathbf{f}_m = f(\mathbf{Z})$, the process conditioned on $\mathbf{f}_m = \mathbf{u}$ admits, in distribution, the representation*

$$\underbrace{(f \mid \mathbf{u})(\cdot)}_{\text{posterior}} \stackrel{d}{=} \underbrace{f(\cdot)}_{\text{prior}} + \underbrace{k(\cdot, \mathbf{Z})\mathbf{K}_{m,m}^{-1}(\mathbf{u} - \mathbf{f}_m)}_{\text{update}}. \quad (9)$$

Proof. By Theorem 1, the corollary holds for arbitrary finite-dimensional marginals, so the claim follows. \square

Unlike (1) and (4), Corollary 2 defines a *pathwise update*: rather than conditioning the prior as a distribution, we update the prior as realized in terms of sample paths. As we will soon see, this ability to go from prior to posterior (function) draws without needing to compute posterior covariance matrices (and their square-roots) will be the key to unlocking fast and accurate sampling from GP posteriors.

We are not the first to have realized this fact. This approach to simulating Gaussian conditionals is implicit in Matheron’s pioneering work in the field of geostatistics, where it was subsequently popularized by Journel and Huijbregts (1978). Decades later, (9) was rediscovered in the context of N -body simulations by Hoffman and Ribak (1991). We combine these ideas with modern machine learning methods (such as sparse GPs and random Fourier features) to create a more efficient approach to sampling.

3.1. Pathwise updates in weight- and function-spaces

Rewriting the standard formulae for sparse and exact GP posteriors, respectively, as pathwise updates in accordance with Theorem 1, we obtain

$$\mathbf{f}_* \mid \mathbf{u} \stackrel{d}{=} \mathbf{f}_* + \mathbf{K}_{*,m} \mathbf{K}_{m,m}^{-1}(\mathbf{u} - \mathbf{f}_m) \quad (10)$$

$$\mathbf{f}_* \mid \mathbf{y} \stackrel{d}{=} \mathbf{f}_* + \mathbf{K}_{*,n}(\mathbf{K}_{n,n} + \sigma^2 \mathbf{I})^{-1}(\mathbf{y} - \mathbf{f} - \varepsilon). \quad (11)$$

When sampling from sparse GPs in (10), we draw \mathbf{f}_* and \mathbf{f}_m together from the prior, and independently generate target values $\mathbf{u} \sim q(\mathbf{u})$. When sampling from exact GPs in (11), we again begin by jointly drawing \mathbf{f}_* and \mathbf{f} from the prior. Here however, we no longer need to generate targets $\mathbf{u} = \mathbf{y}$. Instead, we combine \mathbf{f} with noise variates $\varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ such that $\mathbf{f} + \varepsilon$ constitutes a draw from the prior distribution of \mathbf{y} .

Turning to the weight-space setting, the analogous pathwise update given an initial weight vector $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is then

$$\mathbf{w} \mid \mathbf{y} \stackrel{d}{=} \mathbf{w} + \Phi^\top (\Phi \Phi^\top + \sigma^2 \mathbf{I})^{-1}(\mathbf{y} - \Phi^\top \mathbf{w} - \varepsilon). \quad (12)$$

At first glance, it appears that sampling via Theorem 1 does not improve over standard methods. Whereas (12) is of modest practical interest (it allows us to sample at $\mathcal{O}(\min\{\ell, n\}^3)$ cost without resorting to an eigen-decomposition), (10) and (11) are actually more expensive than their standard counterparts.

At the same time, however, Theorem 1 allows us to view GP posteriors from a different perspective. In particular, separating the effect of the prior from that of the data allows us to better diagnose the different sampling scheme’s shortcomings. For function-space approaches, we see that $\mathcal{O}(*^3)$ time complexity is specific to the prior, since the update is linear in $*$. For weight-space methods, we see that erratic extrapolations stem from difficulty representing the data (i.e., the update), since stationary priors are well-behaved under the Fourier basis. Equipped with a better understanding of *why* these methods fail, we now demonstrate *how* to address these issues.

3.2. Pathwise updates with decoupled bases

So far, we have implicitly assumed a unified view of GP posteriors: when sampling in weight-space and in function-space, we sought to generate draws from conditional distributions over weight vectors and function values, respectively. Several recent works (Cheng and Boots, 2017; Salimbeni et al., 2018; Shi et al., 2020) have introduced decompositions that separately represent different aspects of GPs via different bases, such as RKHS subspaces and their orthogonal complements. There, the authors exploit the different bases’ properties to better approximate the overarching process. We will do the same, but our goal will be to efficiently sample from the accompanying posteriors.




Figure 2: Visual overview of decoupled sampling with a weight-space prior (orange) and a function-space update (purple); this example continues from Figure 1. *Left*: 1000 Fourier basis functions $\phi_i(\mathbf{x}) = \cos(\boldsymbol{\theta}_i^\top \mathbf{x} + \tau_i)$ are used to construct a function draw $f(\cdot) = \phi(\cdot)^\top \mathbf{w}$ from an approximate prior, resulting in residuals at each of $m = 8$ inducing locations $\mathbf{z}_j \in \mathbf{Z}$ (black circles). As additional context, residuals at extraneous locations $\mathbf{x} \notin \mathbf{Z}$ are also shown (dashed black). *Middle*: a conditional sample path $f | \mathbf{u}$ (blue) is formed by adding an update consisting of basis functions $\psi_j(\cdot) = k(\cdot, \mathbf{z}_j)$ to f . *Right*: the empirical distribution of sample paths $f | \mathbf{u}$ is compared with that of the sparse GP posterior (dashed black). 2-Wasserstein errors of empirical (closed-form) posteriors were measured against the exact GP’s moments.

Corollary 2 is a pathwise update for Gaussian random variables that doubles as a decomposition of the posterior. To further build on this distinction, we restate Corollary 2 using a weight-space approximation to the prior

$$(f | \mathbf{u})(\cdot) \stackrel{d}{\approx} \underbrace{\sum_{i=1}^{\ell} w_i \phi_i(\cdot)}_{\text{weight-space prior}} + \underbrace{\sum_{j=1}^m v_j k(\cdot, \mathbf{z}_j)}_{\text{function-space update}}, \quad (13)$$

sparse posterior

where we have defined $\mathbf{v} = \mathbf{K}_{m,m}^{-1}(\mathbf{u} - \Phi^\top \mathbf{w})$. The equivalent expression for exact GPs with Gaussian observations is obtained by adding noise $\varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ to $\Phi^\top \mathbf{w}$ and replacing \mathbf{Z} , \mathbf{u} , and $\mathbf{K}_{m,m}^{-1}$ with \mathbf{X} , \mathbf{y} , and $(\mathbf{K}_{n,n} + \sigma^2 \mathbf{I})^{-1}$.

Figure 2 acts as a visual guide for decoupled sampling, showing the progression from prior (6) to posterior (13). Stepping through this example: (i) we draw a function f from an approximate prior, (ii) we construct an update function to account for the residuals $\mathbf{u} - f(\mathbf{Z})$ produced by an independent sample $\mathbf{u} \sim q(\mathbf{u})$, (iii) we add these functions together to obtain a function drawn from an approximate posterior (13) that we may freely evaluate anywhere in \mathcal{X} .

In (13), we obtain an efficient approximator by separately discretizing the prior using Fourier basis functions $\phi_i(\cdot)$ and the update using canonical basis functions $k(\cdot, \mathbf{z}_j)$. While other decompositions exist (see Appendix A), this particular decoupling directly capitalizes upon each basis’ strengths: the Fourier basis is well-suited for representing the prior (Rahimi and Recht, 2008) and the canonical basis is well-suited for representing the data (Burt et al., 2019).

By combining these bases as in (13), we therefore inherit the best of both worlds. As in weight-space methods, we may efficiently approximate draws from the prior using an ℓ -dimensional Bayesian linear model $f(\cdot) = \phi(\cdot)^\top \mathbf{w}$,

where weights \mathbf{w} are standard normal (owing to the assumed stationarity of kernel k).³ As in function-space methods, we may faithfully represent the data since basis functions $k(\cdot, \mathbf{z}_j)$ are in one-to-one correspondence with inducing locations $\mathbf{z}_j \in \mathbf{Z}$. This retention of statistical propriety is evident on the right-hand side of Figure 2: despite using half as many basis functions as the weight-space method (see Figure 1), decoupled sampling’s statistical properties mirror those of the gold standard.

Expanding upon these properties, we note the following intuitive behaviors. The update function’s role of “correcting” for residuals $\mathbf{u} - f(\mathbf{Z})$ subsumes that of representing the posterior mean: replacing the prior draw f with the prior mean $\mathbb{E}[f]$ reduces (13) to the standard expression for the conditional expectation $\mathbb{E}[f | \mathbf{u}]$. Since this task is performed in the canonical basis, the expected value of decoupled sample paths is guaranteed to coincide with that of (sparse) GP’s posterior. As a result, decoupled sampling becomes increasingly well-behaved as the number of training (inducing) locations grows and uncertainty decreases. Conversely, we are guaranteed to revert to the prior as we move away from the data, assuming local basis functions $k(\cdot, \mathbf{z})$ (see the center column of Figure 2).

Decoupled sampling complements these desiderata with function draws’ inherent strengths. The immediate implication here is that decoupled sampling scales linearly with respect to the number of test locations \mathbf{X}_* . A more subtle point is that these functions are *pathwise differentiable* with respect to \mathbf{x} —an affordance with significant consequences when seeking to understand Gaussian processes’ extrema.

³This point was not lost on Hoffman and Ribak (1991), who similarly approximated stationary priors using spectral methods.




Figure 3: Empirical estimates of 2-Wasserstein distances between true posteriors and empirical distributions of 100,000 samples at 1024 test locations \mathbf{X}_* given varying amounts of training data, shown as medians and interquartile ranges (shaded regions) measured over 64 independent trials. Weight-space (orange) and decoupled (blue) sampling utilized a total of $b = m + \ell$ basis functions. Results using $\ell \in \{1024, 4096, 16384\}$ initial bases correspond with {light, medium, dark} tones and $\{\triangle, \diamond, \star\}$ markers. See Appendix C.1 for extended results and comparison with LOVE (Pleiss et al., 2018).

While these insights tell us about decoupled sampling’s qualitative behavior, they do not allow us to make quantitative statements about its purported benefits. To this end, the following section provides a means of objectively comparing different sampling schemes’ statistical properties.

3.3. Error bounds

Due to its use of an approximate prior, decoupled sampling introduces an additional source of error at test time. Anecdotal evidence (see Figure 2) suggests that this sampling error is often small in comparison to the error introduced by inducing point approximations. Here, we study decoupled sampling’s analytic properties to clarify how quality of the approximate prior impacts that of decoupled function draws. We present the results of this analysis below, and reserve proofs and derivations of associated constants for Appendix B. As a convenient shorthand, we refer to the particular decoupled sparse GP approximation introduced in (13) as DSGP.

Gaussian processes are often compared via a suitable notion of similarity on the space of probability distributions (Gibbs and Su, 2002). We focus on the 2-Wasserstein distance between GPs (Mallasto and Feragen, 2017). Unlike various alternatives, Wasserstein distances between an exact GP and finite-dimensional approximations thereof are finite and, therefore, facilitate meaningful performance comparisons. Moreover, 2-Wasserstein distances between finite-dimensional Gaussian marginals can be efficiently computed as a proxy for distances between processes themselves. For DSGP, we may bound this distance as follows.

Proposition 3. *Assume that $\mathcal{X} \subseteq \mathbb{R}^d$ is compact and that stationary kernel k is sufficiently regular for $f \sim \mathcal{GP}(0, k)$ to be almost surely continuous. Let $f | \mathbf{y}$ be the posterior of f , $f^{(s)}$ that of a sparse GP, and $f^{(d)}$ that of a corresponding DSGP defined via an approximate prior $f^{(w)}$. Then we have*

$$\begin{aligned} & W_{2, L^2(\mathcal{X})}(f^{(d)}, f | \mathbf{y}) \\ & \leq \underbrace{W_{2, L^2(\mathcal{X})}(f^{(s)}, f | \mathbf{y})}_{\text{error in the (sparse) posterior}} + C_1 \underbrace{W_{2, C(\mathcal{X})}(f^{(w)}, f)}_{\text{error in the prior}}, \end{aligned} \quad (14)$$

where $W_{2, L^2(\mathcal{X})}$ and $W_{2, C(\mathcal{X})}$ are the 2-Wasserstein distances over $L^2(\mathcal{X})$ and the space of continuous functions $C(\mathcal{X})$ equipped with the supremum norm, respectively.

Proof. Appendix B. □

This bound tells us that the error exhibited by DSGP sample paths cleanly separates into independent terms associated with the sparse GP and the approximate prior. In particular, the way in which error in the prior carries over to the posterior is controlled by the inducing locations \mathbf{Z} , which C_1 depends on, but not by the inducing distribution $q(\mathbf{u})$.

We continue this analysis by studying DSGP’s moments. Since a DSGP’s mean is guaranteed to coincide with that of a sparse GP, we focus on the error it introduces into the posterior covariance. When using RFF to approximate the prior, this error will depend on the ℓ -dimensional basis ϕ given by parameters $\tau \sim U(0, 2\pi)$ and $\theta \sim s(\theta)$, where $s(\cdot)$ denotes the (normalized) spectral density of k . We therefore bound the expectation of this error.

Proposition 4. *Continuing from Proposition 3, let $k^{(f|\mathbf{y})}$, $k^{(w)}$, $k^{(s)}$, $k^{(d)}$ respectively denote the covariance functions of processes $f | \mathbf{y}$, $f^{(w)}$, $f^{(s)}$, $f^{(d)}$. Denoting the supremum norm over continuous functions by $\|\cdot\|_{C(\mathcal{X}^2)}$, it follows that*

$$\begin{aligned} & \mathbb{E}_\phi \|k^{(d)} - k^{(f|\mathbf{y})}\|_{C(\mathcal{X}^2)} \\ & \leq \|k^{(s)} - k^{(f|\mathbf{y})}\|_{C(\mathcal{X}^2)} + \frac{C_2 C_3}{\sqrt{\ell}}, \end{aligned} \quad (15)$$

where the constants C_2 and C_3 are given by Sutherland and Schneider (2015) and in Appendix B, respectively.

Proof. Appendix B. □




Figure 4: Median performances and interquartile ranges of parallel Thompson Sampling (TS) and popular baselines when optimizing d -dimensional functions drawn from GP priors. Function-space TS delivers competitive performance for $d = 2$, but is held back by its inability to efficiently utilize gradient information to combat the curse of dimensionality. RFF-based TS avoids this issue but requires $b \gg m$ basis functions to perform well. TS with decoupling sampling matches or outperforms competing approaches in all observed cases. See Appendix C.2 for additional results and runtime distributions.

Much like DSGPs themselves, the error in the posterior covariance separates into terms associated with the covariance of the sparse GP $k^{(s)}$ and approximate prior $k^{(w)}$. This latter source of error represents discrepancies introduced by using RFF to approximate the prior and decays at a *dimension-free rate* as the number of basis functions ℓ increases. Intuitively, this behavior reflects RFF’s nature as a Monte Carlo estimator of the true covariance. In practice, the number of training points n typically grows faster than the dimensionality d . Hence, purely RFF-based GP posteriors struggle to capitalize upon this property due to variance starvation. Since DSGP does not exhibit this pathology, it fully benefits from this dimension-free rate of convergence.

4. Experiments

We investigate decoupled sampling’s behavior in a series of sample tests accompanied by two practical applications, Thompson sampling and dynamical system simulation. Each of these experiments highlights different properties of decoupled sample paths: uncertainty calibration, reliability and differentiability, and computational savings.⁴

Testing uncertainty calibration with the 2-Wasserstein.

To better understand how the bounds presented in Section 3.3 manifest in the real world, we put the various sampling schemes through numerical experiments that empirically estimated the 2-Wasserstein distance bounded by (14). These tests allow us to see how this distance is affected by factors, such as the number of training points, whose effects are difficult to directly analyze. In each trial, we measured the distance between the true posterior and empirical distributions of samples generated using the various strategies introduced in the paper. To eliminate confounding variables, experiments were run using exact GPs with known hyperparameters (see Appendix C for details).

⁴Code: <https://github.com/j-wilson/GPflowSampling>

Our investigation focuses on each method’s behavior as the number of inducing locations m (equivalently, the number of training points n) increases relative to the number of basis functions employed. For fair comparison, the total number of basis functions $b = m + \ell$ utilized by weight-space and decoupled samplers was held equal, where ℓ denotes an initial allocation. For decoupled sampling, ℓ specifies the number of Fourier features used to approximate the prior.

Figure 3 shows that weight-space sampling tends to deteriorate as m increases relative to b . Variance starvation causes sample paths’ extrapolatory behavior to increasingly misrepresent the posterior. This issue is exacerbated as dimensionality d rises, since we can expect the (randomly chosen) test locations \mathbf{X}_* to lie further and further away from the data.

In contrast, decoupled sampling retains its performance, and may even improve. This reflects the fact that the data is represented in an efficient basis that grows alongside it. For sparse GPs with $m \geq n$ (which includes exact GPs as a special case), we may always represent the data exactly: usually, however, $m \ll n$ inducing locations (i.e., kernel basis functions) suffice (Burt et al., 2019). Since we expect posteriors to contract as training sets expand, the functions drawn from these posteriors behave increasingly similar to their means. Since decoupled sample paths are guaranteed to exhibit the correct means, their statistical properties may improve. This process occurs more slowly in higher-dimensional cases. However, since away from data these function draws revert to the approximate prior, they exhibit constant error when extrapolating—the approximation error of said prior.

Thompson Sampling with robust, differentiable draws.

Thompson Sampling (TS) is a classic strategy for decision-making in the face of uncertainty, whereby a choice $x \in \mathcal{X}$ is selected according to its estimated probability of being




Figure 5: Sparse GP-based simulation of a FitzHugh-Nagumo model neuron subject to evolution noise $\varepsilon_t \sim \mathcal{N}(\mathbf{0}, 10^{-2}\mathbf{I})$ and current injection $I(t) \in \mathbb{R}$. *Left*: True drift function f given a fixed current $I(t) = 0.5$. *Middle*: Medians and interquartile ranges of 1000 voltage traces generated in response to a sinusoidal control signal (dashed black) using iterative (orange) and decoupled (blue) sampling are compared with those of ground truth simulations (gray). *Upper right*: Runtime comparison of iterative and decoupled sampling: the former scales cubically, while the latter runs in linear time. *Lower right*: 2-Wasserstein distances between state distributions at times t are approximated using the Sinkhorn algorithm (Cuturi, 2013). The noise-floor (gray) was established using additional ground truth simulations.

optimal (Thompson, 1933). When used as a vehicle for GP-based optimization, TS evaluates a pathwise minimizer

$$\mathbf{x}_{n+1} \in \arg \min_{\mathbf{x} \in \mathcal{X}} (f | \mathbf{y})(\mathbf{x}) \quad (16)$$

of a function drawn $f | \mathbf{y}$ from the posterior. Upon finding this minimizer, \mathbf{x}_{n+1} is evaluated to obtain y_{n+1} , the pair $(\mathbf{x}_{n+1}, y_{n+1})$ is added to the training set, and the process repeats. In practice, this algorithm is (embarrassingly) parallelized by independently drawing $\kappa > 1$ functions and evaluating a minimizer of each one (Hernández-Lobato et al., 2017; Kandasamy et al., 2018).

We compare the performance of parallel TS equipped with the various sampling schemes discussed in Section 3, along with two common baselines. To help eliminate confounding variables, experiments were run using functions drawn from known GP priors with fixed measurement noise $y_i \sim \mathcal{N}(f_i, 10^{-3})$. Across trials, we varied both the dimensionality d of search spaces $\mathcal{X} = [0, 1]^d$ and the number of initial basis functions ℓ . We set $\kappa = d$, but this choice was not found to greatly influence results. The total number of basis functions allocated to weight-space and decoupled samplers was again matched, so that $b = m + \ell$.

Figure 4 shows that different methods of sampling from GP posteriors dramatically influence achieved performance. While all methods suffered from the curse of dimensionality, TS in function-space deteriorates most aggressively, owing to its inability to efficiently exploit gradient information and to the prohibitive cost for generating large sample vectors $\mathbf{f}_* | \mathbf{y}$. Weight-space TS resolves both of these issues and, therefore, performs competitively—so long as $b \gg m$, in which case it accurately approximates the posterior. On the

other hand, TS in weight-space collapses due to variance starvation as m increases relative to b , often performing worse than simpler alternatives.

Decoupled sampling avoids these shortcomings. As function draws, decoupled sample paths $(f | \mathbf{y})(\mathbf{X}_*)$ boast linear time complexity $\mathcal{O}(\ast)$ and can be minimized by pathwise differentiating with respect to \mathbf{X}_* . Moreover, because the canonical basis is able to efficiently represent the data, these sample paths retain their statistical properties even when b is comparable to n or, in the case of sparse GPs, when $b \ll n$.

Simulating dynamical systems in linear time. Model-based simulators are commonly used in cases where real-world data collection proves impractical or impossible. For example, GP surrogates are a key component of state-of-the-art methods for solving the types of continuous control problems seen in robotics (Deisenroth et al., 2015; Kamthe and Deisenroth, 2018). Without loss of generality, we assume that our goal is to model a time-invariant system whose dynamics are governed by a stochastic differential equation, discretized according to the Euler-Maruyama integrator

$$\Delta \mathbf{s}_t = \mathbf{s}_{t+1} - \mathbf{s}_t = f(\mathbf{s}_t, \mathbf{c}_t) \Delta t + \sqrt{\Delta t} \Sigma \varepsilon_t, \quad (17)$$

where \mathbf{s}_t denotes the state at time t , $\mathbf{c}_t \in \mathcal{U} \subseteq \mathbb{R}^c$ a control input, and $\varepsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ a standard normal random vector.

Having trained a (sparse) GP to represent possible drift functions f , we simulate the system’s evolution over time by *unrolling*: given a state-control pair $(\mathbf{s}_t, \mathbf{c}_t)$, we sample a transition $\Delta \mathbf{s}_t$ according to the GP posterior and step as in (17). Since the resulting trajectory $\mathbf{s}_{1:t}$ is determined online, standard approaches to sampling require us to itera-

tively condition on the preceding sample f_t when drawing $f_{t+1} \mid \mathbf{f}_{1:t}$. Use of caching and rank-1 downdates help limit associated costs however, the resulting algorithm’s time complexity still scales cubically in the number of steps t (see Appendix C.3). By virtue of drawing functions, decoupled sampling avoids this machinery and allows us to simulate trajectories in linear time $\mathcal{O}(t)$.

To better understand the practical ramifications of unrolling with decoupled samples, we used a sparse GP to simulate the dynamics of a well-known model of a biological neuron (FitzHugh, 1961; Nagumo et al., 1962). Results are shown in Figure 5. For both sampling schemes, simulated trajectories accurately characterizes the ways in which the system may respond to a given control signal. Their respective costs, however, vary dramatically: simulations that required 10 hours using the iterative approach, owing to cubic costs, ran in 20 seconds using decoupled sampling while achieving similar accuracy.

5. Conclusion

Decomposing Gaussian processes is a general strategy for constructing efficient approximation schemes. We have focused on a particular case, where a posterior is seen as the sum of a prior and an update, and shown how this decoupling can be exploited to efficiently draw functions from said posterior. Even within this choice of decomposition however, optimal treatment of these components will ultimately depend upon the nature of the task at hand. For example, when working with structured covariance matrices, it is sometimes possible to efficiently generate draws from the prior without introducing approximation error (Dietrich and Newsam, 1997; Wilson and Nickisch, 2015). These alternatives can then be combined with ideas discussed in previous sections to achieve the desired balance of speed versus accuracy.

Owing to the generality of our assumptions and simplicity of our proposals, decoupled sampling can be used as a plug-in extension to existing sample-based algorithms driven by (sparse) GPs. Separately representing the prior and the data with bases better suited for sampling allows us to obtain the “best of both worlds” by bringing together previous methods’ strengths. The result of this union, decoupled sampling, draws functions from GPs that may be evaluated in linear time without fear of misrepresenting their posteriors.

Acknowledgments

The authors would like to express their gratitude to Prof. Mikhail Lifshits, without whom our collaboration would never have started. This research was partially supported by “Native towns”, a social investment program of PJSC “Gazprom Neft” and by the Ministry of Science and Higher

Education of the Russian Federation, agreements N^o 075-15-2019-1619 and N^o 075-15-2019-1620. The support of the EPSRC Centre for Doctoral Training in High Performance Embedded and Distributed Systems (reference EP/L016796/1) is gratefully acknowledged.

References

- J. Bergstra and Y. Bengio. Random search for hyperparameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012. Cited on page 15.
- D. R. Burt, C. E. Rasmussen, and M. van der Wilk. Rates of convergence for sparse variational Gaussian process regression. In *International Conference on Machine Learning*, pages 862–871, 2019. Cited on pages 2, 5, 7.
- D. Calandriello, L. Carratino, A. Lazaric, M. Valko, and L. Rosasco. Gaussian process optimization with adaptive sketching: scalable and no regret. In *Conference on Learning Theory*, pages 533–557, 2019. Cited on pages 1, 3.
- C.-A. Cheng and B. Boots. Variational inference for Gaussian process models with linear complexity. In *Advances in Neural Information Processing Systems*, pages 5184–5194, 2017. Cited on page 4.
- J.-P. Chiles and P. Delfiner. *Geostatistics: Modeling Spatial Uncertainty*. Wiley, 2009. Cited on pages 1, 4.
- M. Cuturi. Sinkhorn distances: lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013. Cited on page 8.
- C. de Fouquet. Reminders on the conditioning Kriging. In *Geostatistical Simulations*, pages 131–145. Springer, 1994. Cited on page 1.
- M. P. Deisenroth, D. Fox, and C. E. Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2015. Cited on page 8.
- C. R. Dietrich and G. N. Newsam. Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM Journal of Scientific Computing*, 18:1088–1107, 1997. Cited on page 9.
- A. Doucet. A note on efficient conditional simulation of Gaussian distributions. Technical report, University of British Columbia, 2010. Cited on page 4.
- R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445, 1961. Cited on pages 9, 16.

- M. Ghavamzadeh, S. Mannor, J. Pineau, and A. Tamar. Bayesian reinforcement learning: a survey. *Foundations and Trends in Machine Learning*, 8(5–6):359–483, 2015. Cited on page 1.
- A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002. Cited on page 6.
- P. E. Gill, G. H. Golub, W. Murray, and M. A. Saunders. Methods for modifying matrix factorizations. *Mathematics of Computation*, 28(126):505–535, 1974. Cited on page 18.
- J. Hensman, N. Durrande, and A. Solin. Variational Fourier features for Gaussian processes. *Journal of Machine Learning Research*, 18(151):1–151, 2017. Cited on page 2.
- J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, pages 282–290, 2013. Cited on page 1.
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems*, pages 918–926, 2014. Cited on page 3.
- J. M. Hernández-Lobato, J. Requeima, E. O. Pyzer-Knapp, and A. Aspuru-Guzik. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In *International Conference on Machine Learning*, pages 1470–1479, 2017. Cited on page 8.
- Y. Hoffman and E. Ribak. Constrained realizations of Gaussian fields: a simple algorithm. *The Astrophysical Journal*, 380:L5–L8, 1991. Cited on pages 1, 4, 5.
- D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993. Cited on page 15.
- A. G. Journel and C. J. Huijbregts. *Mining geostatistics*. Academic Press London, 1978. Cited on pages 1, 4.
- S. Kamthe and M. P. Deisenroth. Data-efficient reinforcement learning with probabilistic model predictive control. In *International Conference on Artificial Intelligence and Statistics*, pages 1701–1710, 2018. Cited on page 8.
- K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos. Parallelised Bayesian optimisation via Thompson Sampling. In *Artificial Intelligence and Statistics*, pages 133–142, 2018. Cited on page 8.
- M. Kuss and C. E. Rasmussen. Gaussian processes in reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 751–758, 2004. Cited on page 2.
- A. Mallasto and A. Feragen. Learning from uncertain curves: The 2-Wasserstein metric for Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 5660–5670, 2017. Cited on page 6.
- M. Mutny and A. Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. In *Advances in Neural Information Processing Systems*, pages 9005–9016, 2018. Cited on pages 1, 3.
- J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the Institute of Radio Engineers*, 50(10):2061–2070, 1962. Cited on pages 9, 16.
- G. Pleiss, J. R. Gardner, K. Q. Weinberger, and A. G. Wilson. Constant-time predictive distributions for Gaussian processes. In *International Conference on Machine Learning*, pages 4114–4123, 2018. Cited on pages 6, 15.
- J. Quiñero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6(12):1939–1959, 2005. Cited on page 2.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2008. Cited on pages 1, 3, 5.
- C. E. Rasmussen and J. Quiñero-Candela. Healing the relevance vector machine through augmentation. In *International Conference on Machine Learning*, pages 689–696, 2005. Cited on page 3.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006. Cited on page 2.
- J. L. Rodgers, W. A. Nicewander, and L. Toothaker. Linearly independent, orthogonal, and uncorrelated variables. *The American Statistician*, 38(2):133–134, 1984. Cited on page 12.
- Y. Saatçi. *Scalable inference for structured Gaussian process models*. PhD thesis, University of Cambridge, 2012. Cited on page 15.
- H. Salimbeni, C.-A. Cheng, B. Boots, and M. P. Deisenroth. Orthogonally decoupled variational Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 8711–8720, 2018. Cited on pages 4, 12.

- B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2001. Cited on page 2.
- M. Seeger. Low rank updates for the Cholesky decomposition. Technical report, 2004. Cited on page 18.
- M. W. Seeger. Bayesian inference and optimal design for the sparse linear model. *Journal of Machine Learning Research*, 9(4):759–813, 2008. Cited on page 3.
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop: a review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015. Cited on pages 1, 3.
- J. Shi, M. K. Titsias, and A. Mnih. Sparse orthogonal variational inference for Gaussian processes. In *Artificial Intelligence and Statistics*, 2020. Cited on pages 4, 12.
- E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pages 1257–1264, 2006. Cited on page 2.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959, 2012. Cited on page 2.
- D. J. Sutherland and J. Schneider. On the error of random Fourier features. In *Uncertainty in Artificial Intelligence*, pages 862–871, 2015. Cited on pages 6, 13, 14.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933. Cited on page 8.
- M. K. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009. Cited on page 2.
- R. Van de Weygaert and E. Bertschinger. Peak and gravity constraints in Gaussian primordial density fields: an application of the Hoffman-Ribak method. *Monthly Notices of the Royal Astronomical Society*, 281(1):84–118, 1996. Cited on page 1.
- K. Wang, G. Pleiss, J. Gardner, S. Tyree, K. Q. Weinberger, and A. G. Wilson. Exact Gaussian processes on a million data points. In *Advances in Neural Information Processing Systems*, pages 14622–14632, 2019. Cited on page 1.
- Z. Wang, C. Gehring, P. Kohli, and S. Jegelka. Batched large-scale Bayesian optimization in high-dimensional spaces. In *Artificial Intelligence and Statistics*, pages 745–754, 2018. Cited on pages 1, 3.
- A. Wilson and H. Nickisch. Kernel interpolation for scalable structured Gaussian processes. In *International Conference on Machine Learning*, pages 1775–1784, 2015. Cited on pages 9, 15.
- J. T. Wilson, F. Hutter, and M. P. Deisenroth. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 9884–9895, 2018. Cited on page 3.
- D. L. Zimmerman. Computationally exploitable structure of covariance matrices and generalized covariance matrices in spatial models. *Journal of Statistical Computation and Simulation*, 32(1-2):1–15, 1989. Cited on page 15.