

Numerically Stable Sparse Gaussian Processes via Minimum Separation using Cover Trees

Alexander Terenin*

University of Cambridge and Imperial College London

David R. Burt*

University of Cambridge and MIT

Artem Artemev*

Imperial College London and Secondmind

Seth Flaxman

University of Oxford

Mark van der Wilk

Imperial College London and University of Oxford

Carl Edward Rasmussen

University of Cambridge and Secondmind

Hong Ge

University of Cambridge

Editor: Mohammad Emtiyaz Khan

Abstract

Gaussian processes are frequently deployed as part of larger machine learning and decision-making systems, for instance in geospatial modeling, Bayesian optimization, or in latent Gaussian models. Within a system, the Gaussian process model needs to perform in a stable and reliable manner to ensure it interacts correctly with other parts of the system. In this work, we study the numerical stability of scalable sparse approximations based on inducing points. To do so, we first review numerical stability, and illustrate typical situations in which Gaussian process models can be unstable. Building on stability theory originally developed in the interpolation literature, we derive sufficient and in certain cases necessary conditions on the inducing points for the computations performed to be numerically stable. For low-dimensional tasks such as geospatial modeling, we propose an automated method for computing inducing points satisfying these conditions. This is done via a modification of the cover tree data structure, which is of independent interest. We additionally propose an alternative sparse approximation for regression with a Gaussian likelihood which trades off a small amount of performance to further improve stability. We provide illustrative examples showing the relationship between stability of calculations and predictive performance of inducing point methods on spatial tasks.

1. Introduction

Gaussian processes are a flexible framework and model class for learning unknown functions. By way of being constructed in the language of Bayesian learning, Gaussian process models provide an ability to incorporate prior information into the model, and assess and propagate uncertainty in a principled manner. This makes them well-suited for a wide variety of areas where these capabilities

*Equal contribution.

Code available at: [HTTPS://GITHUB.COM/AWAV/CONJUGATE-GRADIENT-SPARSE-GP](https://github.com/AWAV/CONJUGATE-GRADIENT-SPARSE-GP).

are important, including statistical applications such as spatial modeling (Cressie, 1992), and decision-making applications such as Bayesian optimization (Snoek et al., 2012), sensor placement (Krause et al., 2008), and active learning (Krause and Guestrin, 2007).

In many settings, the increased availability of data and need to accurately model higher-resolution phenomena has led to a strong interest in working with Gaussian processes at a larger scale. Unfortunately, classical Gaussian process models generally scale cubically with training data size due to the need to solve large linear systems of equations. This mismatch has led to a longstanding and fruitful line of work on scalable Gaussian processes. In the era of GPUs and automatic differentiation, two classes of scalable approximations have been deployed within major Gaussian process software packages, including *GPflow* (Matthews et al., 2017) and *GPyTorch* (Gardner et al., 2018): those based on *inducing point methods* (Quiñero-Candela and Rasmussen, 2005; Titsias, 2009; Hensman et al., 2013), and iterative methods such as the (*preconditioned*) *conjugate gradient* algorithm (Gibbs and Mackay, 1997; Gardner et al., 2018).

Motivated by these advances, in this work we study a complementary question: how can we *guarantee* that the above algorithms run successfully, no matter what kind of data they are provided with? This question is particularly important in areas such as Bayesian optimization where the data is not available in advance, and in latent variable models, where the data going into the Gaussian process is reconstructed from auxiliary observations. We focus chiefly on inducing point methods applied to geospatial or other low-dimensional data, though the techniques we develop also reveal insights on how the performance of conjugate-gradient-based approaches depends on data.

Our work complements existing analyses of numerical stability in Gaussian processes. Foster et al. (2009) study the subset of regressors approximation, and describe in detail which numerical linear algebra routines one should use to implement it in the most stable manner. Basak et al. (2021) study maximum likelihood estimation, observe that stability of log-determinant computations is effectively controlled by the same quantities that control stability of linear solves, and describe practical strategies for initializing and stopping numerical optimization in a stable manner. In contrast, we describe how different Gaussian process approximations and their hyperparameters affect the numerical stability of the linear-algebraic routines which need to be implemented, regardless of the details of how this implementation is done.

Our contributions are as follows. We (i) identify that mathematical tools developed in the interpolation literature (Narcowich and Ward, 1992; Narcowich et al., 1994; Schaback, 1995; Diederichs and Iske, 2019) can be used to show, in a large class of Gaussian process priors, that datasets satisfying a *minimum separation* criterion lead to linear systems which are provably stable independent of data size. Using this insight, we (ii) develop a clustering analogy and propose a technique to guarantee stability in sparse approximations by using a refinement of the *cover tree* data structure (Beygelzimer et al., 2006) to automatically select inducing points that satisfy the needed criterion while staying as close to the true data as possible. This approximation is parameterized by a hyperparameter called the *spatial resolution*, which directly controls the tradeoff between approximation accuracy and computational cost. Following this, we (iii) extend the clustering analogy further, and propose an alternative, non-variational inducing point approximation which trades off a small amount of performance to further improve stability. Finally, we (iv) study error and computational complexity of the proposed methods, and evaluate them on a number of examples.

2. Gaussian Processes

Let X be a set. We say that a random function $f : X \rightarrow \mathbb{R}$ is a *Gaussian process* if, for any finite set of points $\mathbf{x} \in X^N$, the random vector $f(\mathbf{x})$ is multivariate Gaussian. We write $f \sim \text{GP}(\mu, k)$, where $\mu(\mathbf{x}) = \mathbb{E}(f(\mathbf{x}))$ is the *mean function*, and $k(\mathbf{x}, \mathbf{x}') = \text{Cov}(f(\mathbf{x}), f(\mathbf{x}'))$ is the *covariance kernel*, which determine the mean and covariance of the corresponding multivariate Gaussians.

Define the likelihood $\mathbf{y} = f(\mathbf{x}) + \boldsymbol{\varepsilon}$ where $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. Given training data \mathbf{x}, \mathbf{y} , if we place a Gaussian process prior $f \sim \text{GP}(0, k)$, then the posterior $f \mid \mathbf{y}$ is a Gaussian process with

$$\mathbb{E}(f \mid \mathbf{y}) = \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \boldsymbol{\Sigma})^{-1}\mathbf{y} \quad \text{Cov}(f \mid \mathbf{y}) = \mathbf{K}_{(\cdot,\cdot)} - \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \boldsymbol{\Sigma})^{-1}\mathbf{K}_{\mathbf{x}(\cdot)} \quad (1)$$

where (\cdot) and (\cdot') denote arbitrary sets of test locations and for sets a, b we use \mathbf{K}_{ab} to denote the kernel matrix formed by evaluating k at points in $a \times b$. We assume throughout that $\boldsymbol{\Sigma}$ is diagonal. It is also possible to express the posterior in the form of *pathwise conditioning* (Wilson et al., 2020, 2021) as

$$(f \mid \mathbf{y})(\cdot) = f(\cdot) + \mathbf{K}_{(\cdot)\mathbf{x}}(\mathbf{K}_{\mathbf{x}\mathbf{x}} + \boldsymbol{\Sigma})^{-1}(\mathbf{y} - f(\mathbf{x}) - \boldsymbol{\varepsilon}) \quad (2)$$

where equality holds in distribution. These expressions describe the true posterior, whose computational costs are, classically, $\mathcal{O}(N^3)$ owing to the Cholesky decomposition used to handle the matrix inverse. To alleviate this, we now consider approximations.

2.1 Inducing Points

Many of the most widely-used Gaussian process approximations are based on the idea of *inducing points* (Csato and Opper, 2002; Seeger et al., 2003; Quionero-Candela and Rasmussen, 2005). We present the variational formulation of Titsias (2009) and Hensman et al. (2013), using a pathwise conditioning construction (Wilson et al., 2020, 2021). Let $\mathbf{z} \in X^M$, and define the *variational posterior*

$$(f \mid \mathbf{u})(\cdot) = f(\cdot) + \mathbf{K}_{(\cdot)\mathbf{z}}\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}(\mathbf{u} - f(\mathbf{z})) \quad \mathbf{u} \sim \mathcal{N}(\mathbf{m}, \mathbf{S}) \quad (3)$$

where \mathbf{m}, \mathbf{S} and \mathbf{z} are approximation parameters. We find these by minimizing the Kullback–Leibler divergence of the approximation from the true posterior. By the chain rule for Kullback–Leibler divergences, this quantity reduces to an easy-to-evaluate divergence between finite-dimensional Gaussians (Matthews et al., 2016). Given \mathbf{z} , which one can interpret as a learned set of approximate \mathbf{x} -values, the optimal values for \mathbf{m} and \mathbf{S} can be solved for analytically in closed form: see Titsias (2009) for details.

This approximation only requires factorizing a smaller kernel matrix, leading to a cost of $\mathcal{O}(M^3)$, a significant improvement if $M \ll N$. Following Hensman et al. (2013), the approximation also inherits variational methods’ compatibility with stochastic optimization and mini-batch methods.

2.2 Gaussian Processes for Geospatial Modelling

Gaussian processes models are widely used in spatial statistics both for interpolation, and to allow for spatially correlated noise in spatial regression models (Cressie, 1992). In low-dimensional settings, such as geospatial data, inducing point approximations can lead to provably near-linear computational scalability, particularly when input points, x_i , are sampled from within a bounded region, and the Gaussian process prior is smooth enough—see Burt et al. (2019, 2020). This makes them a reasonable choice for many applications.

In geospatial modeling, specialized methods including multi-resolution approaches (Katzfuss, 2017) or sparse-precision approaches (Vecchia, 1988; Lindgren et al., 2011; Katzfuss and Guinness, 2021) often result in accurate approximations, especially when the domain is much larger than the scale over which the function being modelled varies. See for instance Heaton et al. (2019) for a comparison of a large number of approaches on problems in this class. In such settings, low-rank calculations very similar to those used in inducing point methods remain central to the construction of a number of more tailored approaches, including in particular multi-resolution methods (Katzfuss, 2017) and combinations of low-rank and compactly supported methods (Sang and Huang, 2012). A better understanding of numerical stability in inducing point methods is therefore an important step to better understanding when these more sophisticated methods designed specifically for spatial data will be numerically stable.

2.3 Numerical Stability

On a computer, the calculations needed to deploy a Gaussian process, whether exact or approximate, must be performed in floating-point arithmetic. Due to the introduction of roundoff error, algorithms which solve linear systems $\mathbf{A}^{-1}\mathbf{b}$ can fail if the system's solution is too sensitive to the numerical values of \mathbf{A} or \mathbf{b} . We assume throughout that \mathbf{A} is symmetric positive definite, as will be the case for the kernel matrices of interest. The key quantity used to understand how numerically stable a linear system is the associated *condition number*

$$\text{cond}(\mathbf{A}) = \lim_{\varepsilon \rightarrow 0} \sup_{\|\delta\| \leq \varepsilon \|\mathbf{b}\|} \frac{\|\mathbf{A}^{-1}(\mathbf{b} + \delta) - \mathbf{A}^{-1}\mathbf{b}\|_2}{\varepsilon \|\mathbf{A}^{-1}\mathbf{b}\|_2} = \|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 = \frac{\lambda_{\max}(\mathbf{A})}{\lambda_{\min}(\mathbf{A})} \quad (4)$$

of the matrix defining the system. Here, λ_{\max} and λ_{\min} are the maximum and minimum eigenvalues, respectively, and $\|\cdot\|_2$ denotes the Euclidean norm and the corresponding induced operator norm. A linear system's condition number quantifies how difficult it is to solve numerically. For a given floating-point precision, if $\text{cond}(\mathbf{A})$ is small enough and the size of \mathbf{A} is not too large, then Cholesky factorization is guaranteed to succeed and return an accurate matrix square root.

Result 1. *Let \mathbf{A} be a symmetric positive definite matrix of size $N \times N$. Assume that $N > 10$, that*

$$\text{cond}(\mathbf{A}) \leq \frac{1}{2^{-t} \times 3.9N^{3/2}} \quad (5)$$

where t is the length of the floating point mantissa, and that $3N2^{-t} < 0.1$. Then floating point Cholesky factorization will succeed, producing a matrix \mathbf{L} satisfying

$$\mathbf{L}\mathbf{L}^T = \mathbf{A} + \mathbf{E} \quad \|\mathbf{E}\|_2 \leq 2^{-t} \times 1.38N^{3/2} \|\mathbf{A}\|_2. \quad (6)$$

Proof. This follows from Kiełbasiński (1987, Corollary 2), and the relationship between the Euclidean and Frobenius matrix norms. See also Wilkinson (1966, Theorem 2). \square

For single precision arithmetic, we have $2^{-t} \approx 10^{-7.2}$, and for double-precision arithmetic, we have $2^{-t} \approx 10^{-16}$. We therefore see that well-conditioned systems of equations lead to numerically stable Cholesky factorizations. Moreover, when an iterative algorithm is used to solve a well-conditioned system, the iteration is often guaranteed to converge quickly, leading to computational benefits. In particular, if one solves the linear system with the conjugate gradient algorithm (Golub and Van Loan, 1996), then the algorithm is guaranteed to converge in logarithmically many steps.

Result 2. *Let \mathbf{A} be a positive semi-definite matrix of size $N \times N$, and let $\varepsilon > 0$. Then, in exact arithmetic, the conjugate gradient algorithm for solving $\mathbf{A}^{-1}\mathbf{b}$ converges to within ε of the true solution, with respect to the norm induced by \mathbf{A} , in*

$$\mathcal{O}\left(\sqrt{\text{cond}(\mathbf{A})} \log \frac{\text{cond}(\mathbf{A}) \|\mathbf{b}\|}{\varepsilon}\right) \quad (7)$$

total steps.

Proof. The claim follows from Golub and Van Loan (1996, Theorem 10.2.6) by a direct calculation given in Appendix A.2. \square

Working with stable linear systems is therefore crucial for both inducing point methods which rely on Cholesky factorizations, and for algorithms based on conjugate gradients. In the latter case, *preconditioning* is often used to improve stability and accelerate convergence. Preconditioning can be viewed as a way to partially solve the linear system, resolving its unstable components. State-of-the-art preconditioners in many areas of applied mathematics, such as multigrid methods for elliptic partial differential equations (E, 2011; Xu and Zikatanov, 2017), rely on detailed properties of their respective settings for construction and analysis. A key step in designing provably effective preconditioners for Gaussian processes is therefore to understand how the model and data influence their resulting linear systems: we focus on this and defer ideas on preconditioning to future work.

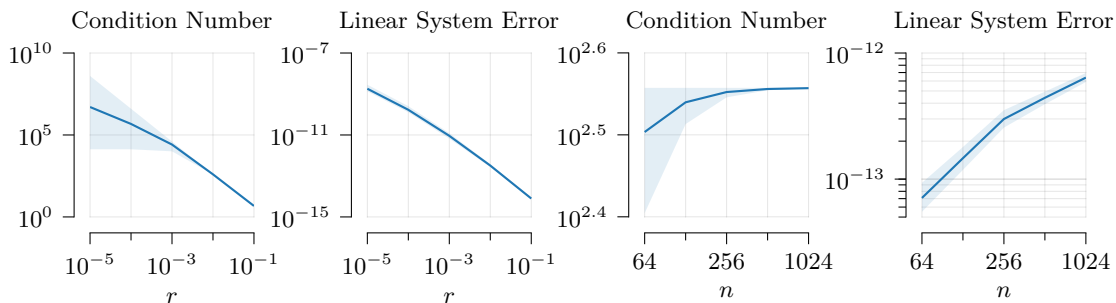


Figure 1: Here we illustrate the condition numbers and resulting linear system error for the Kac–Murdock–Szegő matrix, which is the kernel matrix of an exponential kernel on a regularly-spaced one-dimensional grid. We vary the spacing r and respective correlation $\rho = \exp(-r)$ of neighboring points, and the size n of the matrix, with $\rho = 0.999$ used in cases where n varies, and $n = 256$ used in cases where ρ varies. We plot the condition number, computed numerically using an eigenvalue factorization, along with its theoretical lower and upper bounds. Then, we generate random vectors $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, compute $\mathbf{u} = \mathbf{K}_{\mathbf{x}\mathbf{x}}\mathbf{v}$, solve for $\mathbf{v} = \mathbf{K}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{u}$ numerically, and plot the median error norm over 10^4 samples, along with 25% and 75% quantiles. We observe that condition numbers asymptotically grow as $\rho \rightarrow 1$, but not as $N \rightarrow \infty$, and increase hand-in-hand with numerical error.

2.4 Instability in Gaussian Process Models

To begin, we first observe that, generically, condition numbers of kernel matrices need not be well-behaved. To illustrate what can go wrong, consider the one-dimensional exponential kernel

$$k(x, x') = \exp(-|x - x'|). \quad (8)$$

Suppose that x is a time series which lies on a regular one-dimensional grid with spacing r , so that $x_i = ri$. Then the resulting kernel matrix is a *Kac–Murdock–Szegő matrix* (Trench, 2001; Dow, 2002), meaning it takes the form

$$\mathbf{K}_{\mathbf{x}\mathbf{x}} = \begin{pmatrix} 1 & \rho & \rho^2 & \dots & \rho^{n-1} \\ \rho & 1 & \rho & \dots & \rho^{n-2} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \rho^{n-1} & \rho^{n-2} & \rho^{n-3} & \dots & 1 \end{pmatrix} \quad (9)$$

for $\rho = \exp(-r)$. To compute a posterior Gaussian process under exact observations, the condition number of the matrix we need to invert to compute the posterior satisfies (Trench, 2001, p. 9) the inequality

$$\frac{(1 + \rho)^2}{(1 - \rho)^2} \leq \text{cond}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \leq \frac{(1 + \rho)^2 + 2\rho\epsilon}{(1 - \rho)^2 - 2\rho\epsilon} \quad (10)$$

where $\epsilon = \frac{\pi^2}{(N+1)^2}$ and for the upper bound we have assumed $(1 - \rho)^2 > 2\rho\epsilon$. For large N , the value ϵ is small so the upper and lower bounds essentially match. As $r \rightarrow 0$, we have $\rho \rightarrow 1$, and the condition number diverges to infinity for any given value of N . On the other hand, if the spread of the data, r is fixed, then regardless of how much data is collected, ρ and the condition number of $\mathbf{K}_{\mathbf{x}\mathbf{x}}$ both remain bounded. In cases where data is observed under a Gaussian likelihood, one instead needs to invert $\mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{\Sigma}$. If $\mathbf{\Sigma}$ is diagonal, the minimum eigenvalue of $\mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{\Sigma}$ is bounded below by the minimum diagonal entry of $\mathbf{\Sigma}$, which we will refer to as σ^2 by analogy to the homoscedastic noise case when $\mathbf{\Sigma} = \sigma^2\mathbf{I}$. Unfortunately, this is not enough to ensure stability: if one samples the time series more densely, then as $r \rightarrow 0$ and $N \rightarrow \infty$, the condition number diverges to infinity for all

values of σ^2 due to the growth of the largest eigenvalue of $\mathbf{K}_{\mathbf{x}\mathbf{x}}$. We now observe that the behavior illustrated by this example is neither specific to the exponential kernel, nor to data with algebraic structure arising from grids, nor limited to one dimension.

Proposition 3. *Let k be a continuous stationary kernel, and let the entries in \mathbf{x} be independently sampled from a uniform distribution on some compact subset $X \subseteq \mathbb{R}^d$. Define the covariance operator*

$$\mathcal{K} : L^2(X) \rightarrow L^2(X) \qquad \mathcal{K} : \phi \mapsto \int_X \phi(x)k(x, \cdot) dx \qquad (11)$$

acting on the Hilbert space of (equivalence classes of) square-integrable functions. Then the eigenvalues of \mathcal{K} are countable, non-negative, admit a maximum, and can be ordered to form a non-increasing sequence whose limit is zero. For every $n > 0$, as $N \rightarrow \infty$, the n^{th} largest eigenvalue of $\frac{1}{N}\mathbf{K}_{\mathbf{x}\mathbf{x}}$ converges almost surely to the corresponding eigenvalue of \mathcal{K} . As a consequence, $\text{cond}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \rightarrow \infty$ almost surely.

Proof. This follows from Koltchinskii and Giné (2000, Corollary 3.3), after a bit of technical reformulation and some calculations, which we perform in Appendix A. \square

Proposition 3 implies that λ_{\max} grows linearly with N , while Cauchy’s Interlacing Theorem together with an elementary argument shows that $\lambda_{\min} \rightarrow 0$, causing the condition number of $\mathbf{K}_{\mathbf{x}\mathbf{x}}$ to grow without bound. From the one-dimensional example and Proposition 3, we therefore see that to control the condition numbers of kernel matrices, we need to ensure they are sufficiently far away from this ill-conditioned limiting case. We conclude that if the data is sampled too closely relative to the kernel’s length scale, computing a Cholesky factorization of $\mathbf{K}_{\mathbf{x}\mathbf{x}} + \Sigma$ numerically is not possible for large matrices. We proceed to explore this in more general cases in the sequel.

3. Numerical Stability in Scalable Gaussian Process Approximations

To build towards Gaussian process approximations which are numerically stable, we study what properties of models and data lead to numerical stability of kernel matrices. We focus in particular on how their condition numbers depend on the number of input points under general regimes.

3.1 Numerical Stability in Sparse Approximations via Minimum Separation

In both exact and sparse Gaussian processes, the linear systems that need to be solved arise from kernel matrices $\mathbf{K}_{\mathbf{x}\mathbf{x}}$ and $\mathbf{K}_{\mathbf{z}\mathbf{z}}$, which depend on the prior and data. In sparse Gaussian processes, however, \mathbf{z} is not fixed: it is a variational parameter, and is usually selected to maximize approximation accuracy. We now study whether we can use this freedom to select \mathbf{z} in a way that ensures stability, while maintaining performance. To this end, we introduce the following notion.

Definition 4. *Let \mathbf{z} be the inducing points. Define the SEPARATION RADIUS*

$$\text{sep}(\mathbf{z}) = \min_{i \neq j} \|z_i - z_j\|. \qquad (12)$$

We now show that separation is closely connected with numerical stability. For this, we need a mild regularity condition on the kernel.

Assumption 5 (Spatial decay). *Let k be a kernel on $X \subseteq \mathbb{R}^d$. We say that k has SPATIAL DECAY if there is a decreasing function $\psi : [0, \infty) \rightarrow [0, \infty)$ such that for all $x, x' \in X$ we have*

$$|k(x, x')| \leq \psi(\|x - x'\|) \qquad \psi(m) = \mathcal{O}\left(\frac{1}{m^d \log(m)^2}\right). \qquad (13)$$

In words, spatial decay requires that if we consider points far apart, the covariance between the function at those points must tend to zero sufficiently quickly with the distance between them. This assumption is satisfied for a very large set of kernels, including potentially non-stationary kernels, as well as the squared exponential kernel and all Matérn kernels, whose decay is controlled above and below by products of an exponential with polynomials. We now state the main results.

Proposition 6. *Let $X \subseteq \mathbb{R}^d$, and let k satisfy spatial decay. Then there is a constant $C_{\max}^{k,\delta}$ such that for any M and any \mathbf{z} of size M with $\text{sep}(\mathbf{z}) \geq \delta > 0$, we have*

$$\lambda_{\max}(\mathbf{K}_{\mathbf{z}\mathbf{z}}) \leq C_{\max}^{k,\delta}. \quad (14)$$

Proof. Appendix A, Proposition 24. □

From a technical perspective, the claim is proven by combining Gershgorin’s Circle Theorem with a packing argument. Arguments of this kind have a long history in the interpolation literature—see for instance Narcowich and Ward (1992), Narcowich et al. (1994), Schaback (1995), and Diederichs and Iske (2019). In cases where the Gaussian process approximation requires inversion of $\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda}$, from this we can immediately conclude a condition number bound.

Corollary 7. *Under the conditions of Proposition 6, for diagonal $\mathbf{\Lambda}$, letting Λ_{\max} and Λ_{\min} denote its respective maximum and minimum entry, we have*

$$\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda}) \leq \frac{C_{\max}^{k,\delta} + \Lambda_{\max}}{\Lambda_{\min}}. \quad (15)$$

Proof. Combine Proposition 6 with the triangle inequality for operator norm, noting that for positive semi-definite matrices the operator norm is equal to the largest eigenvalue. □

This generalizes immediately to the case where $\mathbf{\Lambda}$ is not diagonal, as long as Λ_{\max} and Λ_{\min} are replaced with its maximum and minimum eigenvalues. Since in many practical cases $\mathbf{\Lambda}$ is diagonal, we retain this presentation for simplicity. In the noiseless case, we can also conclude a condition number bound, but now additionally require stationarity to do so. We will need the following result.

Result 8. *Let $X \subseteq \mathbb{R}^d$, and let k be stationary and continuous. Then there is a constant $C_{\min}^{k,\delta} > 0$ such that for any M and any \mathbf{z} of size M with $\text{sep}(\mathbf{z}) \geq \delta > 0$, we have*

$$\lambda_{\min}(\mathbf{K}_{\mathbf{z}\mathbf{z}}) \geq C_{\min}^{k,\delta}. \quad (16)$$

Proof. Wendland (2004, Theorem 12.3). See Appendix A for a sketch of the main argument. □

We can combine this the eigenvalue upper bound to bound the condition number, which we present for completeness.

Corollary 9. *Under the conditions of Proposition 6 and Result 8, there is a constant $C_{\text{cond}}^{k,\delta} > 1$ such that for any M and any \mathbf{z} of size M with $\text{sep}(\mathbf{z}) < \delta$, we have*

$$\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}}) \leq C_{\text{cond}}^{k,\delta}. \quad (17)$$

Proof. Combine Proposition 6 and Result 8 with $\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}}) = \frac{\lambda_{\max}(\mathbf{K}_{\mathbf{z}\mathbf{z}})}{\lambda_{\min}(\mathbf{K}_{\mathbf{z}\mathbf{z}})}$, taking $C_{\text{cond}}^{k,\delta} = \frac{C_{\max}^{k,\delta}}{C_{\min}^{k,\delta}}$. □

These results therefore reveal what properties the set of inducing points \mathbf{z} needs to have in order to yield numerically stable linear systems. Separation distance and spatial decay rule out the bad behavior shown in Proposition 3, by preventing the limiting covariance operator from controlling properties of the corresponding kernel matrices. The analysis shown is based on the techniques presented by Wendland (2004, Ch. 12) for proving Result 8, which originate in the polynomial

interpolation literature (Narcowich and Ward, 1992; Schaback, 1995). For non-stationary kernels, mirroring Proposition 6, we conjecture that a similar bound to Result 8 for λ_{\min} holds for a wide class of kernels, but requires a different argument which avoids the use of Fourier analysis. To conclude, we observe that for most kernels, if one allows for an arbitrarily large number of inducing points and no noise is added to the kernel matrix, then the minimum separation condition is necessary to ensure bounded condition numbers.

Proposition 10. *Let $z_m \in X$, $m = 1, \dots, \infty$, be a sequence, and let k be a Lipschitz continuous kernel on $X \times X$. Define $\mathbf{K}_{zz}^{(M)} \in \mathbb{R}^{M \times M}$ to be the kernel matrix formed by evaluating k at z_1, \dots, z_M . If there exists a $C > 0$ such that for all M , $\text{cond}(\mathbf{K}_{zz}^{(M)}) \leq C$, then z_m satisfy minimum separation.*

Proof. Appendix A.5. □

3.2 Inducing Points for Geospatial Data via Cover Trees

To obtain an algorithmic solution that guarantees numerical stability, we study how to construct inducing points which satisfy a user-specified minimum separation radius. Since separation is defined purely in terms of the inducing points, we need a second criterion to quantify how well the inducing points summarize a given dataset. A clear choice is to place the inducing points as close to the data as possible.

Definition 11. *Let \mathbf{x} be the data, and \mathbf{z} be the inducing points. Define the SPATIAL RESOLUTION*

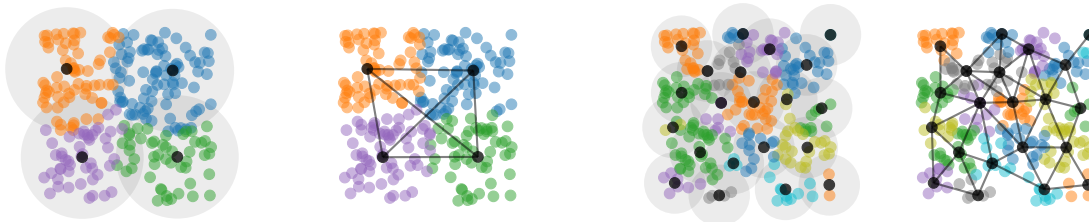
$$\text{res}_{\mathbf{x}}(\mathbf{z}) = \max_{i=1, \dots, N} \min_{j=1, \dots, M} \|x_i - z_j\|. \quad (18)$$

In simple terms, the spatial resolution is the maximum distance from a data point to the closest inducing point. This notion is closely related to *fill distance* which appears in the the analysis of posterior contraction rates (Kanagawa et al., 2018), but instead involves both a dataset \mathbf{x} and inducing points \mathbf{z} . Among all sets of inducing points with a given spatial resolution, finding ones which satisfy separation is always possible. More precisely, if we take the inducing points \mathbf{z} to be the centers of balls which make up a maximal separated subset of radius $\text{sep}(\mathbf{z}) = \varepsilon$, the resulting choice automatically has spatial resolution $\text{res}_{\mathbf{x}}(\mathbf{z}) = \varepsilon$. A similar observation applies if one takes \mathbf{z} to be an optimal covering of \mathbf{x} —see Appendix A for details. This shows that the notions of separation distance and spatial resolution are closely connected with optimal covering and packing numbers—see Vershynin (2018, Section 4.2).

We now study numerical techniques for computing inducing points that simultaneously satisfy separation and covering guarantees. Our goal is to obtain a practical near-linear-time method producing a set of inducing points with small spatial resolution, to ensure accurate approximations of the posterior, and large separation to ensure stability. To achieve this, we propose an approach we term the *R-neighbor cover tree*, which is a refinement of the *cover tree* data structured originally proposed by Beygelzimer et al. (2006) for efficient nearest neighbor search, and studied further by Kibriya and Frank (2007) and Izbicki and Shelton (2015). This approach is well-suited to cases where the inputs are low-dimensional such as geospatial data. We note that the proposed algorithm will also perform well in cases when the dataset has low *intrinsic dimension*—see Beygelzimer et al. (2006)—even if the dataset is embedded in a high dimensional space.

A cover tree is a tree where nodes are points associated with metric balls. The root node is a ball which covers all of the data \mathbf{x} . Each parent node contains the centers of its child nodes, which consist of balls of smaller radius than the parent. At each level, the tree enforces separation and covering properties similar to those in Definition 4 and Definition 11, with the separation distance and spatial resolution proportional to the radius of balls at that depth. Beygelzimer et al. (2006) present a depth-first algorithm for building such a tree.

Since we are interested in inducing point selection with a fixed spatial resolution ε , rather than nearest neighbors search, we seek a different set of guarantees than those provided by Beygelzimer



(a) Place covering nodes (b) Compute R -neighbors (c) Repeat steps (a) and (b) recursively

Figure 2: Two iterations of Algorithm 12 on a tree with $L = 3$ total levels. Given a region, the algorithm first (a) computes a covering of the region by picking data points one-by-one from those not yet covered. Then, it (b) computes the R -neighbors of each node and uses this to efficiently compute a Voronoi tessellation of the region. Finally, the algorithm (c) repeats the process recursively.

et al. (2006), or by improved versions such as that of Izbicki and Shelton (2015). The most important difference is that, in the inducing point case, we want separation to hold *globally* in every level of the tree, and not only for children with a common parent. To do this, we present a modified breadth-first construction, which is of independent interest, and conceptually works as follows.

1. Initialize the root node $z_{1,1}$ at the mean of the data, and assign all data to it.
2. Loop over tree depth ℓ , starting at $\ell = 0$.
 - (a) Loop over parent nodes $z_{\ell,p}$.
 - i. Select a point z' from the parent's assigned data.
 - ii. *Optional*: compute the local average of the parent node's assigned data around z' , and set z' to be this average, as long as it is not too close to another node.
 - iii. Create a child node $z_{\ell+1,c}$ centered at z' .
 - iv. Reassign all points near $z_{\ell+1,c}$ in nodes at level ℓ to $z_{\ell+1,c}$.
 - (b) *Optional*: reassign all points at level $\ell + 1$ according to the Voronoi partition of $z_{\ell+1,c}$.

A formal description is given in Algorithm 12. The key insight behind this top-down, breadth-first construction is that, for a dataset with a given intrinsic dimensionality, it can be performed in near-linear time. The reason for this is that every step is *local*, and can be computed by searching only the children of nodes sufficiently close to the current parent node. More precisely, we show that if the parent node's radius is at level ℓ with radius R and the total depth of the tree is L , then only nodes within a distance of $4(1 - 1/2^{(L-\ell)})R \leq 4R$ need to be searched. The number of such nodes is bounded, can be tracked recursively, and is controlled by the intrinsic dimensionality of the data.

The two optional steps do not alter the algorithm's complexity or its guarantees, but improve practical performance. The local averaging step we propose is similar to the Lloyd's iteration of K -means, and allows the tree to use fewer nodes by placing them in-between data points, rather than exactly on top of data points. The Voronoi repartitioning step, originally proposed by Izbicki and Shelton (2015), makes the tree better-balanced. We now prove the algorithm works as intended.

Theorem 13. *For a given target spatial resolution $\varepsilon > 0$, and dataset with $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ and $d_{\max} = \max_i \|x_i - \bar{x}\|$, the cover tree inducing point algorithm, with or without the optional steps, produces a tree with $L = \lceil \log_2 \frac{d_{\max}}{\varepsilon} \rceil + 1$ levels, terminates in*

$$\mathcal{O}\left(P_{B(4)}^{\text{ext},1/2} \left(P_{B(1)}^{\text{ext},1/2}\right)^2 N \log \frac{d_{\max}}{\varepsilon}\right) \quad (19)$$

Algorithm 12 Cover Tree Inducing Points

- 1: **input** spatial resolution $\varepsilon > 0$ and dataset \mathbf{x} . Define notation $B_z(R) = \{x \in X : \|z - x\| \leq R\}$.
 - 2: Initialize root node $z_{0,1} = \frac{1}{N} \sum_{i=1}^N x_i$, assigned data $\mathcal{A}_{1,1} = \mathbf{x}$, R -neighbors $\mathcal{R}_{1,1} = \{1\}$, as well as constants $d_{\max} = \max_{i=1, \dots, N} \|z_{1,1} - x_i\|$, $L = \lceil \log_2 \frac{d_{\max}}{\varepsilon} \rceil$, $M_0 = 1$, and $R_0 = 2^L \varepsilon$.
 - 3: **for** tree depth level $\ell = 1, \dots, L$ **do**
 - 4: Initialize number of nodes $M_\ell = 0$, child node index $c = 1$, and radius $R_\ell = \frac{1}{2} R_{\ell-1}$.
 - 5: **for** parent node index $p = 1, \dots, M_{\ell-1}$ **do**
 - 6: Initialize child nodes $\mathcal{C}_{\ell-1,p} = \emptyset$ and, *optionally*, assigned data copy $\mathcal{A}'_{\ell-1,p} = \mathcal{A}_{\ell-1,p}$.
 - 7: **while** assigned data $\mathcal{A}_{\ell-1,p} \neq \emptyset$ **do**
 - 8: Choose an arbitrary data point $\zeta \in \mathcal{A}_{\ell-1,p}$.
 - 9: *Optional*: compute local average $\zeta' = \frac{1}{|\mathcal{Z}|} \sum_{z \in \mathcal{Z}} z$, where $\mathcal{Z} = \mathcal{A}_{\ell-1,p} \cap B_\zeta(R_\ell)$.
 - 10: **if** $\min_{z \in \mathcal{Z}'} \|\zeta - z\| > R_\ell$ where $\mathcal{Z}' = \bigcup_{r \in \mathcal{R}_{\ell-1,p,\ell,c}} \mathcal{C}_{\ell-1,p}$ **then** set $\zeta = \zeta'$.
 - 11: Create child node $z_{\ell,c} = \zeta$.
 - 12: **for** parent R -neighbor index r in $\mathcal{R}_{\ell-1,p}$ **do**
 - 13: Update assigned data $\mathcal{A}_{\ell,c} = \mathcal{A}_{\ell,c} \cup (\mathcal{A}_{\ell-1,r} \cap B_{z_{\ell,c}}(R_\ell))$, $\mathcal{A}_{\ell-1,r} = \mathcal{A}_{\ell-1,r} \setminus B_{z_{\ell,c}}(R_\ell)$.
 - 14: Update children $\mathcal{C}_{\ell-1,p} = \mathcal{C}_{\ell-1,p} \cup \{c\}$, parent $p_{\ell,c} = p$, index $c = c + 1$, and $M_\ell = M_\ell + 1$.
 - 15: **for** child node index $\varsigma = 1, \dots, c$ **do**
 - 16: Compute child R -neighbors $\mathcal{R}_{\ell,\varsigma} = \bigcup_{r \in \mathcal{R}_{\ell-1,p,\ell,c}} \mathcal{C}_{\ell-1,r} \cap B_{z_{\ell,c}}(4(1 - 1/2^{L-\ell})R_\ell)$.
 - 17: *Optional*: recompute $\mathcal{A}_{\ell,\varsigma} = \{x \in \bigcup_{r \in \mathcal{R}_{\ell-1,p,\ell-1,c}} \mathcal{A}'_{\ell,r} : \|x - z_\varsigma\| \leq \|x - z_{\varsigma'}\|, \forall \varsigma' \neq \varsigma\}$.
 - 18: **return** $\mathbf{z} = \{z_{\ell,m} : \ell = L\}$.
-

steps, where $B(r)$ denotes a ball of radius r , and $P_X^{\text{ext},\delta}$ is the external packing number, namely the maximum number of disjoint balls of radius δ that one can choose so that their centers lie in X . Moreover, the algorithm guarantees that the nodes at each level ℓ satisfy

$$\text{res}_{\mathbf{x}}(\mathbf{z}_\ell) \leq 2^{L-\ell} \varepsilon \qquad \text{sep}(\mathbf{z}_\ell) \geq 2^{L-\ell} \varepsilon. \quad (20)$$

Proof. Appendix C. □

This gives an efficient and practical algorithm for choosing inducing points which guarantees the properties needed in Corollary 9 hold. The algorithm's output is fairly-sharp: compared to an optimal covering, the cover tree construction may require a slightly larger number of inducing points, but achieves the same separation guarantee, irrespective of the number of levels. Using this, we obtain the following numerical stability guarantee.

Corollary 14. *Let \mathbf{x} be a dataset of size N , and let $\varepsilon > 0$. Let \mathbf{z} be a set of inducing points computed from \mathbf{x} by the cover tree algorithm with target spatial resolution ε . Let k be stationary, continuous, and satisfy spatial decay. Then $\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}})$ is bounded independently of N .*

Proof. Combine Theorem 13 with Corollary 9. □

For kernels which possess a length scale parameter, the desired spatial resolution in a given problem might depend on the value of the length scale. This parameter is often set by training the model via maximum marginal likelihood, and can therefore change as training progresses. In such settings, one can use the hierarchical structure of the tree to dynamically adjust the number of inducing points as the length scale changes. This can be done by switching to inducing points obtained from an intermediate level of the tree if the length scale becomes coarser, or by computing an extra level of the tree if it becomes finer. The simplest way to do this is to adjust the model after initial training is complete, since this avoids non-smooth changes in the inducing point approximation during training, and fine-tune the model after the update if necessary.

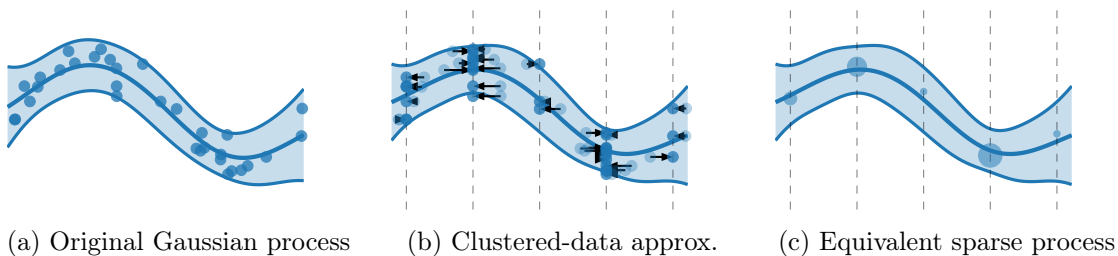


Figure 3: The idea behind the clustered-data approximation of Proposition 15. To construct this approximation, we move each data point in \mathbf{x} to its nearest cluster, transforming (a) to (b). We can then reinterpret and represent (b) in a sparse manner by merging repeated data points, replacing \mathbf{y} with the cluster means \mathbf{u} , and re-weighting each data point according to cluster size—obtaining (c), which is equal in distribution to (b).

3.3 Additional Stability via the Clustered-data Inducing Point Approximation

In the preceding sections, we studied properties of data needed to obtain a stability guarantee for inducing point methods. We now study a complementary question: if selecting \mathbf{z} using a cover tree, can the inducing point approximation be modified to improve stability further? To proceed, we begin by defining a low-rank analog (Bui et al., 2017; Panos et al., 2018; Adam et al., 2021) of the Opper and Archambeau (2009) approximation

$$(f | \mathbf{u})(\cdot) = f(\cdot) + \mathbf{K}_{(\cdot)\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}(\mathbf{u} - f(\mathbf{z}) - \boldsymbol{\epsilon}) \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda}) \quad (21)$$

where \mathbf{z} , diagonal noise matrix $\mathbf{\Lambda}$, and the inducing mean value \mathbf{u} are now variational parameters. Since its optimization objective is non-convex in the variational parameters, this approximation is generally used with natural gradients (Adam et al., 2021) or stochastic optimization (Panos et al., 2018; van der Wilk et al., 2022). Similar to the earlier inducing point approximation of Titsias (2009) and Hensman et al. (2013), it recovers the true posterior when $M = N$, $\mathbf{z} = \mathbf{x}$, $\mathbf{\Lambda} = \mathbf{\Sigma}$, and $\mathbf{u} = \mathbf{y}$.

This approximation is more numerically stable thanks to the diagonal matrix $\mathbf{\Lambda}$ (Panos et al., 2018; van der Wilk et al., 2022), which lower bounds the minimum eigenvalue of $\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda}$ by $\min_{i=1,\dots,M} \Lambda_{ii} + \lambda_{\min}(\mathbf{K}_{\mathbf{z}\mathbf{z}}) > 0$. For finite datasets, this is a strict improvement in stability over approximations involving only $\mathbf{K}_{\mathbf{z}\mathbf{z}}$. Since $\mathbf{\Lambda}$ effectively controls the width of the posterior error bars, the degree to which this benefit manifests itself depends on the data, with more uncertainty leading to more numerically stable linear systems.

Before studying convergence further, we observe that this approximation’s relationship with the true posterior is deeper than might first meets the eye. For a given set of inducing points \mathbf{z} and dataset \mathbf{x} , define the *nearest inducing point clustering* and *cluster size* maps

$$\text{cl}(x) = \arg \min_{z_j, j=1,\dots,M} \|x - z_j\| \quad N_{\text{cl}}(x) = |\{x' \in \mathbf{x} : \text{cl}(x) = \text{cl}(x')\}|. \quad (22)$$

With these notions, we show that, under certain choices of the variational parameters, the Opper–Archambeau approximation is exactly the correct posterior under a perturbed dataset, where \mathbf{x} is replaced with $\text{cl}(\mathbf{x})$ and all other data and model parameters are unchanged.

Proposition 15. Let $u_j = \frac{1}{N_{\text{cl}}(z_j)} \sum_{\text{cl}(x_i)=z_j} y_i$. The Bayesian models defined by $f \sim \text{GP}(0, k)$ and

$$y_i | f \sim \mathcal{N}(f(\text{cl}(x_i)), \sigma^2) \quad u_i | f \sim \mathcal{N}\left(f(z_i), \frac{\sigma^2}{N_{\text{cl}}(z_i)}\right). \quad (23)$$

admit respective posterior distributions $f | \mathbf{y}$ and $f | \mathbf{u}$ which are equal in distribution.

Proof. Appendix B. □

In cases where the spatial resolution of Definition 11 is small—which the cover tree guarantees—we can use Proposition 15 to accurately estimate the correct variational parameters without needing to perform optimization. This is done by choosing \mathbf{z} to be the clusters, choosing \mathbf{u} to be the cluster mean, and choosing the entries of the heteroskedastic noise matrix $\mathbf{\Lambda}$ to be equal to the original noise variance divided by the cluster size. We call this the *clustered data inducing point approximation*, and illustrate it in Figure 3.

The explicit description of the approximation made by use of inducing points in Proposition 15 enables us to better understand how to select them algorithmically. To develop this line of thought in more detail, we proceed to study how to choose inducing points to optimize the tradeoff between approximation error and computational costs using conjugate gradients.

In many applications, employing this approximation involves computing its marginal likelihood objective, or related quantities such as prior density evaluations or Kullback–Leibler divergences. One of the noteworthy properties of this approximation is that the Kullback–Leibler divergence between it and the true prior simplifies into

$$D_{\text{KL}}(q \parallel p) = \frac{1}{2} \ln \frac{|\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda}|}{|\mathbf{\Lambda}|} - \frac{1}{2} \text{tr}((\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1} \mathbf{K}_{\mathbf{z}\mathbf{z}}) + \frac{1}{2} \mathbf{v}^T \mathbf{K}_{\mathbf{z}\mathbf{z}} \mathbf{v} \quad (24)$$

where $\mathbf{v} = (\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1} \mathbf{u}$, and the minus sign in the trace term appears due to cancellation with the usual dimension term. The linear systems appearing in this term are more numerically stable than those in most other variational approximations, owing to the presence of $\mathbf{\Lambda}$. This applies to all computations needed to train the Gaussian process, not just the ones for obtaining predictive means and variances.

4. Experiments

We investigate the performance of the clustered-data inducing point approximation on a series of tests designed to illustrate both the behavior of its components and the overall picture. Each experiment focuses on a different aspect of the method, including approximation error, scalability, and behavior under varying quantities of data and inducing points. Full details regarding the experimental setup are given in Appendix D.

4.1 Spatial Resolution and Empirical Approximation Error

To better understand how different spatial resolutions resolve the tradeoff between approximation error and performance, we performed a number of numerical experiments designed to highlight this tradeoff. To quantify approximation error, we chose the Wasserstein distance, because it is finite in the setting of interest, is readily estimatable numerically, and controls expectations of a large class of posterior functionals (Villani, 2008).

To assess these quantities, we generated synthetic data by sampling $N = 1000$ input points uniformly on a hypercube of dimension 1, 2, 4 and 8, and then sampling from the prior process, which uses a squared exponential kernel. We then varied the spatial resolution, computed inducing points using the cover tree algorithm, and empirically estimated the Wasserstein distance between the approximate and true Gaussian processes. To better contextualize results, since the cover tree construction produces a variable number of inducing points, we also recorded the number of inducing points produced for spaces of different dimension.

Results can be seen in Figure 4. Immediately, we see that behavior of the cover tree depends on the dimension of the space the data lives on. In higher dimension, more inducing points are needed to construct a separated covering with the same spatial resolution. On the other hand, we see

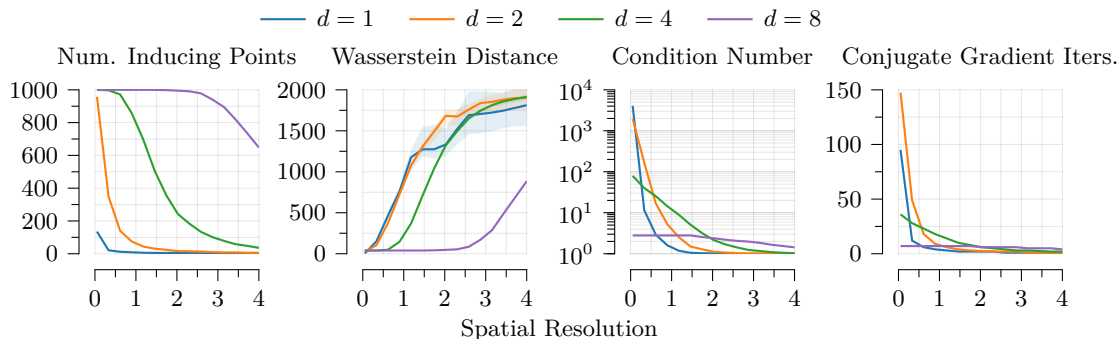


Figure 4: We illustrate how spatial resolution affects approximation accuracy and computational cost. We see that increasing the spatial resolution leads to using fewer inducing points, dropping rapidly in low dimension, and slower in higher dimensions. On the other hand, the error in approximation, as measured by Wasserstein distance between the approximate posterior and exact posterior, increases as the spatial resolution increases. Computational costs follow the opposite relationship: finer spatial resolutions are less stable and more expensive to compute, as evidenced by condition numbers and the number of iterations needed for conjugate gradients to converge.

spatial resolution directly controls approximation accuracy in Wasserstein distance, with finer spatial resolutions more accurate than coarser ones, and the effect being similar in different dimensions.

Next, to assess the consequence of different spatial resolutions on stability and corresponding difficulty of solving the resulting linear systems, we assess empirical performance of conjugate gradients in the settings of interest. We adopt the same data-generation process used in the preceding variant, but this time focus on the linear systems. We again vary spatial resolution, but now compute the numerical condition numbers of the resulting kernel matrices, and run the conjugate gradient algorithm to convergence, obtaining the number of iterations needed for convergence.

In Figure 4, we see that finer spatial resolutions generally result in more expensive-to-solve and less-stable linear systems. The effect varies according to dimension, and is least pronounced in $d=8$: in high dimensions, when sampling uniformly on a hypercube, the inputs tend to end up far away from each other, leading to less extreme eigenvalues. In lower dimensions, where the posterior is more concentrated, the effect of spatial resolution on condition numbers is more pronounced. In all cases, higher condition numbers lead to more iterations conjugate gradient iterations needed. In total, we see that spatial resolution is a directly interpretable quantity which facilitates the tradeoff between approximation accuracy with stability and computational cost.

4.2 Comparison to Alternative Inducing Point Selection Schemes

Next, to understand how different inducing point selection methods perform, we performed an experiment to assess the tradeoffs in terms of performance, stability, and computational cost. We compare five methods for selecting inducing points: the cover tree approach described in Algorithm 12, the online inducing point selection algorithm proposed by Galy-Fajou and Opper (2021), the partial pivoted Cholesky approach studied by Foster et al. (2009), along with baselines consisting of ordinary K -means, K -means++ (Arthur and Vassilvskii, 2007), optimizing the inducing points directly by minimizing the respective Kullback–Leibler divergence, and uniform random sampling from the training data.

To assess performance and stability, we evaluated the methods in terms of their produced kernel matrix condition numbers, and three sparse Gaussian process regression performance metrics, namely root mean squared error, negative log predictive density, and the evidence lower bound. Together, these give a view of both training and test-set performance, as well as variational approximation error.

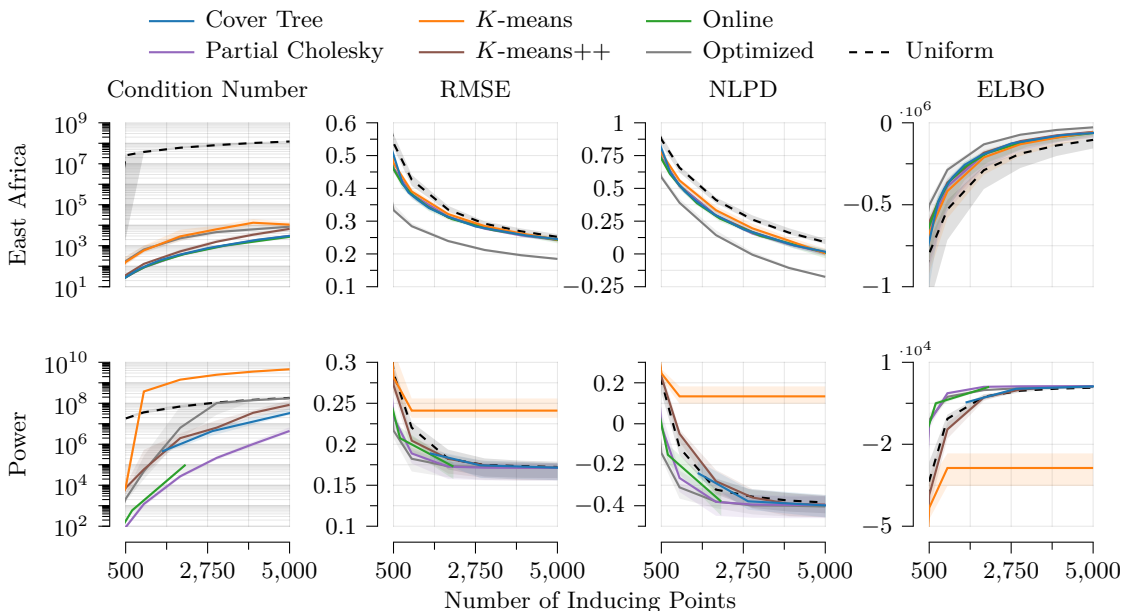


Figure 5: We compare how different inducing point selection algorithms vary according to performance and stability, using the geospatial *East Africa* dataset, and higher-dimensional *Power* dataset. The number of inducing points produced by algorithms whose output is variable is shown rounded to the nearest increment. Note that the computational complexity of these methods differs: cover tree is $\mathcal{O}(N \log N)$, online and the K -means variants are $\mathcal{O}(NM)$, partial pivoted Cholesky is $\mathcal{O}(NM^2)$, and optimizing the inducing points is $\mathcal{O}(NM^2)$ per optimization iteration. We see that using more inducing points results in higher performance, but reduced stability, with the precise effect varying according to dataset, dimensionality, and inducing point selection method.

We chose the sparse approximation of Titsias (2009) owing to its widespread use, working with the *GPFlow*, whose linear algebra is implemented in a stable manner using the V -method suggested by Foster et al. (2009). We use the squared exponential kernel and a jitter value of 10^{-6} . We evaluated the different methods on two datasets: the two-dimensional *East Africa* geospatial dataset studied by Wan et al. (2002), Weiss et al. (2014), and Ton et al. (2018), and the four-dimensional *Combined Cycle Power Plant* dataset from the UCI Machine Learning Repository (Tüfekci, 2014; Dua and Graff, 2017). These give a view of performance in both geospatial and non-geospatial settings, while ensuring all methods run successfully and can therefore be meaningfully compared.

Results can be seen in Figure 5. Immediately, we see that different methods produce different tradeoffs between performance and stability. These tradeoffs need not be achieved in an optimal manner: one might intuitively expect the ordinary K -means baseline to produce inducing points that are well-separated due to its clustering nature, and therefore result in stable linear systems. This doesn’t happen: instead, K -means tends to simultaneously achieve the worst performance and worst numerical stability among all methods except possibly uniform subsampling. In comparison, K -means++, which uses a sophisticated initialization scheme that promotes separation, produces much better performance and stability than K -means, though still not as well as methods with specified separation guarantees. This shows that one must explicitly encourage separation to obtain good performance-stability tradeoffs via such approaches.

The strongest predictive performance is achieved by optimizing the inducing points directly by minimizing the respective Kullback–Leibler divergence using gradient descent. The resulting inducing points, however, can be as numerically unstable as those produced by K -means or uniform

subsampling. Since optimizing inducing point locations requires one to specify an initialization, this suggests it may be advantageous to choose this initialization carefully, using one of the stable inducing point selection algorithms, to minimize potential numerical failures during the optimization process.

On the geospatial example, again somewhat-surprisingly, all methods with stability guarantees produced near-identical condition numbers and performance for a given number of inducing points. The main difference with these methods is therefore their computational cost, which in low dimension are $\mathcal{O}(N \log N)$ for the cover tree, $\mathcal{O}(NM)$ for the online approach and K -means++, and $\mathcal{O}(NM^2)$ for partial pivoted Cholesky. The online method’s complexity can, for some kernels, be improved to $\mathcal{O}(N \log N)$ by for instance precomputing a tree-based data structure that enables fast nearest-neighbor lookups, in which case the resulting variant starts to resemble the final loop of the cover tree construction algorithm. Overall, we conclude, in this setting, that significant numerical stability improvements can be obtained without reducing performance—and, using the cover tree, without paying excessive computational costs.

On the higher-dimensional non-geospatial example, no method of inducing point selection achieves a clear improvement on top of uniform subsampling of the data in terms of performance. On the other hand, the cover tree, online method, K -means++, and partial pivoted Cholesky all achieve better numerical stability compared to K -means, uniform sampling, and optimization, with partial pivoted Cholesky being the most stable. Due to the curse of dimensionality, the $\mathcal{O}(NM^2)$ computational costs of partial pivoted Cholesky are similar or favorable to the other approaches, making this method competitive. In this setting, we therefore conclude that significant numerical stability improvements can be obtained—however, unlike in the lower-dimensional setting, the computational costs needed to do so can become expensive.

In summary, we conclude from the comparisons that different approaches’ strengths are most effective in different settings. In low-dimensional geospatial problems, among all non-iterative methods, the cover tree algorithm provides good performance with guaranteed stability, while running in near-linear time. The online approach performs comparably, and can be more appropriate for settings where data arrives sequentially. K -means++, which promotes but does not enforce separation, performs slightly worse. In high-dimensional problems, the partial pivoted Cholesky algorithm results in strong performance, provided the desired number of inducing points is small enough that running it is practical. One can obtain the best performance by optimizing the inducing points directly, at cost of having to run an iterative procedure, and producing worse stability. In both classes of problems, our experiments suggest that methods with stability guarantees can offer a better performance-stability tradeoff compared to naïve methods.

4.3 Reliability in Floating Point Precision: Geospatial Illustrative Example

One of the appealing properties of using numerically stable approaches in variational inference is that, in principle, they can enable one to use a larger set of inducing points than customarily used. This is for two reasons: (i) the quadratic cost of conjugate gradients and (ii) the ability to save memory by effectively running in floating point precision. To illustrate that using more inducing points can improve performance on applied examples, we train the variational approximation of Titsias (2009) and with the clustered-data approximation of Section 3.3 on the East Africa land surface temperature dataset (Wan et al., 2002; Weiss et al., 2014; Ton et al., 2018). In both cases, we use floating-point precision, and the same inducing points obtained using a cover tree.

One of the defining characteristics of this dataset is that it has a relatively small number of data points per unit length scale ball, due to the fine-scale variation in land surface temperature. This renders approximations based on a sparse set of inducing points less accurate, unless the number of inducing points approaches the order of the size of the dataset. By varying the spatial resolution from relatively long ($\varepsilon = 0.09$) to very short ($\varepsilon = 0.03$) we see both qualitative and quantitative improvement, shown in Figure 6.

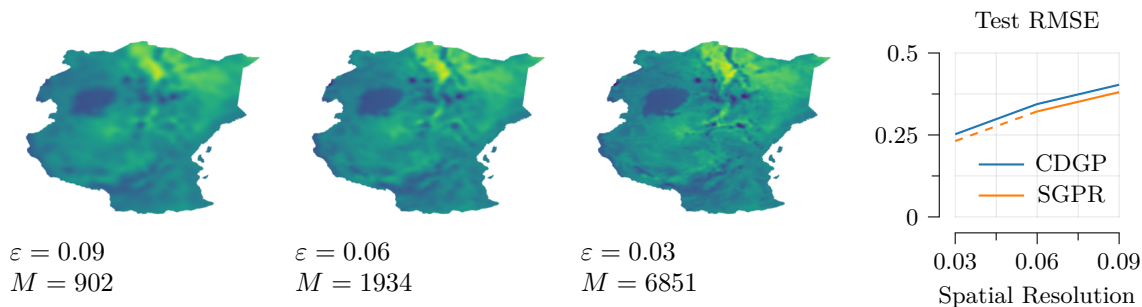


Figure 6: Here we illustrate the effect of spatial resolution ε and resulting number of inducing points M on the stability and performance of the Gaussian process model for the clustered-data approximation (blue), and sparse variational Gaussian process baseline (orange). We see that coarser spatial resolutions result in blurrier predictions and, in turn, higher RMSE. While the sparse variational Gaussian process baseline achieves better performance on coarser resolutions, its Cholesky factorization eventually fails, requiring us to increase jitter—this is displayed by the dashed line. The clustered-data approximation with conjugate gradients, on the other hand, runs reliably in all cases, achieving slightly worse but overall comparable performance.

When using larger sets of inducing points and running in floating point precision, numerical stability becomes increasingly crucial for running successfully. For the variational approximation of Titsias (2009), this manifests itself as Cholesky factorization failure. A common heuristic for alleviating this in the Gaussian process literature is to add *jitter* to the kernel matrix: this is done by replacing \mathbf{K}_{zz} with $\mathbf{K}_{zz} + \varepsilon\mathbf{I}$ for a small value $\varepsilon > 0$. This increases all of the eigenvalues of the matrix by ε , and ensures that the smallest eigenvalue cannot be arbitrarily close to 0. We found that using a larger jitter value compared to the *GPflow* recommended default was necessary when working with the finest spatial resolution of $\varepsilon = 0.03$, which produced $M = 6851$ inducing points, indicating that this separation distance alone was too small to ensure stability in this example.

In comparison, the clustered-data approximation does not require jitter, since it only involves inversion of the matrices $\mathbf{K}_{zz} + \mathbf{\Lambda}$ which include the diagonal matrix $\mathbf{\Lambda}$. This makes it a more reliable approach, but results in slightly worse predictive performance. In settings where the Gaussian process is trained in an automated manner without human oversight, or where data is collected online and not available in advance, reliability can be a bigger concern than performance. In such cases, our results indicate that the clustered-data approximation can be an alternative worth considering.

In this example, numerical stability is first and foremost a consequence of the minimum kernel matrix eigenvalue $\lambda_{\min}(\mathbf{K}_{zz})$ being too close to zero. Both introducing jitter and using the clustered-data approximation can alleviate this, but cannot help with issues arising from the maximum eigenvalue being too large, which must eventually occur if many inducing points accumulate in a bounded region. Our experiment indicates that this does not occur on the scales considered, but might instead occur when working with substantially larger sets of inducing points, for instance in the presence of sparse or otherwise structured kernel matrices (Durrande et al., 2019), which can often be inverted at substantially larger scale.

4.4 Numerical Stability and Dataset Size

The preceding sections show that, in the examples considered, numerical stability tends to deteriorate as approximation accuracy increases, as measured by the number of inducing points. On the other hand, in many settings, the more data one has, the more inducing points one needs to accurately represent the true posterior. One can therefore also think of the number of inducing points as measuring the expressive capacity of the variational family. This naturally raises the question: what

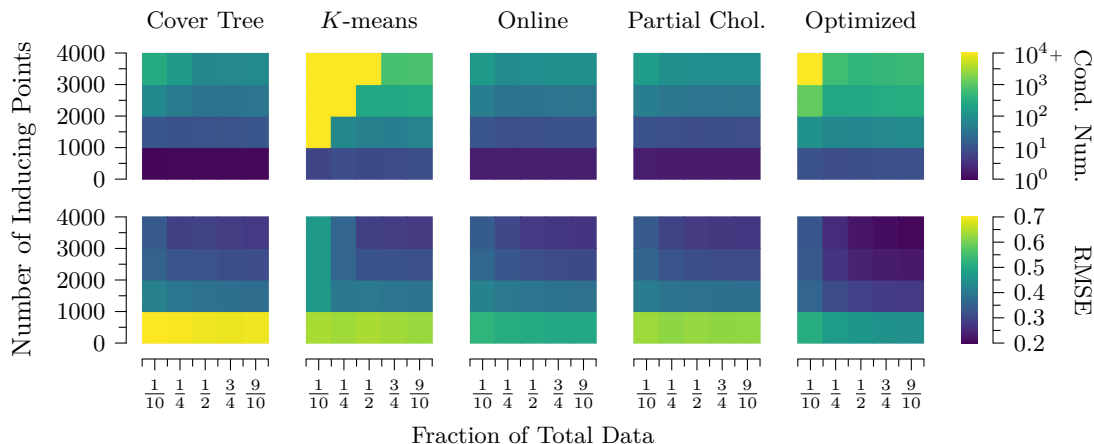


Figure 7: We illustrate how numerical stability and performance change when data size and approximation accuracy are varied simultaneously. Viewing a single plot top-to-bottom, we see that for a fixed data size, increasing the number of inducing points produces better performance and very slightly improves stability. Viewing the same plot left-to-right, we see that for a fixed level of variational approximation expressiveness, using more inducing inducing points produces better performance but decreases stability. Viewing the plot diagonally, we see that increasing data size and variational approximation expressiveness simultaneously results in better performance but worse stability, with different methods choosing tradeoffs in favor of either performance or stability, depending on the details of how they are designed.

happens to numerical stability when the number of data points grows, and expressive capacity grows with it?

To understand this, we decided to examine how stability and predictive performance change depending on data size and the number of inducing points simultaneously. For this, we randomly sampled subsets from the East Africa dataset used in the preceding illustrative example, and computed both the respective kernel matrix condition number, and predictive root-mean-squared error.

Results can be seen in Figure 7. For a fixed data size (Figure 7, top-to-bottom), we see that using more inducing points improves performance, but makes stability worse, with the tradeoff determined by the individual method: numerically stable methods favor stability, whereas optimizing the inducing points favors performance but quickly leads to unstable linear algebra. On the other hand, for a fixed number of inducing points (Figure 7, left-to-right), using more data always improves performance, and can actually slightly improve stability. In light of Corollary 9, this makes sense, because using more data provides a larger region within which the algorithms can place a fixed number of inducing points, allowing them to be spaced further apart from one another.

If we grow both the number of data points and the number of inducing points at the same time while keeping their ratio approximately fixed (Figure 7, diagonal, bottom-left to top-right), we see that numerical stability deteriorates as data size and the number of inducing points increase simultaneously. The severity of the tradeoff depends on the algorithm used: stable methods provide a tradeoff which favors numerical stability, while optimized inducing points favor performance and become unstable faster. *K-means* achieves the worst of both worlds, having a level of performance comparable to stable methods, and a level of numerical stability comparable to optimizing the inducing points. This highlights the value of methods designed to produce a favorable performance-stability tradeoff.

5. Conclusion

We study numerical stability of variational approximations used for Gaussian process models. We survey and synthesize a number of results which, in total, enable one to conclude that the *minimum separation distance* between inducing points controls stability properties of the numerical linear algebra needed to train variational approximations. We introduce two techniques for ensuring stability: inducing points constructed by a *cover tree*, and the *clustered-data approximation*. We provide examples illustrating the tradeoff between performance, stability, and dataset size. Compared to standard inducing point placement methods, we find that approaches that promote minimum separation produce a favorable performance-stability tradeoff.

Acknowledgments

DRB was supported by the Qualcomm Innovation Fellowship. SF was supported by the EPSRC (EP/V002910/2). We further acknowledge support from Huawei Research and Development, the Imperial College COVID-19 Research Fund, and UK Research and Innovation.

References

- V. Adam, P. Chang, M. E. E. Khan, and A. Solin. Dual Parameterization of Sparse Variational Gaussian Processes. In *Advances in Neural Information Processing Systems*, 2021. Cited on page 11.
- D. Arthur and S. Vassilvitskii. *K*-means++: The advantages of careful seeding. In *ACM Symposium on Discrete Algorithms*, 2007. Cited on page 13.
- S. Basak, S. Petit, J. Bect, and E. Vazquez. Numerical issues in maximum likelihood parameter estimation for Gaussian process interpolation. In *International Conference on Machine Learning, Optimization, and Data Science*, pages 116–131. Springer, 2021. Cited on page 2.
- A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *International Conference on Machine Learning*, pages 97–104, 2006. Cited on pages 2, 8.
- T. D. Bui, J. Yan, and R. E. Turner. A unifying framework for Gaussian process pseudo-point approximations using power expectation propagation. *The Journal of Machine Learning Research*, 2017. Cited on page 11.
- D. R. Burt, C. E. Rasmussen, and M. van der Wilk. Convergence of Sparse Variational Inference in Gaussian Processes Regression. *Journal of Machine Learning Research*, 2020. Cited on page 3.
- 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*, 2019. Cited on page 3.
- N. Cressie. *Statistics for Spatial Data*. Wiley, 1992. Cited on pages 2, 3.
- L. Csató and M. Opper. Sparse on-line Gaussian processes. *Neural Computation*, 2002. Cited on page 3.
- B. Diederichs and A. Iske. Improved estimates for condition numbers of radial basis function interpolation matrices. *Journal of Approximation Theory*, 2019. Cited on pages 2, 7.

- M. Dow. Explicit inverses of Toeplitz and associated matrices. *Australian and New Zealand Industrial and Applied Mathematics Journal*, 2002. Cited on page 5.
- D. Dua and C. Graff. UCI Machine Learning Repository, 2017. URL: [HTTP://ARCHIVE.ICS.UCI.EDU/ML](http://archive.ics.uci.edu/ml). Cited on page 14.
- N. Durrande, V. Adam, L. Bordeaux, S. Eleftheriadis, and J. Hensman. Banded Matrix Operators for Gaussian Markov Models in the Automatic Differentiation Era. In *Artificial Intelligence and Statistics*, 2019. Cited on page 16.
- W. E. *Principles of Multiscale Modeling*. Cambridge University Press, 2011. Cited on page 4.
- L. Foster, A. Waagen, N. Aijaz, M. Hurley, A. Luis, J. Rinsky, C. Satyavolu, M. J. Way, P. Gazis, and A. Srivastava. Stable and Efficient Gaussian Process Calculations. *Journal of Machine Learning Research*, 2009. Cited on pages 2, 13, 14.
- T. Galy-Fajou and M. Opper. Adaptive Inducing Points Selection For Gaussian Processes. In *Workshop on Continual Learning at the International Conference on Machine Learning*, 2021. Cited on page 13.
- J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson. Gpytorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. *Advances in Neural Information Processing Systems*, 2018. Cited on page 2.
- M. Gibbs and D. Mackay. Efficient Implementation of Gaussian Processes. Technical report, Cavendish Laboratory, University of Cambridge, 1997. Cited on pages 2, 31.
- G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996. Cited on pages 4, 23.
- D. H. Greene and D. E. Knuth. *Mathematics for the Analysis of Algorithms*. Springer, 1990. Cited on page 33.
- M. J. Heaton, A. Datta, A. O. Finley, R. Furrer, J. Guinness, R. Guhaniyogi, F. Gerber, R. B. Gramacy, D. Hammerling, M. Katzfuss, et al. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 24:398–425, 2019. Cited on page 3.
- J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, 2013. Cited on pages 2, 3, 11.
- M. F. Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics – Simulation and Computation*, 1989. Cited on pages 30, 31.
- M. Izbicki and C. Shelton. Faster cover trees. In *International Conference on Machine Learning*, 2015. Cited on pages 8, 9.
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv:1807.02582*, 2018. Cited on page 8.
- M. Katzfuss. A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*, 2017. Cited on page 3.

- M. Katzfuss and J. Guinness. A General Framework for Vecchia Approximations of Gaussian Processes. *Statistical Science*, 2021. Cited on page 3.
- A. M. Kibriya and E. Frank. An empirical comparison of exact nearest neighbour algorithms. In *European Conference on Principles of Data Mining and Knowledge Discovery*, 2007. Cited on page 8.
- A. Kielbasiński. A note on rounding-error analysis of Cholesky factorization. *Linear Algebra and its applications*, 88, 1987. Cited on page 4.
- V. Koltchinskii and E. Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, 2000. Cited on pages 6, 22.
- A. Krause and C. Guestrin. Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach. In *International Conference on Machine learning*, 2007. Cited on page 2.
- A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 2008. Cited on page 2.
- F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011. Cited on page 3.
- A. G. d. G. Matthews, J. Hensman, R. Turner, and Z. Ghahramani. On sparse variational methods and the Kullback–Leibler divergence between stochastic processes. In *Artificial Intelligence and Statistics*, 2016. Cited on page 3.
- A. G. d. G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. GPflow: A Gaussian Process Library using TensorFlow. *Journal of Machine Learning Research*, 2017. Cited on pages 2, 35.
- F. Narcowich, N. Sivakumar, and J. Ward. On Condition Numbers Associated with Radial-Function Interpolation. *Journal of Mathematical Analysis and Applications*, 1994. Cited on pages 2, 7.
- F. J. Narcowich and J. D. Ward. Norm estimates for the inverses of a general class of scattered-data radial-function interpolation matrices. *Journal of Approximation Theory*, 1992. Cited on pages 2, 7, 8, 25, 27.
- M. Opper and C. Archambeau. The variational Gaussian approximation revisited. *Neural Computation*, 2009. Cited on page 11.
- A. Panos, P. Dellaportas, and M. K. Titsias. Fully scalable Gaussian processes using subspace inducing inputs. *arXiv:1807.02537*, 2018. Cited on page 11.
- J. Quiñonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 2005. Cited on pages 2, 3.
- H. Sang and J. Z. Huang. A full scale approximation of covariance functions for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2012. Cited on page 3.
- R. Schaback. Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 1995. Cited on pages 2, 7, 8, 25, 27.

- M. W. Seeger, C. K. Williams, and N. D. Lawrence. Fast forward selection to speed up sparse Gaussian process regression. In *Artificial Intelligence and Statistics*, 2003. Cited on page 3.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, 2012. Cited on page 2.
- M. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, 2009. Cited on pages 2, 3, 11, 14–16.
- J.-F. Ton, S. Flaxman, D. Sejdinovic, and S. Bhatt. Spatial mapping with Gaussian processes and nonstationary Fourier features. *Spatial Statistics*, 28:59–78, 2018. Cited on pages 14, 15.
- W. F. Trench. Properties of some generalizations of Kac–Murdock–Szegő matrices. *Contemporary Mathematics*, 2001. Cited on page 5.
- P. Tüfekci. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. *International Journal of Electrical Power and Energy Systems*, 2014. Cited on page 14.
- M. van der Wilk, A. Artemev, and J. Hensman. Improved Inverse-Free Variational Bounds for Sparse Gaussian Processes. In *Symposium on Advances in Approximate Bayesian Inference*. 2022. Cited on page 11.
- R. van Handel. *Probability in high dimension*. Princeton University, 2014. Cited on page 25.
- A. V. Vecchia. Estimation and Model Identification for Continuous Spatial Processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 1988. Cited on page 3.
- R. Vershynin. *High-dimensional probability*. Cambridge University Press, 2018. Cited on pages 8, 25, 34.
- C. Villani. *Optimal Transport: old and new*. Springer, 2008. Cited on page 12.
- Z. Wan, Y. Zhang, Q. Zhang, and Z.-l. Li. Validation of the land-surface temperature products retrieved from Terra Moderate Resolution Imaging Spectroradiometer data. *Remote Sensing of Environment*, 2002. Cited on pages 14, 15.
- D. J. Weiss, P. M. Atkinson, S. Bhatt, B. Mappin, S. I. Hay, and P. W. Gething. An effective approach for gap-filling continental scale remotely sensed time-series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2014. Cited on pages 14, 15.
- H. Wendland. *Scattered Data Approximation*. Cambridge University Press, 2004. Cited on pages 7, 25–27.
- J. H. Wilkinson. A Priori Error Analysis of Algebraic Processes. In *Proceedings of the International Congress of Mathematicians*, pages 119–129, 1966. Cited on page 4.
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Efficiently Sampling Functions from Gaussian Process Posteriors. In *International Conference on Machine Learning*, 2020. Cited on page 3.
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise Conditioning of Gaussian Processes. *Journal of Machine Learning Research*, 2021. Cited on page 3.
- J. Xu and L. Zikatanov. Algebraic Multigrid Methods. *Acta Numerica*, 2017. Cited on page 4.

Appendix A. Eigenvalues and Condition Numbers

Here we analyze the behavior eigenvalues of kernel matrices, and how this affects numerical linear algebra.

A.1 Covariance Operator Limits

Proposition 3. *Let k be a continuous stationary kernel, and let the entries in \mathbf{x} be independently sampled from a uniform distribution on some compact subset $X \subseteq \mathbb{R}^d$. Define the covariance operator*

$$\mathcal{K} : L^2(X) \rightarrow L^2(X) \qquad \mathcal{K} : \phi \mapsto \int_X \phi(x)k(x, \cdot) dx \qquad (11)$$

acting on the Hilbert space of (equivalence classes of) square-integrable functions. Then the eigenvalues of \mathcal{K} are countable, non-negative, admit a maximum, and can be ordered to form a non-increasing sequence whose limit is zero. For every $n > 0$, as $N \rightarrow \infty$, the n^{th} largest eigenvalue of $\frac{1}{N}\mathbf{K}_{\mathbf{x}\mathbf{x}}$ converges almost surely to the corresponding eigenvalue of \mathcal{K} . As a consequence, $\text{cond}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \rightarrow \infty$ almost surely.

Proof. Define the (random) matrix $\tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}}$ to be zero on the diagonal and equal to $\tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}}$ on the off-diagonal. Let Π denotes the set of all permutations of the natural numbers. For an operator \mathcal{A} with discrete spectrum, we view the spectrum as an element of ℓ^p/Π , denoted by $\lambda(\mathcal{A})$. By convention, we extend the spectrum of a symmetric, positive definite matrix \mathbf{M} to an element of ℓ^p/Π by appending zeros, denoted by $\lambda(\mathbf{M})$. For $p \in (1, \infty]$ define the the ℓ^p rearrangement distance by,

$$d_p(x, y) = \inf_{\pi \in \Pi} \left(\sum_{i=1}^{\infty} |x_i - y_{\pi(i)}|^p \right)^{1/p}, \qquad (25)$$

with the unusual convention that $p = \infty$ corresponds to the element-wise supremum. One can show that d_p defines a pseudo-metric on ℓ^p and a metric on ℓ^p/Π : the only property that is tricky to check is triangle inequality, which is shown in Koltchinskii and Giné (2000, p. 116), in the case $p = 2$, which immediately generalizes to other p . Since k is stationary, $\tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}} = \mathbf{K}_{\mathbf{x}\mathbf{x}} - c\mathbf{I}$ for some $c > 0$. Therefore,

$$d_{\infty} \left(\lambda \left(\frac{1}{N} \tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}} \right), \lambda \left(\frac{1}{N} \mathbf{K}_{\mathbf{x}\mathbf{x}} \right) \right) = \frac{c}{N} \rightarrow 0. \qquad (26)$$

Since k is bounded, the conditions of Koltchinskii and Giné (2000, Corollary 3.3) are satisfied, so

$$d_2 \left(\lambda \left(\frac{1}{N} \tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}} \right), \lambda(\mathcal{K}) \right) \rightarrow 0. \qquad (27)$$

and therefore by the standard (ℓ^2, ℓ^{∞}) inequality

$$d_{\infty} \left(\lambda \left(\frac{1}{N} \tilde{\mathbf{K}}_{\mathbf{x}\mathbf{x}} \right), \lambda(\mathcal{K}) \right) \rightarrow 0 \qquad (28)$$

which implies that by the triangle inequality

$$d_{\infty} \left(\lambda \left(\frac{1}{N} \mathbf{K}_{\mathbf{x}\mathbf{x}} \right), \lambda(\mathcal{K}) \right) \rightarrow 0. \qquad (29)$$

We have

$$\left| \lambda_1 \left(\frac{1}{N} \mathbf{K}_{\mathbf{x}\mathbf{x}} \right) - \lambda_1(\mathcal{K}) \right| \leq d_{\infty} \left(\lambda \left(\frac{1}{N} \mathbf{K}_{\mathbf{x}\mathbf{x}} \right), \lambda(\mathcal{K}) \right) \rightarrow 0, \qquad (30)$$

so by the reverse triangle inequality

$$\lambda_1(\mathbf{K}_{\mathbf{x}\mathbf{x}}) = N\lambda_1(\mathcal{K}) + o(N). \qquad (31)$$

As \mathcal{K} is not the zero operator and is positive definite, this is $\Omega(N)$. On the other hand, there exists an N_0 such that for all $n > N_0$, we have $\lambda_n(\mathcal{K}) \leq \varepsilon/2$. Also, there exists an N'_0 such that for all $n \geq N'_0$

$$d_\infty(\lambda(\frac{1}{N}\mathbf{K}_{\mathbf{x}\mathbf{x}}), \lambda(\mathcal{K})) \leq \varepsilon/2. \quad (32)$$

Choose $M_0 = \max(N_0, N'_0)$. Then for all $n \geq M_0$,

$$\frac{1}{N}\lambda_N(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \leq d_\infty(\lambda(\frac{1}{N}\mathbf{K}_{\mathbf{x}\mathbf{x}}), \lambda(\mathcal{K})) + \lambda_n(\mathcal{K}) \leq \varepsilon \quad (33)$$

and so $\lambda_N(\mathbf{K}_{\mathbf{x}\mathbf{x}}) = o(N)$. In fact, we can say much more: as the data is compactly supported, for any $\varepsilon > 0$ and for all N sufficient large there exists a $1 \leq i, j \leq N$ such that $k(x_i, x_j) \geq (1 - \varepsilon)k(x_i, x_i)$. By Cauchy's Interlacing Theorem,

$$\lambda_N(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \leq \lambda_2 \begin{pmatrix} k(x_i, x_i) & k(x_i, x_j) \\ k(x_j, x_i) & k(x_j, x_j) \end{pmatrix} \leq k(x_i, x_i)\varepsilon. \quad (34)$$

As ε is arbitrary, we conclude $\lambda_N(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \rightarrow 0$. \square

A.2 Convergence of Conjugate Gradients

Result 16. *Let \mathbf{v}^* satisfy $\mathbf{A}\mathbf{v}^* = \mathbf{b}$. For any $t \in \mathbb{N}$ and $\mathbf{v}_0 \in \mathbb{R}^n$ and \mathbf{v}_t denote the solution found by running conjugate gradients on this system of equations with initial vector \mathbf{v}_0 , then*

$$\|\mathbf{v}^* - \mathbf{v}_t\|_{\mathbf{A}} \leq 2 \left(1 - \frac{2}{\sqrt{\text{cond}(\mathbf{A})} + 1} \right)^t \|\mathbf{v}^* - \mathbf{v}_0\|_{\mathbf{A}}. \quad (35)$$

Proof. Golub and Van Loan (1996, Theorem 10.2.6). \square

Corollary 17. *A sufficient condition for $\|\mathbf{v}^* - \mathbf{v}_t\|_{\mathbf{A}} \leq \varepsilon$ is*

$$t \geq \left(\log \left(1 + \frac{2}{\sqrt{\text{cond}(\mathbf{A})} + 1} \right) \right)^{-1} \log \frac{2\|\mathbf{v}^* - \mathbf{v}_0\|_{\mathbf{A}}}{\varepsilon} = \left(\frac{\sqrt{\text{cond}(\mathbf{A})}}{2} + \mathcal{O}(1) \right) \log \frac{2\|\mathbf{v}^* - \mathbf{v}_0\|_{\mathbf{A}}}{\varepsilon} \quad (36)$$

where the implicit constant is between $\frac{1}{2}$ and $\frac{3}{2}$.

Proof. The first inequality is immediate from inverting the bound in Result 16. The asymptotic formula can be derived directly from applying the inequalities $1 - \frac{1}{x} \leq \log x \leq x - 1$ to yield

$$\frac{\sqrt{\text{cond}(\mathbf{A})} + 1}{2} \leq \left(\log \left(1 + \frac{2}{\sqrt{\text{cond}(\mathbf{A})} + 1} \right) \right)^{-1} \leq \frac{\sqrt{\text{cond}(\mathbf{A})} + 3}{2}. \quad (37)$$

\square

A.3 Upper Eigenvalue Estimates

Here we prove the eigenvalue upper bound under separation and spatial decay via a packing argument. We begin with a basic eigenvalue estimate.

Result 18 (Gershgorin's Circle Theorem). *Let \mathbf{K} be a positive semi-definite symmetric matrix with eigenvalues $\lambda(\mathbf{K}) = \{\lambda_i\}$. Define the GERSHGORIN RADIUS $R_i = \sum_{j \neq i} |k_{ij}|$ to be the sum of the off-diagonal entries of \mathbf{K} . Then $\lambda(\mathbf{K}) \in \bigcup_i [k_{ii} - R_i, k_{ii} + R_i]$.*

Proof. Golub and Van Loan (1996, Theorem 8.1.3). \square

We can use this result to get upper eigenvalue estimates as

$$\lambda_{\max} \leq \max_i k_{ii} + R_i \quad (38)$$

given a set of points that are separated and a kernel satisfying spatial decay. We now introduce these notions.

Definition 19. We say that a set of points $\mathbf{x} \subset \mathbb{R}^d$ are δ -SEPARATED if, for every $x_i \neq x_j$ we have $\|x_i - x_j\| \geq \delta > 0$.

Definition 20. We define the ball of radius δ -centered at $x \in \mathbb{R}^d$ to be the open set,

$$B_{x,\delta} = \{x' \in X : \|x - x'\| < \delta\}. \quad (39)$$

We write $B_\delta = B_{0,\delta}$.

Assumption 5 (Spatial decay). Let k be a kernel on $X \subseteq \mathbb{R}^d$. We say that k has SPATIAL DECAY if there is a decreasing function $\psi : [0, \infty) \rightarrow [0, \infty)$ such that for all $x, x' \in X$ we have

$$|k(x, x')| \leq \psi(\|x - x'\|) \quad \psi(m) = \mathcal{O}\left(\frac{1}{m^d \log(m)^2}\right). \quad (13)$$

The argument for the maximum eigenvalue bound is essentially that if we have spatial decay and uniformly bounded variance, then the kernel matrix has bounded diagonal and bounded Gershgorin radius, both independent of N .

Definition 21 (Covering and packing numbers). Let $X \subset \mathbb{R}^d$ and let $\delta > 0$. Define the following.

1. A subset $C \subseteq C' \subseteq \mathbb{R}^d$ is an EXTERNAL COVER of X in C' if for every $x \in X$ there is a $c \in C$ such that $\|x - c\| \leq \delta$.
2. A subset $C \subseteq X$ is an INTERNAL COVER of X if for every $x \in X$ there is a $c \in C$ such that $\|x - c\| \leq \delta$.
3. A subset $P \subseteq X$ is an EXTERNAL PACKING of X if any distinct $x, x' \in P$ satisfy $\|x - x'\| > \delta$.
4. A subset $P \subseteq X$ is an INTERNAL PACKING of X if it is an external packing, and moreover satisfies $B_{x,\delta} \subseteq X$ for all $x \in P$.

Define the internal and external covering numbers $C_X^{\text{int},\delta}$ and $C_{X,C'}^{\text{ext},\delta}$ to be the smallest cardinality of such sets, and the internal and external packing numbers $P_X^{\text{int},\delta}$ and $P_X^{\text{ext},\delta}$ to be the largest cardinality of such sets.

Lemma 22. Let $A_{m,\delta} = B_{0,(m+1)\delta} \setminus B_{0,m\delta} = \{x \in \mathbb{R}^d : m\delta \leq \|x\| < (m+1)\delta\}$ be an annulus with external radius $(m+1)\delta$ and internal radius $m\delta$. Then we have

$$\text{vol}(B_\delta) = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2} + 1)} \delta^d \quad \text{vol}(A_{m,\delta}) = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2} + 1)} ((m+1)^d - m^d) \delta^d. \quad (40)$$

We also have for $\beta \geq 2$ as well as $\beta = \frac{1}{2}$ and $\beta = \frac{3}{2}$ that

$$(\beta + 2)^d - \beta^d \leq 5^d \beta^{d-1}. \quad (41)$$

Proof. The first part of the claim is standard. For (41), suppose first that $\beta \geq 2$, and expand the first term on the left hand side with the Binomial Theorem. The lead term cancels with the second term on the left hand side, and we are left with a total of $2^d - 1$ non-negative terms, among which, using the assumption that $\beta \geq 2$, the largest is $2\beta^{d-1}$. Combining these estimates gives $(\beta + 2)^d - \beta^d \leq (2^{d+1} - 2)\beta^{d-1}$, which using $2^{d+1} - 2 \leq 5^d$ gives the inequality. We now handle the remaining cases: for $\beta = \frac{1}{2}$ we have $(\frac{3}{2})^d - (\frac{1}{2})^d \leq 5^d (\frac{1}{2})^{d-1}$, and for $\beta = \frac{3}{2}$, we have $(\frac{7}{2})^d - (\frac{3}{2})^d \leq 5^d (\frac{3}{2})^{d-1}$, both by elementary rearrangement of terms. The claim follows. \square

Lemma 23. *For an integer $m \geq 1$, we have*

$$\left(\frac{1}{\delta}\right)^d \frac{\text{vol}(A_{m,\delta})}{\text{vol}(B_\delta)} \leq C_{A_{m,\delta}}^{\text{ext},\delta} \leq P_{A_{m,\delta}}^{\text{ext},\delta} \leq 5^d \left(\frac{2}{\delta}\right)^d \left(m - \frac{1}{2}\right)^{d-1}. \quad (42)$$

and $P_{A_{0,\delta}}^{\text{ext},\delta} = 1$.

Proof. The first part of the claim follows immediately from van Handel (2014, Lemma 5.13). To bound $P_{A_{m,\delta}}^{\text{ext},\delta}$ for $m \geq 1$, use Vershynin (2018, Proposition 4.2.12), along with the fact that the Minkowski sum of an annulus with a ball is another annulus, and the previous lemma, to bound

$$P_{A_{m,\delta}}^{\text{ext},\delta} \leq \frac{\text{vol}(B_{(m+\frac{3}{2})\delta} \setminus B_{(m-\frac{1}{2})\delta})}{\text{vol}(B_{\delta/2})} = \frac{(m+\frac{3}{2})^d - (m-\frac{1}{2})^d}{\left(\frac{\delta}{2}\right)^d} \leq 5^d \left(\frac{2}{\delta}\right)^d \left(m - \frac{1}{2}\right)^{d-1}. \quad (43)$$

For the case $m = 0$, $A_{0,\delta} = B_\delta$ is non-empty, so $P_{A_{0,\delta}}^{\text{ext},\delta} \geq 1$. On the other hand, for any $x, x' \in B_\delta$, $\|x - x'\| < \delta$, hence they cannot be part of the same δ -packing. \square

Proposition 24. *Let $X = \mathbb{R}^d$, and let k satisfy spatial decay. Then there is a constant $C_{\max}^{k,\delta}$ such that for any N and any δ -separated \mathbf{x} of size N we have*

$$\lambda_{\max}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \leq C_{\max}^{k,\delta} \quad (44)$$

Proof. Let ψ be the upper-bounding function. For an arbitrary x_i , we have

$$\lambda_{\max}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \leq k_{ii} + R_i \leq \sum_{j=1}^N \psi(\|x_i - x_j\|) \quad (45)$$

so it suffices to bound this sum. For that, we have

$$\sum_{j=1}^N \psi(\|x_i - x_j\|) = \sum_{m=0}^{\infty} \sum_{x_j \in A_{m,\delta}} \psi(\|x_i - x_j\|) \leq \sum_{m=0}^{\infty} P_{A_{m,\delta}}^{\text{ext},\delta} \psi(m\delta). \quad (46)$$

The first equality uses that $\{A_{m,\delta}\}_{m=0}^{\infty}$ forms a partition of \mathbb{R}^d . The second uses that for $x_j \in A_{m,\delta}$, we have $\psi(\|x_i - x_j\|) \leq \psi(m\delta)$, and the maximum number of x_j in each $A_{m,\delta}$ is bounded by the external packing number $P_{A_{m,\delta}}^{\text{ext},\delta}$ by the definition of a δ -separated set. By the preceding lemma, spatial decay, and the integral test for infinite series, we conclude

$$\sum_{m=0}^{\infty} P_{A_{m,\delta}}^{\text{ext},\delta} \psi(m\delta) \leq \psi(0) + 5^d \left(\frac{2}{\delta}\right)^d \sum_{m=1}^{\infty} \left(m - \frac{1}{2}\right)^{d-1} \psi(m\delta) = C_{\max}^{k,\delta} < \infty. \quad (47)$$

\square

A.4 Lower Eigenvalue Estimates

For lower bounds, we apply Wendland (2004, Theorem 12.3), which in the stationary case follows from our assumptions. This argument, which builds on the work of Narcowich and Ward (1992) and Schaback (1995), is somewhat involved. To aid understanding, rather than simply quote the result, below we present a sketch of its key steps.

The idea is to invoke the lower estimate given by Gershgorin's Circle Theorem—Result 18—which was previously used for the upper bounds. However, applying this estimate directly leads in general to the inequality $\lambda_{\min}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \geq \tilde{C}_{\min}^{k,\delta}$ where potentially $\tilde{C}_{\min}^{k,\delta} \leq 0$, which is vacuous. To overcome this, the idea is to carefully construct a different covariance kernel ϕ whose eigenvalues lower-bound those of k , and which admit better Gershgorin estimates. This is done via the following lemma.

Lemma 25. *Let k and ϕ be stationary kernels with respective spectral densities ρ and ς . Suppose that $\varsigma \leq \rho$. Then for all \mathbf{x} we have*

$$\lambda_{\min}(\Phi_{\mathbf{x}\mathbf{x}}) \leq \lambda_{\min}(\mathbf{K}_{\mathbf{x}\mathbf{x}}). \quad (48)$$

Proof. For any $\mathbf{v} \in \mathbb{R}^N$ and by Bochner's Theorem—see Wendland (2004, Theorem 6.6)—we have

$$\mathbf{v}^T \mathbf{K}_{\mathbf{x}\mathbf{x}} \mathbf{v} = \sum_{i=1}^N \sum_{j=1}^N v_i k(0, x_i - x_j) v_j \quad (49)$$

$$= \sum_{i=1}^N \sum_{j=1}^N v_i v_j \int_{\mathbb{R}^d} e^{2\pi i \langle \omega, x_i - x_j \rangle} d\rho(\omega) \quad (50)$$

$$= \int_{\mathbb{R}^d} \sum_{i=1}^N v_i e^{2\pi i \langle \omega, x_i \rangle} \sum_{j=1}^N v_j e^{2\pi i \langle \omega, x_j \rangle} d\rho(\omega) \quad (51)$$

$$= \int_{\mathbb{R}^d} \left| \sum_{j=1}^N v_j e^{2\pi i \omega x_j} \right|^2 d\rho(\omega) \quad (52)$$

$$\geq \int_{\mathbb{R}^d} \left| \sum_{j=1}^N v_j e^{2\pi i \omega x_j} \right|^2 d\varsigma(\omega) \quad (53)$$

$$= \mathbf{v}^T \Phi_{\mathbf{x}\mathbf{x}} \mathbf{v} \quad (54)$$

where, in the third inequality, we use that since v_j is real, it is equal to its complex conjugate. Taking an infimum over \mathbf{v} on the unit sphere on the left hand side gives the result. \square

With this established, we can obtain lower eigenvalue estimates for any kernel ϕ whose Gershgorin radius is small enough, by invoking separation and using a packing argument similar to the one used to prove Proposition 24. This is done as follows.

Lemma 26. *Let $X = \mathbb{R}^d$, and let ϕ decay sufficiently fast. Then there is a constant $C_{\min}^{\phi, \delta} > 0$ such that for any N any δ -separated \mathbf{x} of size N we have*

$$\lambda_{\min}(\Phi_{\mathbf{x}\mathbf{x}}) \geq C_{\min}^{\phi, \delta}. \quad (55)$$

Proof. For arbitrary x_i , letting R_i be the Gershgorin radius, we have

$$\lambda_{\min}(\Phi_{\mathbf{x}\mathbf{x}}) \geq \Phi_{ii} - R_i \quad (56)$$

so it suffices to upper-bound R_i sufficiently tightly. Applying the packing argument of Proposition 6, we get

$$\sum_{j \neq i} \phi(x_i - x_j) \leq d \left(\frac{3}{\delta} \right)^d \sum_{m=1}^{\infty} m^{d-1} \phi(m\delta) \quad (57)$$

where we have substituted the lower-bounding kernel ϕ in to the argument, modified slightly to exclude the diagonal part, which corresponds to $m = 0$. The claim follows, provided

$$d \left(\frac{3}{\delta} \right)^d \sum_{m=1}^{\infty} m^{d-1} \phi(m\delta) < \phi(0) \quad (58)$$

which holds provided that ϕ decays sufficiently fast. \square

From here, all that remains is to show that for any stationary kernel k , it is always possible to choose a ϕ that decays sufficiently fast. Wendland (2004) shows that for every kernel k there is a constant $0 < M_k < \infty$ for which taking the spectral density ς of ϕ to be the convolution of an indicator function of the ball B_{M_k} with itself, namely

$$\varsigma \propto \mathbb{1}_{B_{M_k}} * \mathbb{1}_{B_{M_k}} \quad (59)$$

with proportionality scaling chosen to ensure the spectral lower-bound needed for Lemma 25 holds, suffices to yield a non-vacuous bound. The idea behind this choice, which dates back at least to ideas of Narcowich and Ward (1992) and Schaback (1995) in the polynomial interpolation literature, is to leverage the uncertainty principle property of the Fourier transform. Specifically, any kernel concentrated enough to decay sufficiently fast in the spatial domain must possess a sufficiently diffuse spectral measure. The simplest choice is to take $\varsigma \propto \mathbb{1}_{B_{M_k}}$, but this fails due to issues with smoothness, which are fixed by applying a convolution and rescaling appropriately, yielding the presented choice.

At this point, the main technical work that remains is to calculate the form of the resulting kernel ϕ , and deduce how to choose the constant M_k and proportionality scaling to ensure it both lower-bounds k and gives a non-vacuous estimate. For that, we refer to Wendland (2004, Ch. 12.2).

Result 27. *Let $X = \mathbb{R}^d$, and let k be stationary. Then there is a constant $C_{\min}^{k,\delta}$ such that for any N and any δ -separated \mathbf{x} of size N we have*

$$\lambda_{\min}(\mathbf{K}_{\mathbf{x}\mathbf{x}}) \geq C_{\min}^{k,\delta}. \quad (60)$$

Proof. Wendland (2004, Theorem 12.3). □

A.5 Settings where Separation is a Necessary Condition

Let k be a kernel. We say that k is a *Lipschitz continuous kernel* if k is Lipschitz in both arguments separately. We extend the notion of minimum separation from finite vectors to infinite sequences in the obvious pairwise manner.

Proposition 10. *Let $z_m \in X$, $m = 1, \dots, \infty$, be a sequence, and let k be a Lipschitz continuous kernel on $X \times X$. Define $\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(M)} \in \mathbb{R}^{M \times M}$ to be the kernel matrix formed by evaluating k at z_1, \dots, z_M . If there exists a $C > 0$ such that for all M , $\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(M)}) \leq C$, then z_m satisfy minimum separation.*

Proof. We argue by contradiction: suppose minimum separation does not hold. We will show that there does not exist a C such that, for all M , $\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(M)}) \leq C$. By Cauchy's Interlacing Theorem, any submatrix of $\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(M)}$ has condition number no larger than $\text{cond}(\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(M)})$. Hence, it suffices to find a sequence of submatrices \mathbf{K}_m , for $m = 1, \dots, \infty$, such that the rows and columns of \mathbf{K}_m are contained in the rows and columns of $\mathbf{K}_{\mathbf{z}\mathbf{z}}^{(m)}$, and show that this sequence has unbounded condition number. We therefore consider such sequences, consisting of 2×2 matrices, in the remainder of the proof.

We now split into two cases. The first case handles when the variance of the kernel can differ by an arbitrary factor at different points in input space. This is handled as a separate case, as the technique in the second, arguably more intuitive case, will break down if the variance can be arbitrarily close to zero.

Case I: for any $C > 0$, there exists an i, j with $\frac{k(z_i, z_i)}{k(z_j, z_j)} > C$. Let \mathbf{K}_m be the matrix formed by

$$z_i^{(m)} = \arg \max_{i' \leq m} k(z_{i'}, z_{i'}) \quad z_j^{(m)} = \arg \min_{j' \leq m} k(z_{j'}, z_{j'}) \quad (61)$$

where we take the smallest such i, j if the maximum and minimum are not unique. Then

$$\text{cond}(\mathbf{K}_m) \geq \text{cond} \begin{pmatrix} k(z_i^{(m)}, z_i^{(m)}) & 0 \\ 0 & k(z_j^{(m)}, z_j^{(m)}) \end{pmatrix} = \frac{k(z_i^{(m)}, z_i^{(m)})}{k(z_j^{(m)}, z_j^{(m)})}. \quad (62)$$

The inequality can be verified by inspection using the formula for eigenvalues of a 2×2 matrix. By the assumption in this case, this ratio tends to infinity with m , contradicting that the condition number is bounded.

Case II: there exists a $C > 0$ such that for all pairs i, j we have $\frac{k(z_i, z_i)}{k(z_j, z_j)} \leq C$. By the assumption in this case, for all i , we have

$$0 < \frac{1}{C}k(z_1, z_1) \leq k(z_i, z_i) \leq Ck(z_1, z_1). \quad (63)$$

Define $a = \frac{1}{C}k(z_1, z_1)$. Take $z_i^{(m)}, z_j^{(m)}$ to be any pair such that $\|z_i^{(m)} - z_j^{(m)}\| \leq \|z_{i'} - z_{j'}\|$ for any pair $i', j' \leq m$, and \mathbf{K}_m to be the submatrix associated to this pair of points. By Lipschitzness of the kernel, there exists an L such that

$$|k(z_i^{(m)}, z_i^{(m)}) - k(z_i^{(m)}, z_j^{(m)})| \leq L \text{sep}_m \quad |k(z_i^{(m)}, z_i^{(m)}) - k(z_k^{(m)}, z_j^{(m)})| \leq 2L \text{sep}_m. \quad (64)$$

where sep_m denotes the separation distance of (z_1, \dots, z_m) . Combining these estimates gives the Frobenius norm bound

$$\left\| \frac{1}{k(z_i, z_i)} \mathbf{K}_m - \mathbf{1} \right\|_F \leq \frac{L \text{sep}_m \sqrt{6}}{k(z_i, z_i)} \leq \frac{L \text{sep}_m \sqrt{6}}{a}, \quad (65)$$

where $\mathbf{1}$ is the 2×2 matrix of all ones. As the condition number of $\frac{1}{k(z_i, z_i)} \mathbf{K}_m$ and \mathbf{K}_m are the same, we study the condition number of $\frac{1}{k(z_i, z_i)} \mathbf{K}_m$. The operator norm is less than the Frobenius norm, so by reverse triangle inequality applied to the operator norm

$$\lambda_1 \left(\frac{1}{k(z_i, z_i)} \mathbf{K}_m \right) \geq 2 - \left\| \frac{1}{k(z_i, z_i)} \mathbf{K}_m - \mathbf{1} \right\|_2 \geq 2 - \frac{L \text{sep}_m \sqrt{6}}{a} \quad (66)$$

and

$$\lambda_2 \left(\frac{1}{k(z_i, z_i)} \mathbf{K}_m \right) \leq \frac{L \text{sep}_m \sqrt{6}}{a}. \quad (67)$$

By our assumption that sep_m is not bounded below, this ratio of these tends to infinity with m . \square

A.6 Effect of Dimension on Condition Number

We now consider the simple case of independent inputs in \mathbb{R}^d with independent and identically distributed dimensions. The key observation is that for large d , by concentration of measure, we should most pairwise distances to concentrate around the average distance. This means that there will not be many tightly clustered sets of data points. Instead, the data will be automatically well-separated with high probability in this specific setting.

Let ρ be the distribution of each input. We work with uniformly generated inputs for Figure 4, specifically $\rho = \text{U}(-2\pi, 2\pi)$. For any $i \neq j$ we have

$$\|x_i - x_j\|_2^2 = \sum_{\ell=1}^d \Delta_\ell^2, \quad (68)$$

with $\Delta_\ell = (\xi_{i,\ell} - \xi_{j,\ell})^2$ where $\xi_{i,\ell} \sim \rho$ are all independent, hence Δ_ℓ are also all independent.

If we consider the isotropic squared exponential kernel with length scale \sqrt{d} , which is used for generating Figure 4, the off-diagonal entries of the kernel matrix are independently and identically distributed random variables of the form

$$k_{ij} = \exp \left(-\frac{1}{2d} \sum_{\ell=1}^d \Delta_\ell^2 \right). \quad (69)$$

By the Law of Large Numbers, as $d \rightarrow \infty$, we have

$$\frac{1}{2d} \sum_{\ell=1}^d \Delta_{\ell}^2 \rightarrow \frac{1}{2} \mathbb{E}(\Delta_{\ell}^2), \quad (70)$$

and so by the Continuous Mapping Theorem we get

$$k_{ij} \rightarrow \exp\left(-\frac{1}{2} \mathbb{E}(\Delta_{\ell}^2)\right) \quad (71)$$

for large d . On the other hand, the diagonal entries of the kernel matrix are all surely 1.

In the case of Figure 4, we have $(x_i)_{\ell} \sim U(-2\pi, 2\pi)$. Hence,

$$\mathbb{E}(\Delta_{\ell}^2) = \frac{4}{3} \pi^2, \quad (72)$$

so for d large, we have $k_{ij} \rightarrow \exp\left(-\frac{2\pi^2}{3}\right)$. For fixed N and large d , we have the approximation

$$\mathbf{K}_{\mathbf{x}\mathbf{x}} \rightarrow \left(1 - \exp\left(-\frac{2\pi^2}{3}\right)\right) \mathbf{I} + \exp\left(-\frac{2\pi^2}{3}\right) \mathbf{1}\mathbf{1}^T. \quad (73)$$

The spectrum of the matrix on the right hand side is

$$\lambda(\mathbf{K}_{\mathbf{x}\mathbf{x}}) = \left(1 + (N-1) \exp\left(-\frac{2\pi^2}{3}\right), 1 - \exp\left(-\frac{2\pi^2}{3}\right), \dots, 1 - \exp\left(-\frac{2\pi^2}{3}\right)\right). \quad (74)$$

Since, in our experiments, $\exp\left(-\frac{2\pi^2}{3}\right) \approx \frac{1}{N}$, the condition number of $\mathbf{K}_{\mathbf{x}\mathbf{x}}$ will be small. Since the entries of \mathbf{A} are small, this implies the condition number of $\mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{A}$ will be small. While this argument only holds asymptotically, one might expect the Central Limit Theorem to apply relatively quickly since the inputs are light-tailed, thus explaining observed behavior for, say, $d = 8$ in Figure 4.

Appendix B. Inducing Points

Here we develop the inducing-point-specific theory used in the manuscript.

B.1 Clustering-based Inducing Points

Proposition 15. *Let $u_j = \frac{1}{N_{\text{cl}(z_j)}} \sum_{\text{cl}(x_i)=z_j} y_i$. The Bayesian models defined by $f \sim \text{GP}(0, k)$ and*

$$y_i | f \sim \text{N}(f(\text{cl}(x_i)), \sigma^2) \quad u_i | f \sim \text{N}\left(f(z_i), \frac{\sigma^2}{N_{\text{cl}(z_i)}}\right). \quad (23)$$

admit respective posterior distributions $f | \mathbf{y}$ and $f | \mathbf{u}$ which are equal in distribution.

Proof. The likelihood can be written

$$p(y_i | f(x_i)) \propto \exp\left(-\sum_{i=1}^N \frac{(y_i - f(\text{cl}(x_i)))^2}{2\sigma^2}\right) = \exp\left(-\sum_{j=1}^M \sum_{\text{cl}(x_i)=z_j} \frac{(y_i - f(z_j))^2}{2\sigma^2}\right) \quad (75)$$

and the claim follows by noting that by Gaussianity, each $\left(u_j, \frac{1}{N_{\text{cl}(z_j)}}\right)$ pair is a sufficient statistic for all y_i whose indices satisfy $\text{cl}(x_i) = z_j$. \square

B.2 Stochastic Maximum Marginal Likelihood

We train hyperparameters using maximum marginal likelihood via doubly stochastic gradient descent. The minimization problem resembles the usual variational inference problem in most inducing point approaches, and is

$$\arg \min_{q_f \in \mathbb{Q}} D_{\text{KL}}(q_f \parallel \pi_f) - \frac{N}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \mathbb{E}_{f \sim q_f} (\mathbf{y} - f(\mathbf{x}))^T (\mathbf{y} - f(\mathbf{x})) \quad (76)$$

where \mathbb{Q} is the space of all Gaussian processes of the form

$$(f \mid \mathbf{u})(\cdot) = f(\cdot) + \mathbf{K}_{(\cdot)\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}(\mathbf{u} - f(\mathbf{z}) - \boldsymbol{\varepsilon}) \quad (77)$$

and $\mathbf{z} = \text{cl}(\mathbf{x})$, $u_j = \frac{1}{N_{\text{cl}(z_j)}} \sum_{\text{cl}(x_i)=z_j} y_i$, and $\Lambda_{ii} = \frac{\sigma^2}{N_{\text{cl}(z_i)}}$. The only parameters which are minimized numerically are the kernel hyperparameters. We simplify the variational expectation term using the identity

$$\mathbb{E}(\mathbf{y} - f(\mathbf{x}))^T (\mathbf{y} - f(\mathbf{x})) = \sum_{i=1}^N \mathbb{E}(y_i - f(x_i))^2 = \sum_{i=1}^N (y_i - \mathbb{E}(f(x_i)))^2 + \text{Var}(f(x_i)). \quad (78)$$

To train the process, we need to do two things: compute the prior Kullback–Leibler divergence, and compute the likelihood term. For the latter, we sample a mini-batch of data, and use it compute an unbiased estimator of the sum by evaluating the inner terms on the mini-batch. For a mini-batch \mathbf{x}' , the mean and variance terms are

$$\mathbb{E}(f(\mathbf{x}')) = \mathbf{K}_{\mathbf{x}'\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{u} \quad \text{Var}(f(\mathbf{x}')) = \mathbf{K}_{\mathbf{x}'\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{x}'} \quad (79)$$

which can be computed using matrix-vector products whose cost, in the case of the variance, is quadratic for each data point in the mini-batch. For the prior term, we have the identity

$$2D_{\text{KL}}(q_f \parallel \pi_f) = \ln \frac{|\mathbf{K}_{\mathbf{z}\mathbf{z}}|}{|\mathbf{K}_{\mathbf{z}\mathbf{z}} - \mathbf{K}_{\mathbf{z}\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}|} - d \quad (80)$$

$$+ \text{tr}(\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}(\mathbf{K}_{\mathbf{z}\mathbf{z}} - \mathbf{K}_{\mathbf{z}\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}})) \quad (81)$$

$$+ \mathbf{u}^T(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}\mathbf{K}_{\mathbf{z}\mathbf{z}}^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{u} \quad (82)$$

$$= \ln \frac{|\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda}|}{|\mathbf{\Lambda}|} - \text{tr}((\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}) \quad (83)$$

$$+ \mathbf{u}^T(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{u}. \quad (84)$$

This means that all linear systems that need to be solved involve the well-conditioned matrix $(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}$. Unfortunately, $(\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}$ has a matrix-valued rather than vector-valued right-hand-side, and thus requires M total matrix-vector products, leading to cubic costs, even with conjugate gradients. To avoid this, we apply Hutchinson’s trace estimator, which we now describe.

To illustrate Hutchinson’s trace estimator, let \mathbf{A} an arbitrary matrix, for instance $\mathbf{A} = (\mathbf{K}_{\mathbf{z}\mathbf{z}} + \mathbf{\Lambda})^{-1}\mathbf{K}_{\mathbf{z}\mathbf{z}}$. We are interested in evaluating $\text{tr}(\mathbf{A})$ in cases where we do not have direct access to \mathbf{A} , but can compute matrix-vector products with \mathbf{A} . Hutchinson (1989) observed that for any vector \mathbf{v} such that $\mathbb{E}(\mathbf{v}\mathbf{v}^T) = \mathbf{I}$ we have

$$\text{tr}(\mathbf{A}) = \text{tr}(\mathbb{E}(\mathbf{v}\mathbf{v}^T)\mathbf{A}) = \mathbb{E}(\mathbf{v}^T\mathbf{A}\mathbf{v}). \quad (85)$$

This motivates one to employ the unbiased Monte Carlo estimator

$$\frac{1}{L} \sum_{\ell=1}^L \mathbf{v}_{\ell}^T \mathbf{A} \mathbf{v}_{\ell} \quad (86)$$

where the \mathbf{v}_ℓ are independent and identically distributed. Typically, \mathbf{v}_ℓ is chosen to be Rademacher distributed, where each entry of \mathbf{v} is independent and takes the values $\{-1, 1\}$ with equal probability. This choice of \mathbf{v} leads to the minimum variance—see Hutchinson (1989, Proposition 1).

In the case $\mathbf{A} = (\mathbf{K}_{zz} + \mathbf{\Lambda})^{-1}\mathbf{K}_{zz}$, we sample the Rademacher random variables, solve the linear system $(\mathbf{K}_{zz} + \mathbf{\Lambda})^{-1}\mathbf{v}$, then multiply the result by $\mathbf{K}_{zz}\mathbf{v}$. The complexity of the matrix-vector products required for both operations is quadratic.

The log-determinant term is handled similarly. Rather than evaluate it directly, we follow Gibbs and Mackay (1997) and apply stochastic estimation on its gradient. Using that

$$\frac{\partial |\mathbf{K}_{zz} + \mathbf{\Lambda}|}{\partial \theta} = \text{tr} \left(\frac{\partial (\mathbf{K}_{zz} + \mathbf{\Lambda})}{\partial \theta} (\mathbf{K}_{zz} + \mathbf{\Lambda})^{-1} \right) \quad (87)$$

we apply Hutchinson’s trace estimator as in the case of the trace discussed above.

Appendix C. Cover Trees

To prove the key claim, we formalize the cover tree as a labeled rooted tree in the sense of a directed graph, defined inductively by the steps of the algorithm. We distinguish individual nodes from their locations in space, which we view as labels.

Definition 28. Let $T = (V, E, v_0)$ be a rooted tree, with root node $v_0 \in V$. Define the TREE ORDER \prec by $v \prec v'$ if the unique path from v_0 to v' passes through v . Define the PARENT of a node to be

$$\text{pa} : V \setminus \{v_0\} \rightarrow V \quad (88)$$

which maps each v to the unique vertex $\text{pa}(v) \prec v$ for which $(v, \text{pa}(v)) \in E$, which is well-defined because T is a rooted tree. Define the CHILDREN of a node to be

$$\text{ch} : V \rightarrow 2^V \quad \text{ch}(v) = \text{pa}^{-1}(\{v\}) \quad (89)$$

which is the set-theoretic preimage of the parent function. Define the LEVEL

$$\text{lv} : V \rightarrow \mathbb{N}_0 \quad (90)$$

of a node to be function which maps each v to the length $\text{lv}(v)$ of the path from v_0 to v , which is also well-defined because T is a rooted tree.

With these notions, we can define an intermediate state of the cover tree algorithm.

Definition 29. We say that a COVER TREE is a tuple (T, z, \mathcal{A}) , where T is a rooted tree and

$$z : V \rightarrow \mathbf{x} \quad \mathcal{A} : V \rightarrow 2^{\mathbf{x}} \quad (91)$$

are, respectively, the SPATIAL LOCATION and ASSIGNED DATA of each node.

The cover tree algorithm, then, is a function which maps a cover tree into another cover tree, by constructing the next level. Our goal is to prove that this algorithm terminates in finite time, producing a rooted tree whose nodes satisfy the desired properties. We analyze a more general form of the algorithm parameterized by a constant $a > 1$ and sequence $b_\ell > 0$, where a determines the rate at which the radius of each ball decays, and b_ℓ determines the size of an R -neighbor ball at level ℓ , with Algorithm 12 corresponding to $a = 2$ and $b_\ell = 4(1 - 1/2^\ell)R_0$, which will be derived as the optimal sequence for the given choice of a .

Definition 30. Let $V_0 = \{v_0\}$, $E_0 = \emptyset$, $z_0(v_0) = \frac{1}{N} \sum_{i=1}^n x_i$, and $\mathcal{A}_0 = \mathbf{x}$. Let $(T_\ell, z_\ell, \mathcal{A}_\ell)$ be the output of one step of Algorithm 12, defined inductively starting from $\ell = 0$.

From this, we can define the inducing points produced by the algorithm.

Definition 31. For a given tree T , define the LEVEL- ℓ NODES and LEVEL- ℓ INDUCING POINTS

$$V_\ell(T) = \{v \in V : \text{lv}(v) = \ell\} \quad \mathbf{z}_\ell(T) = \{z(v) : v \in V : \text{lv}(v) = \ell\}. \quad (92)$$

We first verify that the nodes and hence inducing points at a fixed level remain fixed as the algorithm proceeds.

Lemma 32. For $\ell' \geq \ell$, we have

$$V_\ell(T_{\ell'}) = V_\ell(T_\ell) \quad \text{sep}(\mathbf{z}_\ell(T_{\ell'})) = \text{sep}(\mathbf{z}_\ell(T_\ell)) \quad \text{res}_\mathbf{x}(\mathbf{z}_\ell(T_{\ell'})) = \text{res}_\mathbf{x}(\mathbf{z}_\ell(T_\ell)). \quad (93)$$

Proof. The second and third claims follow from the first one, which is immediate by the construction of Algorithm 12. \square

Next, we check that the separation and spatial resolution properties hold at the root node.

Lemma 33. We have

$$\text{sep}(\mathbf{z}_0(T_0)) = \infty \quad \text{res}_\mathbf{x}(\mathbf{z}_0(T_0)) \leq \max_{x \in \mathbf{x}} \|x - z(v_0)\|. \quad (94)$$

Proof. Because $V_0 = \{v_0\}$, the supremum in the definition of sep is taken over an empty set, and the first claim follows. For the second claim, recall that $z(v_0) = \frac{1}{|\mathbf{x}|} \sum_{x' \in \mathbf{x}} x'$. For any $x \in \mathbf{x}$,

$$\|x - z(v_0)\| = \frac{1}{|\mathbf{x}|} \left\| \sum_{x' \in \mathbf{x}} x - x' \right\| \leq \frac{1}{|\mathbf{x}|} \sum_{x' \in \mathbf{x}} \|x - x'\| \leq \max_{x' \in \mathbf{x}} \|x - x'\|. \quad (95)$$

Taking a maximum over $x \in \mathbf{x}$ completes the proof. \square

We now define the notion of R -neighbors.

Definition 34. Define $R_\ell = \frac{R_{\ell-1}}{a}$ inductively, where $R_0 = \max_{x \in \mathbf{x}} \|x - z(v_0)\|$. For $b > 0$, define the R_ℓ -NEIGHBORS of a node inductively by

$$\mathcal{R}(v_0) = \{v_0\} \quad \mathcal{R}(v) = \left\{ v' \in \bigcup_{s \in \mathcal{R}(\text{pa}(v))} \text{ch}(s) : \|z(v) - z(v')\| \leq b_\ell R_\ell \right\} \quad (96)$$

The next step is to deduce what sequences b_ℓ are valid for a given value of a .

Lemma 35. Let

$$b_\ell = \frac{2(1 + c(\frac{1}{a})^\ell)}{1 - \frac{1}{a}}. \quad (97)$$

for $c \in \{-1, -a^L\}$. Then, letting

$$\mathcal{R}'(v) = \{v' \in V_\ell(T_\ell) : \|z(v) - z(v')\| \leq b_\ell R_\ell\} \quad (98)$$

be the set of R_ℓ -nearby nodes, we have

$$\mathcal{R}'(v) = \mathcal{R}(v). \quad (99)$$

Proof. First, note for all $v \in V_\ell(T_\ell)$ that $\mathcal{R}(v) \subseteq \mathcal{R}'(v)$, since

$$\bigcup_{s \in \mathcal{R}(\text{pa}(v))} \text{ch}(s) \subseteq V_\ell(T_\ell) \quad (100)$$

since the set of children of the parents of a node are contained in the same level as the original node. It therefore suffices to show $\mathcal{R}'(v) \subseteq \mathcal{R}(v)$. The proof of this proceeds by induction on depth. In the base case $\ell = 0$ there is a single node v_0 , and $v_0 \in \mathcal{R}(v_0) = \{v_0\}$, so the proposition is true. We take the inductive hypothesis that

$$\mathcal{R}'(v'') \subseteq \mathcal{R}(v'') \quad (101)$$

for all $v'' \in V_{\ell-1}(T_{\ell-1}) = V_{\ell-1}(T_\ell)$, where the equality follows from Lemma 32. We now show the claim holds for $v \in V_\ell(T_\ell)$. Let

$$v' \in \mathcal{R}'(v) \quad (102)$$

for which we would like to show that $v' \in \mathcal{R}(v)$. The strategy will be to consider these nodes' parents, and show that $\text{pa}(v') \in \mathcal{R}(\text{pa}(v))$, in which case the result follows from the definition of \mathcal{R} . By the combination of the inductive hypothesis and first part of the proof, we have $\mathcal{R}'(v'') = \mathcal{R}(v'')$ for all $v'' \in V_{\ell-1}$, so it suffices to show

$$\|z(\text{pa}(v')) - z(\text{pa}(v))\| \leq b_{\ell-1}R_{\ell-1}. \quad (103)$$

We therefore upper bound the distance between the locations of the parents of the nodes in question. We have

$$\|z(\text{pa}(v')) - z(\text{pa}(v))\| \leq \|z(\text{pa}(v')) - z(v')\| + \|z(\text{pa}(v)) - z(v)\| + \|z(v') - z(v)\| \quad (104)$$

$$\leq \|z(\text{pa}(v')) - z(v')\| + \|z(\text{pa}(v)) - z(v)\| + b_\ell R_\ell \quad (105)$$

$$\leq 2R_{\ell-1} + b_\ell R_\ell. \quad (106)$$

$$= 2R_{\ell-1} + \frac{b_\ell}{a}R_{\ell-1}. \quad (107)$$

The first inequality by the triangle inequality. The second inequality follows by definition of v' . The third inequality uses that the locations of the children of a node are contained in the assigned data of the node, and that all points contained in the assigned data of a node in level $\ell - 1$ must be within a distance of $R_{\ell-1}$ of the label, or in equations that for all $v'' \in V_{\ell-1}(T_\ell)$ and $x \in \mathcal{A}_{\ell-1}(v)$ we have

$$\|z(v'') - x\| \leq R_{\ell-1}. \quad (108)$$

Combining (107) and (103) gives non-homogenous linear recurrence relation

$$2 + \frac{1}{a}b_\ell \leq b_{\ell-1}. \quad (109)$$

The solution can be derived using standard results on recurrence relations (Greene and Knuth, 1990, Chapter 2), which gives

$$b_\ell = \frac{2(1 + c(\frac{1}{a})^\ell)}{1 - \frac{1}{a}} \quad (110)$$

for some constant c to be determined by the initial conditions—one can easily check by plugging in the formula that this satisfies the recurrence. We require that $b_\ell \geq 0$ for $1 \leq \ell \leq L$, where L is the maximum tree depth. The optimal choice of initial conditions given this constraint, in the sense of minimizing the size of neighborhoods while retaining guarantees, come from selecting $b_0 = 0$ or $b_L = 0$. The former gives $c = -1$, while the latter gives $c = -a^L$. This completes the proof. \square

We now argue that the spatial resolution of the algorithm satisfies the claimed inequality.

Lemma 36. *We have*

$$\text{res}_{\mathbf{x}}(z_{\ell}(T_{\ell})) \leq R_{\ell}. \quad (111)$$

Proof. By induction on depth, one can show that for every level, $\{\mathcal{A}_{\ell}(v)\}_{v \in V_{\ell}}$ is a partition of $\bigcup_{i=1}^N \{x_i\}$, where x_i are entries of \mathbf{x} . For $x \in X$, let $v_{x,\ell}$ denote the unique $v \in V_{\ell}$ such that $x \in \mathcal{A}_{\ell}(v_{x,\ell})$. Then

$$\min_{v \in V_{\ell}} \|x - z(v)\| \leq \|x - z(v_{x,\ell})\| \leq R_{\ell}, \quad (112)$$

where the second inequality uses that the active set is contained in the ball of radius R_{ℓ} centered at $z(v_{x,\ell})$, which holds by definition of Algorithm 12, specifically Line 13. If optional Voronoi repartitioning is performed, we further upper bound $\|x - z(v_{x,\ell})\|$ by $\|x - z(v'_{x,\ell})\|$, where $v'_{x,\ell}$ denotes the vertex to which x was assigned prior to the Voronoi step. The results follows from taking a maximum over $x \in \mathbf{x}$. \square

Next, we argue the separation distance inequality holds.

Lemma 37. *We have*

$$\text{sep}(z_{\ell}(T_{\ell})) \geq R_{\ell}. \quad (113)$$

Proof. We prove the contrapositive: suppose that there exist a $v, v' \in V_{\ell}$ such that $\|z(v) - z(v')\| \leq R_{\ell}$. By the reverse triangle inequality, this yields

$$R_{\ell} \geq \|z(v) - z(v')\| \quad (114)$$

$$\geq \|z(\text{pa}(v)) - z(\text{pa}(v'))\| - \|z(\text{pa}(v)) - z(v)\| - \|z(\text{pa}(v')) - z(v')\|. \quad (115)$$

On the other hand, we have,

$$\|z(\text{pa}(v)) - z(v)\| \leq R_{\ell-1} \quad \|z(\text{pa}(v')) - z(v')\| \leq R_{\ell-1} \quad (116)$$

since the assigned data of $\text{pa}(v)$ is contained in the ball of radius $R_{\ell+1}$ centered at v by Line 13 of Algorithm 12, and $z(v)$ is in the convex hull of the assigned data. Rearranging, we conclude that

$$\|z(\text{pa}(v)) - z(\text{pa}(v'))\| \leq 2R_{\ell-1} + R_{\ell} = (2 + 1/a)R_{\ell-1}. \quad (117)$$

From Lemma 35 and since for $a > 1$, $2 + 1/a \leq 2a/(a - 1)$, we conclude $\text{pa}(v') \in \mathcal{R}(\text{pa}(v))$. Without loss of generality, assume that v was added to the tree prior to v' , which is unambiguous because nodes are added sequentially. Then by Line 13 of Algorithm 12, and since $\text{pa}(v') \in \mathcal{R}(\text{pa}(v))$, any node selected in Line 8, and Line 10 if optional Lloyd's averaging is performed, must satisfy

$$\|z(v') - z(v)\| > R_{\ell} \quad (118)$$

which is a contradiction. \square

Lemma 38. *We have*

$$|\text{ch}(v)| \leq P_{B_1}^{\text{ext}, 1/a} \leq (2a + 1)^d. \quad (119)$$

Proof. Suppose $v \in V_{\ell}$. As $\mathcal{A}_{\ell}(v)$ is contained in the ball of radius R_{ℓ} centered at $z(v)$, for any $w \in \text{ch}(v)$, $z(w)$ is contained in this ball. On the other hand, since $w \in V_{\ell+1}$ and since $\text{sep}(z(V_{\ell+1})) \geq R_{\ell+1}$, balls of radius $R_{\ell+1}$ centered at $\text{ch}(v)$ form a $R_{\ell+1}$ (exterior) packing of a sphere of radius $R_{\ell+1}$. Since packing numbers are invariant to rescaling, this is the same as the packing number of the unit sphere by spheres of radius $1/a$. The first inequality then follows from the definition of an exterior packing, and the second from Vershynin (2018, Corollary 4.2.13). \square

Lemma 39. *We have*

$$|\text{ch}(v)| \leq P_{B_1}^{\text{ext}, 1/(4a)} \leq (8a + 1)^d. \quad (120)$$

Proof. The proof is identical to Lemma 38, upon noting that $\mathcal{R}(v)$ is contained in a ball of radius $4R_\ell$ centered at $z(v)$, as given in Line 16 of Algorithm 12. \square

We are now ready to prove the main result.

Theorem 13. *For a given target spatial resolution $\varepsilon > 0$, and dataset with $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ and $d_{\max} = \max_i \|x_i - \bar{x}\|$, the cover tree inducing point algorithm, with or without the optional steps, produces a tree with $L = \lceil \log_2 \frac{d_{\max}}{\varepsilon} \rceil + 1$ levels, terminates in*

$$\mathcal{O}\left(P_{B(4)}^{\text{ext},1/2} \left(P_{B(1)}^{\text{ext},1/2}\right)^2 N \log \frac{d_{\max}}{\varepsilon}\right) \quad (19)$$

steps, where $B(r)$ denotes a ball of radius r , and $P_X^{\text{ext},\delta}$ is the external packing number, namely the maximum number of disjoint balls of radius δ that one can choose so that their centers lie in X . Moreover, the algorithm guarantees that the nodes at each level ℓ satisfy

$$\text{res}_{\mathbf{x}}(\mathbf{z}_\ell) \leq 2^{L-\ell} \varepsilon \quad \text{sep}(\mathbf{z}_\ell) \geq 2^{L-\ell} \varepsilon. \quad (20)$$

Proof. The claim follows by combining the above results. \square

Appendix D. Experimental Details

All experiments were run on a single Nvidia V100 GPU with 32GB RAM in double precision, except for Figure 6 where floating-point precision was used instead. For the Wasserstein distance and condition number experiment of Figure 4, the dataset of size $N = 1000$ was generated by sampling from a GP prior with squared exponential kernel, where inputs were uniformly sampled on the d dimensional cube $[-5, 5]$ with d equal to 1, 2, 4 and 8. The length scale of the kernel was set to $0.5\sqrt{d}$. The parameters of the clustered-data approximation were set via a cover tree with spatial resolutions 0.05–4.0. Experiments for each d were repeated 20 times to assess variability.

In the surface temperature geospatial data illustrative example shown of Figure 6, we ran experiments with using the sparse Gaussian process regression (SGPR) implementation from GPflow (Matthews et al., 2017) and a GPflow-based implementation of the clustered-data Gaussian process using stochastic maximum marginal likelihood training described in Appendix B.2. For Hutchinson’s trace estimator, we used 10 probe vectors. For the two-dimensional East Africa land surface temperature dataset, we split the dataset into training and testing sets of size 55884 and 27525, respectively. Both Gaussian process models were initialized with 0.1 likelihood noise and configured with a squared exponential kernel with automatic relevance determination. The length scales of the kernel were initialized to 1.0. The inducing points were set using a cover tree with the spatial resolution 0.03–0.09. To facilitate comparisons, in the sparse Gaussian process baseline, the inducing points were excluded from the trainable parameter set. For both models hyperparameters were trained using mini-batch stochastic optimization via the Adam optimizer, with constant learning rate 0.01 and batch size 1000. We plot the posterior mean and standard deviation from the clustered-data Gaussian process in Figure 8.

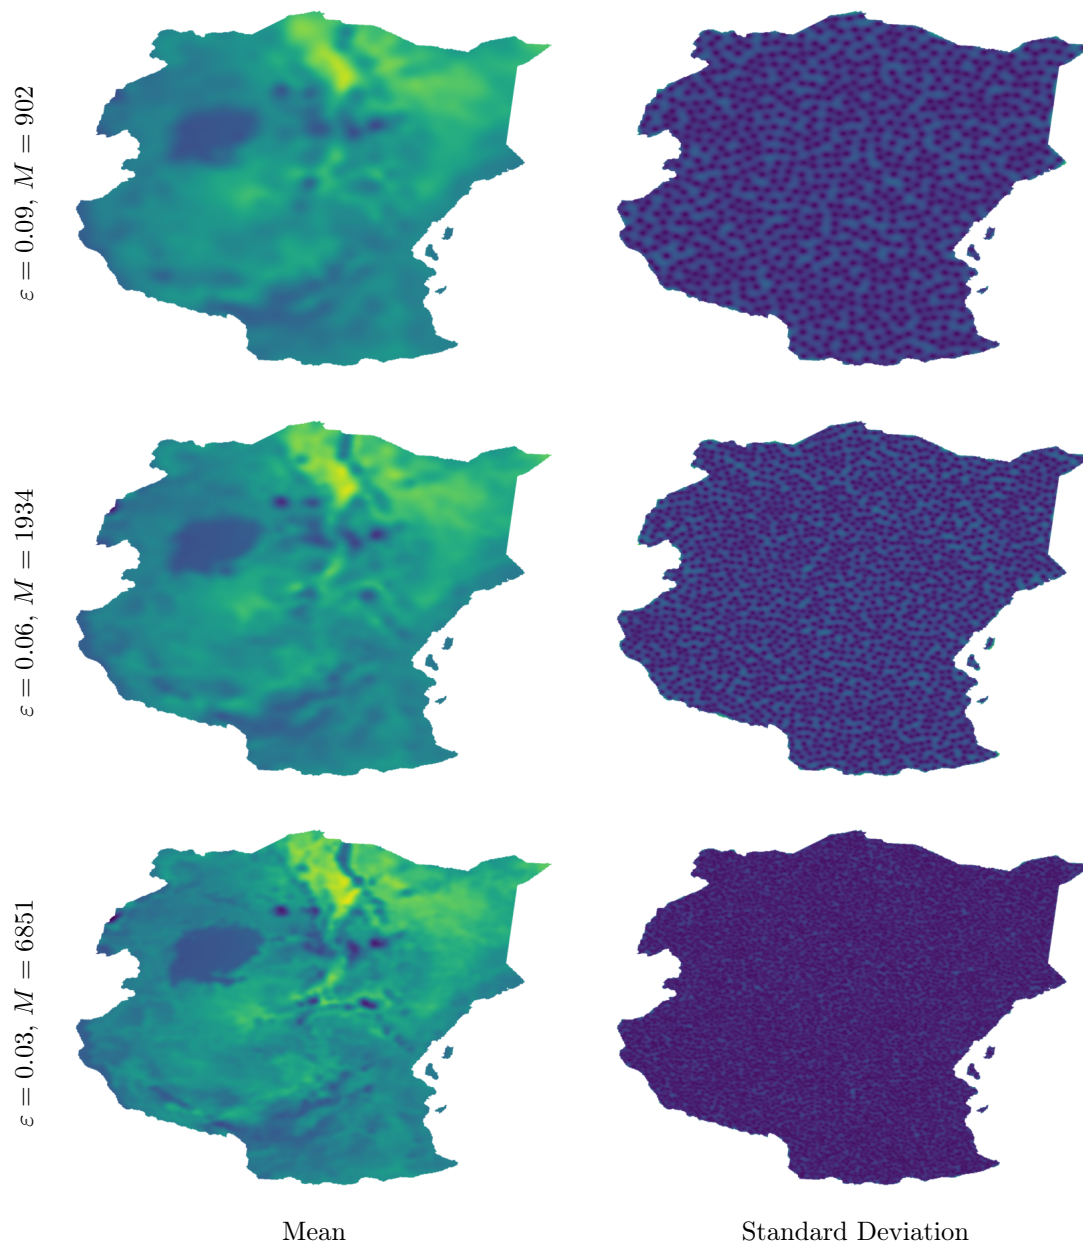


Figure 8: Predictions from the clustered-data Gaussian process approximation, with varied spatial resolutions ε . For each spatial resolution, we also display the resulting number of inducing points M .