
Revisiting the Sample Complexity of Sparse Spectrum Approximation of Gaussian Processes

Quang Minh Hoang

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, PA 15213
qhoang@andrew.cmu.edu

Trong Nghia Hoang

MIT-IBM Watson AI Lab
IBM Research
Cambridge, MA 02142
nghiaht@ibm.com

Hai Pham

Language Technologies Institute
Carnegie-Mellon University
Pittsburgh, PA 15213
htpham@cs.cmu.edu

David P. Woodruff

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, PA 15213
dwoodruf@cs.cmu.edu

Abstract

We introduce a new scalable approximation for Gaussian processes with provable guarantees which hold simultaneously over its entire parameter space. Our approximation is obtained from an improved sample complexity analysis for sparse spectrum Gaussian processes (SSGPs). In particular, our analysis shows that under a certain data disentangling condition, an SSGP’s prediction and model evidence (for training) can well-approximate those of a full GP with low sample complexity. We also develop a new auto-encoding algorithm that finds a latent space to disentangle latent input coordinates into well-separated clusters, which is amenable to our sample complexity analysis. We validate our proposed method on several benchmarks with promising results supporting our theoretical analysis.

1 Introduction

A Gaussian process (GP) [36] is a popular probabilistic kernel method for regression which has found applications across many scientific disciplines. Examples of such applications include meteorological forecasting, such as precipitation and sea-level pressure prediction [2]; sensing and monitoring of ocean and freshwater phenomena such as temperature and plankton bloom [7, 12]; traffic flow and mobility demand predictions over urban road networks [9, 10, 29]; flight delay predictions [15, 19, 20]; and persistent robotics tasks such as localization and filtering [43]. The broad applicability of GPs is in part due to their expressive Bayesian non-parametric nature which provides a closed-form prediction [36] in the form of a Gaussian distribution with formal measures of predictive uncertainty, such as entropy and mutual information criteria [27, 39, 44]. Such expressiveness makes GPs not only useful as predictive methods but also a go-to representation for active learning applications [24, 23, 27, 44] or Bayesian optimization [38, 45, 22, 16] that need to optimize for information gain while collecting training data.

Unfortunately, the expressive power of a GP comes at a cost of poor scalability (i.e., cubic time [36]) in the size of the training data (see Section 2.1 below), hence limiting its use to small datasets. This prevents GPs from being applied more broadly to modern settings with increasingly growing volumes of data [15, 19, 20]. To sidestep this limitation, a prevalent research trend is to impose sparse structural assumptions [33, 34] on the GP’s kernel matrix to reduce its multiplication and inversion cost, which comprises the main bulk of the training and inference complexity. This results in a broad family of

sparse Gaussian processes [15, 17, 19, 28, 37, 40] that are not only computationally efficient but also amenable to various forms of parallelism [8, 29] and distributed computation [1, 13, 20, 21, 18], further increasing their efficiency.

Despite such advantages, the sparsification components at the core of these methods are heuristically designed and do not come with provable guarantees that explicitly characterize the interplay between approximation quality and computational complexity. This motivates us to develop a more robust, theoretically-grounded approximation scheme for GPs that is both provable and amenable to the many fast computation schemes mentioned above. More specifically, our contributions include:

1. An analysis of a new approximation scheme that generates a sparse spectrum approximation of a GP with provable bounds on its sample complexity, which practically becomes significantly small when the input data exhibits a certain clustering structure. Furthermore, the impact of the approximation on the resulting training and inference qualities is also formally analyzed (Section 3.1).
2. A data partitioning algorithm inspired from the above analysis, which learns a cluster embedding that reorients the input distribution while ensuring reconstructability of the original distribution (Section 3.3). We show that using sparse spectrum Gaussian processes (SSGP) [17, 28] on the embedded space requires fewer samples to achieve the same level of approximation quality. This also induces a linear feature map which enables efficient training and inference of GPs.
3. An empirical study on benchmarks that demonstrates the efficiency of the proposed method over existing works in terms of its approximation quality versus computational efficiency (Section 4).

2 Related Work

In this section we provide an overview of Gaussian processes (Section 2.1), followed by a succinct summary of their spectral representations (Section 2.2).

2.1 Gaussian Processes (GPs)

A Gaussian process [36] defines a probabilistic prior over a random function $g(\mathbf{x})$ defined by mean function $m(\mathbf{x}) = 0^1$ and kernel function $k(\mathbf{x}, \mathbf{x}')$. These functions induce a marginal Gaussian prior over the evaluations $\mathbf{g} = [g(\mathbf{x}_1) \dots g(\mathbf{x}_n)]^\top$ on an arbitrary finite subset of inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. Let \mathbf{x}_* be an unseen input whose corresponding output $g_* = g(\mathbf{x}_*)$ we wish to predict. The Gaussian prior over $[g(\mathbf{x}_1) \dots g(\mathbf{x}_n) g(\mathbf{x}_*)]^\top$ implies the following conditional distribution:

$$g_* \triangleq g(\mathbf{x}_*) \mid \mathbf{g} \sim \mathbf{N}\left(\mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{g}, k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{k}_*\right), \quad (1)$$

where $\mathbf{k}_* = [k(\mathbf{x}_*, \mathbf{x}_1) \dots k(\mathbf{x}_*, \mathbf{x}_n)]^\top$ and \mathbf{K} denotes the Gram matrix induced by $k(\mathbf{x}, \mathbf{x}')$ on $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ for which $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. For a noisy observation y perturbed by Gaussian noise such that $y \sim \mathbf{N}(g(\mathbf{x}), \sigma^2)$, Eq. (1) above can be integrated with $\mathbf{N}(\mathbf{g}, \sigma^2 \mathbf{I})$ to yield:

$$g_* \triangleq g(\mathbf{x}_*) \mid \mathbf{y} \sim \mathbf{N}\left(\mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_*\right), \quad (2)$$

which explicitly forms the predictive distribution of a Gaussian process. The defining parameter Θ of $k(\mathbf{x}, \mathbf{x}')$ (see Section 2.2) is crucial to the predictive performance and needs to be optimized via minimizing the negative log likelihood of \mathbf{y} :

$$\ell(\Theta) = \frac{1}{2} \log |\mathbf{K}_\Theta + \sigma^2 \mathbf{I}| + \frac{1}{2} \mathbf{y}^\top (\mathbf{K}_\Theta + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \quad (3)$$

where we now use the subscript Θ to indicate that \mathbf{K} is a function of Θ . In practice, both training Θ and prediction incur $\mathcal{O}(n^3)$ processing cost, which prevents direct use of Gaussian processes on large datasets that might contain more than tens of thousands of training inputs.

2.2 Sparse Spectrum Gaussian Processes

Sparse spectrum Gaussian processes (SSGPs) [14, 17, 28] exploit Theorem 1 below to re-express the Gaussian kernel $k(\mathbf{x}, \mathbf{x}') \triangleq \exp(-0.5 (\mathbf{x} - \mathbf{x}')^\top \Theta^{-1} (\mathbf{x} - \mathbf{x}'))$ (where $\Theta \triangleq \text{diag}[\theta_1^2 \dots \theta_d^2]$) as an integration over a spectrum of cosine functions such that the integrating distribution (over the frequencies that parameterize these functions) is a multivariate Gaussian.

¹For simplicity, we assume a zero mean function since we can always re-center the training outputs around 0.

Theorem 1 (Bochner Theorem). *Let $k(\mathbf{x}, \mathbf{x}')$ denote a Gaussian kernel defined above and let $p(\mathbf{r}) \sim \mathcal{N}(\mathbf{0}, (4\pi^2\Theta)^{-1})$. It follows that:*

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\mathbf{r} \sim p(\mathbf{r})} \left[\cos\left(2\pi\mathbf{r}^\top(\mathbf{x} - \mathbf{x}')\right) \right], \quad (4)$$

where \mathbf{r} is a d -dimensional random variable that parameterizes $\cos(2\pi\mathbf{r}^\top(\mathbf{x} - \mathbf{x}'))$. In practice, \mathbf{r} is often referred to as the spectral frequency.

This allows us to approximate the original Gram matrix \mathbf{K} with a low-rank matrix \mathbf{K}' constructed by a linear kernel $\mathbf{K}'(\mathbf{x}, \mathbf{x}') = \Phi(\mathbf{x})^\top \Phi(\mathbf{x}')$ with feature map $\Phi(\mathbf{x}) = [\phi_1(\mathbf{x}) \dots \phi_{2m}(\mathbf{x})]^\top$ comprising $2m$ basis trigonometric functions [17]. Each pair of odd- and even-index basis functions $\phi_{2i-1}(\mathbf{x}) = \cos(2\pi\mathbf{r}_i^\top \mathbf{x})$ and $\phi_{2i}(\mathbf{x}) = \sin(2\pi\mathbf{r}_i^\top \mathbf{x})$ is parameterized by the same sample of spectral parameter $\mathbf{r}_i \sim p(\mathbf{r})$. For efficient computation, m is often selected to be significantly smaller than n (i.e., the number of training examples). However, to guarantee that $\|\mathbf{K} - \mathbf{K}'\|_2 \leq \lambda$ with probability at least $1 - \delta$, m needs to be as large as $\mathcal{O}(n^2/\lambda^2 \log(n/\delta))$ [31]², which makes the total prediction complexity much worse than that of a full GP.

Alternatively, one can use kernel sketching methods [3, 32, 35] to generate feature maps that scale more favorably with the effective dimension of the kernel matrix, which empirically tends to be on the order of $\mathcal{O}(\log n)$. However, the pitfall of these methods is that without knowing the exact parameter configuration Θ that underlies the data, they cannot sample from the true $p(\mathbf{r})$, which is necessary in their analyses. As such, existing random maps [3, 35] that were generated based on this spectral construction often depend on a parameter initialization and their approximation quality is only guaranteed for that particular parameter setting instead of uniformly over the entire parameter space. This motivates us to revisit the sample complexity of SSGP from a setting which specifically searches for a reorientation of the input distribution such that the reoriented data exhibits a disentangled cluster structure. Such disentanglement provides a more sample-efficient bound as we show in our analysis in Section 3.1 below.

3 Provable Approximation of SSGPs with Improved Sample Complexity

We first show how a sparse spectrum Gaussian process (SSGP) [28] can be approximated well with a provably low sample complexity. This is achieved by revisiting its sample complexity which, unlike prior work [3, 31, 35], explicitly characterizes and accounts for a certain set of data disentanglement conditions. Importantly, our new analysis (Section 3.1) yields practical bounds on both an SSGP’s prediction and model evidence (Section 3.2) that hold with high probability uniformly over the entire parameter space³. Furthermore, our analysis also inspires an encoding algorithm that finds a latent space to disentangle the encoded coordinates of data into well-separated clusters on which a sparse spectrum GP can approximate a GP provably well (Section 3.3). Our experiments show that such a latent space can be found for several real-world datasets (Section 4).

3.1 Practically Improved Sample Complexity for Sparse Spectrum Gaussian Processes

This section derives a new data-oriented feature map to approximate a Gaussian process parameterized with a Gaussian kernel. Unlike existing work which assumes knowledge of the true kernel parameters [3, 32, 35], our derivation remains oblivious to such parameters, and therefore holds universally over their entire candidate space. We assume that the GP prior of interest is of the form $g(\mathbf{x}) \sim \text{GP}(0, k(\mathbf{x}, \mathbf{x}'))$ where $k(\mathbf{x}, \mathbf{x}')$ represents its Gaussian kernel in Section 2.2.

We give our analysis in three parts: (1) the spectral sampling scheme and a notion of approximation loss; (2) a set of practical data conditions which can be either observed from a raw data distribution or approximately imposed on the data via a certain embedding; (3) a theoretical analysis that delivers our key result that establishes an improved sample complexity when our data conditions are met.

3.1.1 Spectral Sampling Scheme and Spectral Loss

We show that $g(\mathbf{x})$ can be approximated by $g'(\mathbf{x}) = \sum_{i=1}^p g_i(\mathbf{x})$ with provable data-oriented guarantees where $g_i(\mathbf{x}) \sim \text{GP}(0, (1/\sqrt{p})k_i(\mathbf{x}, \mathbf{x}'))$. To achieve this, we first establish in Lemma 1

²See Theorem 6.28 in Chapter 6 of [31].

³In contrast, existing literature often generates bounds on either an SSGP’s prediction or its model evidence (for training) for a single parameter configuration, which makes such an analysis only heuristic.

that the induced Gram matrix \mathbf{K} of $k(\mathbf{x}, \mathbf{x}')$ on any dataset can be represented as an expectation over a space of induced Gram matrices $\{\mathbf{K}_i\}_{i=1}^p$ produced by a corresponding space of random kernels $\{k_i(\mathbf{x}, \mathbf{x}')\}_{i=1}^p$.

Lemma 1. *Let $k(\mathbf{x}, \mathbf{x}')$ and \mathbf{K} denote a Gaussian kernel parameterized by Θ (Section 2.2) and its induced Gram matrix on an arbitrary set of training inputs, respectively. There exists a space \mathcal{K} of random kernels $\kappa(\mathbf{x}, \mathbf{x}')$ and a Θ -independent distribution ρ over \mathcal{K} for which $\mathbf{K} = \mathbb{E}_\rho[\mathbf{K}_\kappa]$ where \mathbf{K}_κ denotes the induced Gram matrix of κ on the same set of training inputs.*

This follows directly from Theorem 1 above which states that $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[\cos(2\pi\mathbf{r}^\top(\mathbf{x}-\mathbf{x}'))]$ where $\mathbf{r} \sim \mathbf{N}(\mathbf{0}, (4\pi^2\Theta)^{-1})$. We can choose $\kappa(\mathbf{x}, \mathbf{x}'; \epsilon) = \cos(\epsilon^\top\Theta^{-0.5}(\mathbf{x}-\mathbf{x}'))$ where $\epsilon \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ which implies $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_\epsilon[\kappa(\mathbf{x}, \mathbf{x}'; \epsilon)]$. Thus, $\mathbf{K} = \mathbb{E}_\epsilon[\mathbf{K}_\epsilon]$ where the Θ -independent parameter ϵ indexes κ and \mathbf{K}_ϵ is the induced Gram matrix of κ . Leveraging the result of Lemma 1, a naïve analysis [31] using worst-case concentration bounds to derive a conservative estimate for a sufficient number of samples would require a prohibitively expensive sample complexity of $\mathcal{O}(n^2 \log n)$.

Such analyses, however, often ignore the input distribution, which can be used to sample more selectively, thereby significantly reducing the sample complexity. This is demonstrated below in Theorem 2 which shows that when the input distribution exhibits a certain degree of compactness and separation (as defined in Conditions 1-3), we only require $\mathcal{O}((\log^2 n/\lambda^2) \log \log(n/\delta))$ sampled kernels $\{\kappa_i\}_{i=1}^p$ indexed by $\{\epsilon_i\}_{i=1}^p$ to produce an average Gram matrix $\mathbf{K}' = \frac{1}{p} \sum_{i=1}^p \mathbf{K}_{\epsilon_i}$ that is sufficiently close to \mathbf{K} in spectral norm (see Definition 1) with probability at least $1 - \delta$.

Definition 1 (Spectral Closeness). *Given $\lambda > 0$, the symmetric matrices \mathbf{K} and \mathbf{K}' are λ -close if $\|\mathbf{K} - \mathbf{K}'\|_2 \leq \lambda$ where $\|\mathbf{K} - \mathbf{K}'\|_2 = \lambda_{\max}(\mathbf{K} - \mathbf{K}')$ denotes the largest eigenvalue of $\mathbf{K} - \mathbf{K}'$.*

Thus, parameterizing the GP prior with \mathbf{K}' instead of \mathbf{K} allows us to derive an upper bound on the expected difference between their induced model evidence (for learning kernel parameters) and prediction losses (for testing) with respect to the same parameter setup (Theorem 3). Theorem 3 importantly exploits the fact that the bound in Theorem 2 holds universally over the entire space of parameters, which allows us to bound the prediction difference between the original and approximated GPs with respect to their own optimized parameters (that are not necessarily the same).

3.1.2 Practical Conditions on Data Distributions

We now outline key practical data conditions, which can be satisfied approximately via an encoding algorithm that transforms the input data into a latent space where such conditions are met. These conditions are necessary for deriving a practically improved sample complexity in Section 3.1.3.

Condition 1. For each parameter configuration $\Theta = \text{diag}[\theta_1^2, \dots, \theta_d^2]$, there exists a mixture distribution $\mathcal{M}(\mathbf{x}; \boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_b), \boldsymbol{\pi} = (\pi_1, \dots, \pi_b), \mathbf{c} = (\mathbf{c}_1, \dots, \mathbf{c}_b))$ with at most $b = \mathcal{O}(\log n)$ Gaussian components $\mathbf{N}(\mathbf{x}; \mathbf{c}_i, \gamma_i^2 \Theta^{-1})$ over the data space with the mixing weights $\pi_i \propto 2^{\frac{i}{2}}$ and variances $\gamma_i = \mathcal{O}(\frac{1}{\sqrt{d}})$ that generate the observed data in d -dimensional space.

Condition 2. The i^{th} Gaussian component as defined in Condition 1 above was used to generate $2^{\frac{i}{2}}$ data points of the observed dataset. This can be substantiated easily with high probability given the above setup in Condition 1 that assigns selection probability $\pi_i \propto 2^{\frac{i}{2}}$ to the i^{th} -component.

Condition 3. For each parameter configuration $\Theta = \text{diag}[\theta_1^2, \dots, \theta_d^2]$, the mixture distribution of data in Condition 1 has sufficiently separated cluster centers. That is, for all $i \neq j$:

$$\left\| \Theta^{-1/2} (\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 > \frac{3}{2} \log \left(\frac{2^a}{2^a - 1} \right) \quad \text{where} \quad a = \frac{1}{\log 2} \log \left(\frac{n^4}{n^4 - \lambda^4} \right). \quad (5)$$

These conditions impose that the observed data can be separated into a number of clusters with exponentially growing sizes and concentration (see the small variances defined in Condition 1 and the imposed sizes of Condition 2). Intuitively, this means data points that belong to clusters with high concentration are responsible for kernel entries with high values whereas those in clusters with low concentration generate entries with low values. This is easy to see since high concentration reduces the distance between data points, thus increasing their kernel values and vice versa.

Furthermore, as imposed by Condition 2, clusters with high concentration also have denser population and induce kernel entries with high value. In addition, Condition 3 requires that clusters are well-separated, which implies that a large number of kernel entries are small and therefore can be

approximated cheaply. Together, these conditions form the foundations of our reduced complexity analysis for SSGP in Theorem 2. Interestingly, we show that such conditions also inspire the development of a probabilistic algorithm that finds an encoding of the input that (approximately) satisfies these conditions while preserving the statistical properties of the input (Section 3.3). This results in an improved sample complexity for SSGPs in practice (see Section 3.1.3).

3.1.3 Main Results

To understand the intuition why an improved sample complexity can be obtained, we note that when data is partitioned in clusters with different concentrations and sizes, the kernel entries are also partitioned into multiple value-bands with narrow width (i.e., low variance). Exploiting this, we can calibrate a significantly lower sample complexity for each band using concentration inequalities that improve with lower variance [11, 25].

Then, to combine these in-band sample complexities efficiently, we further exploit the data conditions in Section 3.1.2 to show that statistically, value bands with smaller width also tend to be populated more densely⁴. This allows us to aggregate these in-band sample costs into an overall sample complexity with low cost. In practice, this also inspires an embedding algorithm (Section 3.3) that transforms the data in such a way that the distribution of their induced kernel entries will be denser in narrower bands, which is advantageous in our analysis.

Formally, let \mathcal{C} be the set of all kernel entries indexed by (u, v) in the Gram matrix \mathbf{K} such that \mathbf{x}_u and \mathbf{x}_v belong to the same cluster and \mathcal{C}' be its complement. Also, let \mathcal{C} be partitioned⁵ into b value-bands $\kappa_i = \{(u, v) \in \mathcal{C} \mid 1 - \mathbf{O}(2^{1-i}) \leq \mathbf{K}_{uv}^4 \leq 1 - \mathbf{O}(2^{-i})\}$ for $i \in [1 \dots b]$ and let $\kappa_0 = \{(u, v) \in \mathcal{C} \mid \mathbf{K}_{uv}^4 \geq 1 - \mathbf{O}(2^{-b})\}$ be a band that is only populated by very large kernel entries. Theorem 2 below shows that we can construct a λ -spectral approximation of \mathbf{K} with arbitrarily high probability and low sample complexity.

Theorem 2. *For any $1 \geq \delta \geq \mathbf{O}(\exp(b - \sqrt{d}))$, if the training data has n data points and satisfies Conditions 1-3 above with respect to λ , then with probability at least $1 - 2\delta$, the approximation $\mathbf{K}' = (1/p) \sum_{i=1}^p \mathbf{K}_{\epsilon_i}$ where $\epsilon_i \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ is λ -spectral close to \mathbf{K} .*

Proof Sketch. Our proof strategy is outlined below. The formal statements are spelled out in Appendix A.

First, with a proper choice of a clustering partition, the cross-cluster entries in \mathbf{K} are guaranteed to be sufficiently small so as to be well-approximated by zero. We can then show with high probability that any kernel entry that corresponds to a pair of unique data points from the same cluster can be well-approximated with a sample complexity that scales favorably with the cluster's variance. In particular, we show that kernel values induced by data points generated by lower-variance clusters (see Condition 1) will have smaller approximation variances than those generated by data from higher-variance clusters and therefore require fewer samples to produce the same level of approximation.

Second, for certain configurations of mixture weights, Condition 2 asserts that the number of data points from each cluster is inversely proportional to the cluster variance, which implies that a small sample complexity is enough to approximate the majority of kernel entries. More specifically, Lemma 3 shows that when the input points are distributed into clusters with certain choices of variances $\{\gamma_i\}_{i=1}^b$ and at an inversely proportional ratio $\mathbf{O}(\gamma_i^{-1})$, then with high probability, over all clusters, the kernel entries (excluding those on the diagonal) associated with pairs in the i -th cluster belong to their corresponding band κ_i .

Lemma 5 shows that for $p = \mathbf{O}(\log^2 n / \lambda^2 \log(\log n / \delta))$, with probability $1 - \delta/b$, the total approximation error of all kernel entries in the \mathcal{C}_i will be at most $\lambda^2/4b$, which implies with probability $1 - \delta$, the total approximation cost for items in \mathcal{C} is at most $\lambda^2/4$. Next, Lemma 2 establishes that with the above data distribution, \mathcal{C} accounts for $n^2/4$ entries while \mathcal{C}' accounts for $3n^2/4$ entries, which needs to be approximated with error at most $3\lambda^2/4$.

Finally, Lemma 4 shows that when the clusters are sufficiently well-separated (see Condition 3), any kernel value corresponding to an arbitrary data pair with points belonging to different clusters is guaranteed to be smaller than λ^2/n^2 , which then guarantees a total error of at most $3\lambda^2/4$ when they are uniformly approximated with zero. Putting these together yields a total error of λ^2 with probability

⁴The intuition here is that kernel entries in narrower bands are cheaper (in term of sample cost) to approximate.

⁵The exact bounds defining the band can be found in Appendix A.

$1 - 2\delta$, which implies \mathbf{K} and \mathbf{K}' are λ -spectrally close since $\|\mathbf{K} - \mathbf{K}'\|_2 \leq \|\mathbf{K} - \mathbf{K}'\|_F \leq \lambda$. Please see Appendix A for details.

3.2 Approximation Loss for Prediction and Model Evidence

In terms of prediction and model evidence approximation, our result holds simultaneously for all parameter configurations and is thus oblivious to the choice of parameters (see Theorem 3). While existing kernel sketch methods [3, 32] generically achieve near-linear complexity for the approximate feature map⁶, they often require knowledge of the parameters to construct the kernel approximations. In contrast, our result in Theorem 2 can be leveraged to bound the same prediction discrepancy when the original and approximated GPs use their own optimized parameter configurations, as shown in Theorem 4 below. To establish Theorem 4, however, we first establish an intermediate result that bounds the prediction and model evidence in the case when both the original and approximated GPs use the same parameter configurations.

Theorem 3. *Let $\delta < 1$ be a user-specified confidence as defined previously in Theorem 2 and let \mathbf{K}' be an approximation to \mathbf{K} for which $\|\mathbf{K} - \mathbf{K}'\|_2^2 \leq \lambda^2$ with probability $1 - \delta$, uniformly over the entire parameter space. Then, with probability $1 - \delta$, the following hold:*

$$\mathbb{E}[g(\mathbf{x}_*)] = \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{E}[g'(\mathbf{x}_*)] \quad \text{and} \quad \mathbb{V}[g(\mathbf{x}_*)] = \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{V}[g'(\mathbf{x}_*)] \pm \frac{\lambda}{\sigma^2} \quad (6)$$

where σ^2 is the noise of the variance (Eq. (2)), and $g(\mathbf{x}_*)$, $g'(\mathbf{x}_*)$ respectively denote the predictive distributions of the full GP and the approximated GP pertaining to an arbitrary test input \mathbf{x}_* .

Proof. This follows directly from Lemma 7 and Lemma 8 in Appendix B. \square

Finally, Theorem 4 analyzes how close the approximated predictive mean is to the full GP predictive mean when both are evaluated at the optimizer of their respective training objective.

Theorem 4. *Let $\delta < 1$ be a user-specified confidence as defined in Theorem 2. Let \mathbf{K}' denote an approximation to \mathbf{K} for which $\|\mathbf{K} - \mathbf{K}'\|_2^2 \leq \lambda^2$ with probability at least $1 - \delta$ uniformly over the entire parameter space. Let Θ_* and Θ'_* denote the optimal hyperparameters obtained by respectively minimizing the negative log likelihood of a full GP and the approximated GP. With probability $1 - \delta$, the following holds:*

$$\mathbb{E}[g'(\mathbf{x}_*; \Theta'_*)] = (1 \pm \rho(\lambda, \sigma, \Theta_*, \Theta'_*)) \cdot \mathbb{E}[g(\mathbf{x}_*; \Theta_*)] + \wp(\lambda, \sigma, \Theta_*, \Theta'_*) \quad (7)$$

where $\rho(\lambda, \sigma, \Theta_*, \Theta'_*)$ and $\wp(\lambda, \sigma, \Theta_*, \Theta'_*)$ are constants with respect to $\lambda, \sigma, \Theta_*, \Theta'_*$.

Proof. This follows immediately from Lemma 11 in Appendix B, which was built on the result of Theorem 3 above. This completes our loss analysis for SSGPs. \square

3.3 Optimizing Feature Map Complexity

We next present a practical probabilistic embedding algorithm that transforms the input data to meet the requirements of Conditions 1-3. Our method is built on the rich literature of variational auto-encoders (VAE) [26], which is a broad class of deep generative models that combine the rigor of Bayesian methods and rich parameterization of (deep) neural networks to discover (non-linear) low-dimensional embeddings of data while preserving their statistical properties. We first provide a short review on VAEs below, followed by an augmentation that aims to achieve the impositions in Conditions 1-3 above.

3.3.1 Variational Auto-Encoders (VAEs)

Let \mathbf{x} be a random variable with density function $p(\mathbf{x})$. We want to learn a latent variable model $p_\theta(\mathbf{x}, \mathbf{z}) = p(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z})$ that captures this generative process. The latent variable model comprises a fixed latent prior $p(\mathbf{z})$ and a parametric likelihood $p_\theta(\mathbf{x}|\mathbf{z})$. To learn θ , we maximize the variational evidence lower-bound (ELBO) $\mathbf{L}(\mathbf{x}; \theta, \phi)$ of $\log p_\theta(\mathbf{x})$:

$$\mathbf{L}(\mathbf{x}; \theta, \phi) \triangleq \mathbb{E}_{\mathbf{z} \sim q_\phi} \left[\log p_\theta(\mathbf{x}|\mathbf{z}) \right] - \mathbb{KL} \left(q_\phi(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}) \right) \quad (8)$$

⁶[3, 32] achieves a complexity of $\mathcal{O}(nm^2)$ where m scales with the effective dimension of the kernel matrix.

with respect to an arbitrary posterior surrogate $q_\phi(\mathbf{z}|\mathbf{x}) \simeq p_\theta(\mathbf{z}|\mathbf{x})$ over the latent variable \mathbf{z} . The ELBO is always a lower-bound on $\log p_\theta(\mathbf{x})$ regardless of our choice of $q_\phi(\mathbf{z}|\mathbf{x})$. This is due to the non-negativity of the KL divergence as seen in the first part of the above equation.

This can be viewed as a stochastic auto-encoder with $p_\theta(\mathbf{x}|\mathbf{z})$ and $q_\phi(\mathbf{z}|\mathbf{x})$ acting as the encoder and decoder, respectively. Here, θ and ϕ characterize the neural network parameterization of these models. Their learning is enabled via a re-parameterization of $q_\phi(\mathbf{z}|\mathbf{x})$ that enables stochastic gradient ascent.

3.3.2 Re-configuring Data via an Augmenting Variational Auto-Encoder

To augment the above VAE framework [26, 30] to account for the impositions in Conditions 1 and 2, we ideally want to configure the parameterization of the above generative process to guarantee that the marginal posterior $q(\mathbf{z}) = \int_{\mathbf{x}} q(\mathbf{z}|\mathbf{x})p(\mathbf{x})d\mathbf{x}$ will manifest itself in the form of a mixture of Gaussians with the desired concentration and population densities as stated in Condition 1.

However, it is often difficult to make such an imposition directly given that we typically have no prior knowledge of $p(\mathbf{x})$. We instead impose the desired structure on the latent prior $p(\mathbf{z})$ and then penalize the divergence between $q_\phi(\mathbf{z})$ and $p(\mathbf{z})$ while optimizing for the above ELBO in Eq. (8). That is, we parameterize $p(\mathbf{z}) = \pi_1 \mathbf{N}(\mathbf{z}; \mathbf{c}_1, \gamma_1^2 \mathbf{\Theta}^{-1}) + \dots + \pi_b \mathbf{N}(\mathbf{z}; \mathbf{c}_b, \gamma_b^2 \mathbf{\Theta}^{-1})$ where $\pi_i \propto 2^{i/2}$ (see Condition 2), which encodes the desired clustering structure. This is then reflected on the marginal posterior $q(\mathbf{z})$ via augmenting the above ELBO as,

$$\mathbf{L}_\alpha(\mathbf{x}; \theta, \phi) \triangleq \mathbb{E}_{\mathbf{z} \sim q_\phi} \left[\log p_\theta(\mathbf{x}|\mathbf{z}) \right] - \mathbb{KL} \left(q_\phi(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}) \right) - \alpha \mathbb{KL} \left(q(\mathbf{z}) || p(\mathbf{z}) \right), \quad (9)$$

where the penalty term $\alpha \mathbb{KL}(q(\mathbf{z}) || p(\mathbf{z}))$ serves as an incentive to encourage $q(\mathbf{z})$ to assume the same clustering structure as $p(\mathbf{z})$. The parameter α can be manually set to adjust the strength of the incentive. To encourage separation among learned clusters (see Condition 3), we also add an extra penalty term to the above augmented ELBO,

$$\mathbf{L}_{\alpha, \beta}(\mathbf{x}; \theta, \phi) \triangleq \mathbf{L}_\alpha(\mathbf{x}; \theta, \phi) + \beta \sum_{i \neq j} \mathbb{KL} \left(\mathbf{N}(\mathbf{z}; \mathbf{c}_i, \gamma_i^2 \mathbf{\Theta}^{-1} \mathbf{I}) || \mathbf{N}(\mathbf{z}; \mathbf{c}_j, \gamma_j^2 \mathbf{\Theta}^{-1} \mathbf{I}) \right). \quad (10)$$

Once these clusters are learned, we can use the resulting encoding network $q_\phi(\mathbf{z}|\mathbf{x})$ to transform each training input \mathbf{x} into its latent projection and subsequently train an SSGP on the latent space of \mathbf{z} (instead of training it on the original data space). Our previous analysis can then be applied on \mathbf{z} to give the desired sample complexity. The empirical efficiency of the proposed method is demonstrated in Section 4 below. Note that the cost of training the embedding is linear in the number of data points and therefore does not noticeably affect our overall running time.

4 Experiments

Datasets. This section presents our empirical studies on two real datasets: (a) the ABALONE dataset [42] with 3000 data points which was used to train a model that predicts the age of abalone (number of rings on its shell) from physical measurements such as length, diameter, height, whole weight, shucked weight, viscera weight and shell weight; and (b) the GAS SENSOR dataset with 4 million data points [5, 6] which was used to train a model that predicts the CO concentration (ppm) from measurements of humidity, temperature, flow rate, heater voltage and the resistant measures of 14 gas sensors.

In both settings, we compare our revised SSGP method with the traditional SSGP on both datasets to demonstrate its sample efficiency. In particular, our SSGP method is applied on the embedded space of data which was generated and configured using the auto-encoding method in Section 3.3.2 to approximately meet the aforementioned Conditions 1-3.

The detailed parameterization of our entire algorithm⁷ is provided in Appendix C. The traditional SSGP method on the other hand was applied directly to the data space. The prediction root-mean-square-error (RMSE) achieved by each method is reported at different sample complexities in Figure 1 below. All reported performances were averaged over 5 independent runs on a computing server with a Tesla K40 GPU with 12GB RAM.

Results and Discussions. It can be observed from the results that at all levels of sample complexity, the revised SSGP achieves substantially better performance than its vanilla SSGP counterpart. This

⁷Our experimental code is released at https://github.com/hqminh/gp_sketch_nips.

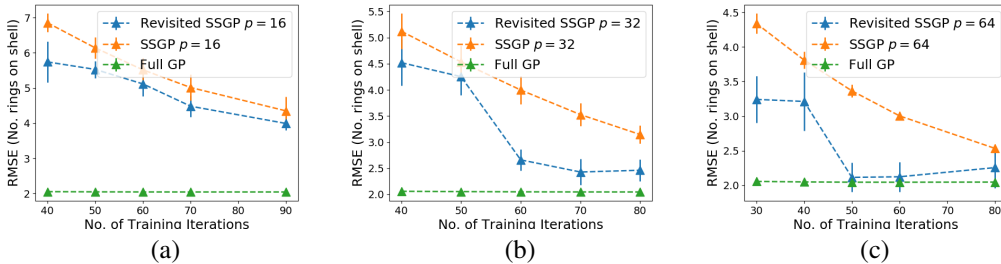


Figure 1: Graphs of performance comparisons between our revised SSGP and the traditional SSGP on the ABALONE dataset [42] at varying sample complexities (see Theorem 2) $p = 16, 32$ and 64 .

is expected since our revised SSGP is guaranteed to require many fewer samples than the vanilla SSGP when the data is reconfigured to exhibit a certain clustering structure (see Conditions 1-3 and Theorem 2). As such, when both are set to operate at the same level of sample complexity, one would expect the revised SSGP to achieve better performance since SSGP generally performs better when its sample complexity is set closer to the required threshold. On the larger GAS SENSOR dataset (which contains approximately 4M data points), we also observe the same phenomenon from the performance comparison graph as shown in Figure 2a below: A vanilla SSGP needs to increase its number of samples to marginally improve its predictive performance while our revised SSGP is able to outperform the former with the least number of samples ($p = 16$).

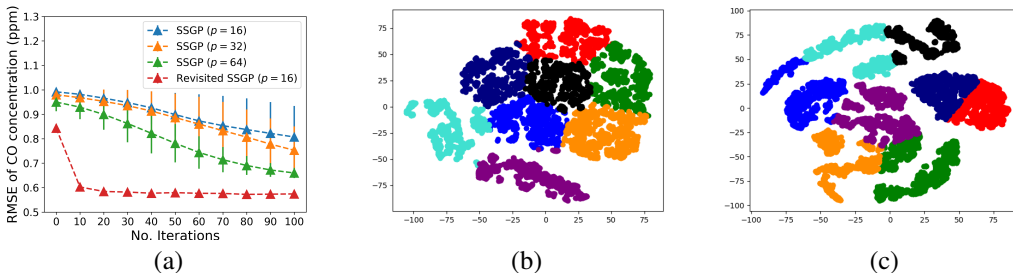


Figure 2: Graphs of (a) performance comparison between our revisited SSGP’s (with sample complexity $p = 16$) and the vanilla SSGP’s (with sample complexity $p = 16, 32, 64$) on the GAS SENSOR dataset [4]; and visualizations of (b) original and (c) reconfigured data distributions of GAS SENSOR data on a 2-dimensional latent space generated by our auto-encoding algorithm in Section 3.3.

Furthermore, a closer look into the data distribution (visualized on a 2D space in Fig. 2b) and the data reconfigured data distribution (visualized on a 2D space in Fig. 2c) also corroborates our hypothesis earlier that a well-separated data partition with high in-cluster concentration (in the form of a mixture of clusters – see Condition 1) can be found (by our embedding algorithm in Section 3.3) to reconfigure our data distribution to (approximately) meet the necessary technical conditions that enable our sample-complexity enhancement analysis (see Section 3.1). Due to limited space, interested readers are referred to Appendix D for more detailed empirical studies and demonstrations.

5 Conclusion

We present a new method and analysis for approximating Gaussian processes. We obtain provable guarantees for both training and inference, which are the first to hold simultaneously over the entire space of kernel parameters. Our results complement existing work in kernel approximation that often assumes knowledge of its defining parameters. Our results also reveal important (practical) insights that allow us to develop an algorithmic handle on the tradeoff between approximation quality and sample complexity, which is achieved via finding an embedding that disentangles the latent coordinates of data. Our empirical results show for many datasets, such a disentangled embedding space can be found, which leads to a significantly reduced sample complexity of SSGP.

6 Statement of Broader Impact

Our work focuses on approximating Gaussian processes using a mixture of practical methods and theoretical analysis to reconfigure data in ways that reduce their approximation complexity. As such, it could have significant broader impact by allowing users to more accurately solve practical problems such as the ones discussed in our introduction, while still providing concrete theoretical guarantees. While applications of our work to real data could result in ethical considerations, this is an indirect (and unpredictable) side-effect of our work. Our experimental work uses publicly available datasets to evaluate the performance of our algorithms; no ethical considerations are raised.

7 Acknowledgement

T. N. Hoang is supported by the MIT-IBM Watson AI Lab, IBM Research. Q. M. Hoang is supported by the Gordon and Betty Moore Foundation’s Data-Driven Discovery Initiative through Grant GBMF4554, by the US National Science Foundation (DBI-1937540), by the US National Institutes of Health (R01GM122935), and by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program. D. Woodruff is supported by National Institute of Health grant 5R01 HG 10798-2, Office of Naval Research grant N00014-18-1-2562, and a Simons Investigator Award.

References

- [1] Rakshit Allamraju and Girish Chowdhary. Communication efficient decentralized gaussian process fusion for multi-uas path planning. In *American Control Conference*, pages 4442–4447, 05 2017.
- [2] T. J. Ansell et al. Daily mean sea level pressure reconstructions for the European-North Atlantic region for the period 1850-2003. *J. Climate*, 19(12):2717–2742, 2006.
- [3] H. Avron, M. Kapralov, C. Musco, C. Musco, A. Velingker, and A. Zandieh. Random fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In *Proc. ICML*, pages 253–262, 2017.
- [4] Javier Burgues. Gas Sensor Array Temperature Modulation Dataset, howpublished= <https://archive.ics.uci.edu/ml/machine-learning-databases/00487/>.
- [5] Javier Burgués, Juan Manuel Jiménez-Soto, and Santiago Marco. Estimation of the limit of detection in semiconductor gas sensors through linearized calibration models. *Analytica chimica acta*, 1013:13–25, 2018.
- [6] Javier Burgués and Santiago Marco. Multivariate estimation of the limit of detection by orthogonal partial least squares in temperature-modulated mox sensors. *Analytica chimica acta*, 1019:49–64, 2018.
- [7] N. Cao, K. H. Low, and J. M. Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proc. AAMAS*, pages 7–14, 2013.
- [8] J. Chen, N. Cao, K. H. Low, R. Ouyang, C. K.-Y. Tan, and P. Jaillet. Parallel Gaussian process regression with low-rank covariance matrix approximations. In *Proc. UAI*, pages 152–161, 2013.
- [9] J. Chen, K. H. Low, C. K.-Y. Tan, A. Oran, P. Jaillet, J. M. Dolan, and G. S. Sukhatme. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In *Proc. UAI*, pages 163–173, 2012.
- [10] Jie Chen, Kian Hsiang Low, and Colin Tan. Gaussian process-based decentralized data fusion and active sensing for mobility-on-demand system. *Robotics: Science and System*, 06 2013.
- [11] H. Chernoff. A measure of asymptotic efficiency for tests of hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–509, 1952.

- [12] J. M. Dolan, G. Podnar, S. Stancliff, K. H. Low, A. Elfes, J. Higinbotham, J. C. Hosler, T. A. Moisan, and J. Moisan. Cooperative aquatic sensing using the telesupervised adaptive ocean sensor fleet. In *Proc. SPIE Conference on Remote Sensing of the Ocean, Sea Ice, and Large Water Regions*, volume 7473, 2009.
- [13] Y. Gal, M. van der Wilk, and C. Rasmussen. Distributed variational inference in sparse Gaussian process regression and latent variable models. In *Proc. NIPS*, 2014.
- [14] Yarin Gal and Richard Turner. Improving the gaussian process sparse spectrum approximation by representing uncertainty in frequency inputs. 2015.
- [15] J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Proc. UAI*, pages 282–290, 2013.
- [16] M. Hoang and C. Kingsford. Optimizing dynamic structures with bayesian generative search. In *International Conference on Machine Learning*, 2020.
- [17] Q. M. Hoang, T. N. Hoang, and K. H. Low. A generalized stochastic variational Bayesian hyperparameter learning framework for sparse spectrum Gaussian process regression. In *Proc. AAAI*, pages 2007–2014, 2017.
- [18] Q. M. Hoang, T. N. Hoang, K. H. Low, and C. Kingsford. Collective model fusion for multiple black-box experts. In *Proc. ICML*, 2019.
- [19] T. N. Hoang, Q. M. Hoang, and K. H. Low. A unifying framework of anytime sparse Gaussian process regression models with stochastic variational inference for big data. In *Proc. ICML*, pages 569–578, 2015.
- [20] T. N. Hoang, Q. M. Hoang, and K. H. Low. A distributed variational inference framework for unifying parallel sparse Gaussian process regression models. In *Proc. ICML*, pages 382–391, 2016.
- [21] T. N. Hoang, Q. M. Hoang, K. H. Low, and J. P. How. Collective online learning of Gaussian processes in massive multi-agent systems. In *Proc. AAAI*, 2019.
- [22] T. N. Hoang, Q. M. Hoang, O. Ruofei, and K. H. Low. Decentralized high-dimensional bayesian optimization with factor graphs. In *Proc. AAAI*, 2018.
- [23] T. N. Hoang, K. H. Low, P. Jaillet, and M. Kankanhalli. Nonmyopic ϵ -Bayes-optimal active learning of Gaussian processes. In *Proc. ICML*, pages 739–747, 2014.
- [24] T. N. Hoang, K. H. Low, P. Jaillet, and M. S. Kankanhalli. Active learning is planning: Nonmyopic ϵ -Bayes-optimal active learning of Gaussian processes. In *Proc. ECML-PKDD Nectar Track*, pages 494–498, 2014.
- [25] W. Hoeffding. Probability inequalities for the sum of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- [26] D. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *Proc. ICLR*, 2013.
- [27] A. Krause and C. Guestrin. Nonmyopic active learning of Gaussian processes: An exploration-exploitation approach. In *Proc. ICML*, pages 449–456, 2007.
- [28] M. Lázaro-Gredilla, J. Quiñonero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal. Sparse spectrum Gaussian process regression. *Journal of Machine Learning Research*, pages 1865–1881, 2010.
- [29] K. H. Low, J. Yu, J. Chen, and P. Jaillet. Parallel Gaussian process regression for big data: Low-rank representation meets Markov approximation. In *Proc. AAAI*, pages 2821–2827, 2015.
- [30] Emile Mathieu, Tom Rainforth, Siddharth Narayanaswamy, and Yee Whye Teh. Disentangling disentanglement in variational autoencoders. In *ICML*, 2019.
- [31] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT press, 2018.

- [32] C. Musco and C. Musco. Recursive sampling for the nystrom method. In *Proc. NIPS*, 2016.
- [33] J. Quiñonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.
- [34] J. Quiñonero-Candela, C. E. Rasmussen, and C. K. I. Williams. Approximation methods for gaussian process regression. *Large-Scale Kernel Machines*, pages 203–223, 2007.
- [35] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Proc. NIPS*, 2007.
- [36] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [37] M. Seeger, C. K. I. Williams, and N. D. Lawrence. Fast forward selection to speed up sparse Gaussian process regression. In *Proc. AISTATS*, 2003.
- [38] J. Snoek, L. Hugo, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Proc. NIPS*, pages 2960–2968, 2012.
- [39] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. ICML*, pages 1015–1022, 2010.
- [40] M. K. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Proc. AISTATS*, 2009.
- [41] L. J. P. van der Maaten and G. E. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [42] Sam Waugh. Abalone Dataset, howpublished= <https://archive.ics.uci.edu/ml/machine-learning-databases/abalone/>.
- [43] N. Xu, K. H. Low, J. Chen, K. K. Lim, and E. B. Özgül. GP-Localize: Persistent mobile robot localization using online sparse Gaussian process observation model. In *Proc. AAAI*, pages 2585–2592, 2014.
- [44] Y. Zhang, T. N. Hoang, K. H. Low, and M. Kankanhalli. Near-optimal active learning of multi-output Gaussian processes. In *Proc. AAAI*, pages 2351–2357, 2016.
- [45] Y. Zhang, T. N. Hoang, K. H. Low, and M. Kankanhalli. Information-based multi-fidelity bayesian optimization. In *NIPS Workshop BayesOpt*, 2017.

A Intermediate Results for Theorem 2

Let $\Delta(\mathbf{x}_u, \mathbf{x}_v) \triangleq |\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) - \mathbf{K}'(\mathbf{x}_u, \mathbf{x}_v)|$ where $\mathbf{K}'(\mathbf{x}_u, \mathbf{x}_v) = (1/p) \sum_{i=1}^p \mathbf{K}_{\epsilon_i}(\mathbf{x}_u, \mathbf{x}_v)$, and where $\epsilon_i \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ as defined in Lemma 1 above. We will first measure the approximation loss across different value-bands of $\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v)$, thereby deriving tight sample bounds for each band. Combining these with the union bound allows us to establish a much cheaper overall sample complexity as compared to the naïve $\mathbf{O}(n^2 \log n)$ bound.

Lemma 2. *Suppose the data distribution follows Conditions 1-3 above. Let $\mathbf{c}(\mathbf{x}_u)$ denote the cluster index of each data point \mathbf{x}_u . Let $\mathcal{C} \triangleq \{u, v \mid \mathbf{c}(\mathbf{x}_u) = \mathbf{c}(\mathbf{x}_v)\}$ and $\mathcal{C}' \triangleq \{u, v \mid \mathbf{c}(\mathbf{x}_u) \neq \mathbf{c}(\mathbf{x}_v)\}$ denote the sets of in-cluster and out-cluster kernel entries, respectively, where $|\mathcal{C}| \simeq \frac{n^2}{4}$ and $|\mathcal{C}'| \simeq \frac{3n^2}{4}$.*

Proof. By Condition 2, since n data points are scattered across b clusters and each cluster i has $2^{i/2}$ points, it follows that:

$$\begin{aligned} n &= \sum_{i=1}^b 2^{\frac{i}{2}} = \frac{\sqrt{2^{b+1}} - \sqrt{2}}{\sqrt{2} - 1} \\ \Rightarrow |\mathcal{C}| &= \sum_{i=1}^b 2^i = 2^{b+1} - 1 = \left(n(\sqrt{2} - 1) + \sqrt{2}\right)^2 - 1 \simeq \frac{n^2}{4} \\ \Rightarrow |\mathcal{C}'| &= n^2 - |\mathcal{C}| \simeq \frac{3n^2}{4}. \end{aligned} \quad (11)$$

This also implies that $b = \mathbf{O}(\log n)$, which is consistent with Condition 1 above. \square

Lemma 3. *Let $\mathcal{C}_i = \{(u, v) \in \mathcal{C} \mid \mathbf{c}(\mathbf{x}_u) = \mathbf{c}(\mathbf{x}_v) = i\}$ for $i \in [1 \dots b]$. Then with probability at least $1 - \delta$, for $\delta \geq \mathbf{O}(\exp(\log n - \sqrt{d}))$, the following holds for all i and $(u, v) \in \mathcal{C}_i$ for which $u \neq v$:*

$$\left(1 - \frac{1}{2^{a+i-1}}\right)^{\frac{1}{4}} \leq \mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) < \left(1 - \frac{1}{2^{a+i}}\right)^{\frac{1}{4}} \text{ where } a = \frac{1}{\log 2} \log \left(\frac{n^4}{n^4 - \lambda^4}\right) \quad (12)$$

Proof. If \mathbf{x}_u and \mathbf{x}_v are both generated from component i of the data distribution as defined in Condition 1, it follows that $\Theta^{-1/2}(\mathbf{x}_u - \mathbf{x}_v) \sim \mathbf{N}(\mathbf{c}_i, \gamma_i^2 \mathbf{I})$. Therefore, by standard chi-squared tail bounds, with probability at least $1 - 2e^{-t}$, we have:

$$\left\| \Theta^{-1/2}(\mathbf{x}_u - \mathbf{x}_v) \right\|_2^2 = \gamma_i^2 d \pm \mathbf{O}\left(\gamma_i^2 \sqrt{dt}\right), \quad (13)$$

where d is the data dimension. Using this, we can then figure out a setting for γ_i^2 such that $\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v)$ follows the above condition in Eq. 12. In particular, set

$$\mathcal{L}(i) = \log \left(\frac{2^{a+i}}{2^{a+i} - 1}\right) \quad \text{and} \quad \mathcal{U}(i) = \log \left(\frac{2^{a+i-1}}{2^{a+i-1} - 1}\right). \quad (14)$$

We can then choose:

$$\gamma_i^2 = \frac{1}{4d} \left(\mathcal{U}(i) + \mathcal{L}(i)\right) \quad \text{and} \quad t = \sqrt{d} \left(\frac{\mathcal{U}(i) - \mathcal{L}(i)}{\mathcal{U}(i) + \mathcal{L}(i)}\right) \simeq \mathbf{O}(\sqrt{d}), \quad (15)$$

so that by plugging these choices in Eq. 13 above, we have with probability at least $1 - 2e^{-t}$:

$$\begin{aligned} \left\| \Theta^{-1/2}(\mathbf{x}_u - \mathbf{x}_v) \right\|_2^2 &\in \left[\frac{1}{2} \log \left(\frac{2^{a+i}}{2^{a+i} - 1}\right), \frac{1}{2} \log \left(\frac{2^{a+i-1}}{2^{a+i-1} - 1}\right) \right] \\ \Rightarrow \mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) &\in \left[\left(1 - \frac{1}{2^{a+i-1}}\right)^{\frac{1}{4}}, \left(1 - \frac{1}{2^{a+i}}\right)^{\frac{1}{4}} \right]. \end{aligned} \quad (16)$$

Now, note that for any δ for which $\delta \geq \mathbf{O}\left(4^b e^{-\sqrt{d}}\right) \geq \mathbf{O}\left(4^i e^{-\sqrt{d}}\right) \forall i \leq b$, we have $\delta/4^i \geq 2e^{-t}$ since $t \simeq \mathbf{O}(\sqrt{d})$. This also means $\delta \geq \mathbf{O}(\exp(\log n - \sqrt{d}))$ since $b = \mathbf{O}(\log n)$.

That is, Eq. (16) and hence, Eq. (12), hold with probability at least $1 - 2e^{-t} \geq 1 - \delta/4^i$ for each entry in \mathcal{C}_i . For each cluster i , even though there are up to 2^i kernel entries, by the triangle inequality it is easy to see that we only need to apply a union bound over at most $2^{i/2}$ (carefully selected) entries (excluding the entries on the diagonal) to meet Eq. (12) with probability at least $1 - 2^i(\delta/4^i) = 1 - \delta/2^i$.

Subsequently, applying a union bound over all clusters gives us that with probability at least $1 - \delta \sum_{i=1}^b 1/2^i \geq 1 - \delta$, all kernel entries within the i -th cluster satisfy Eq. (12) simultaneously for $1 \leq i \leq b$. \square

Lemma 4. For all $(u, v) \in \mathcal{C}' \triangleq \{u, v \mid \mathbf{c}(\mathbf{x}_u) \neq \mathbf{c}(\mathbf{x}_v)\}$, we have $\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) < \left(1 - \frac{1}{2^a}\right)^{\frac{1}{4}}$

where $a = \frac{1}{\log 2} \log \left(\frac{n^4}{n^4 - \lambda^4}\right)$ as defined in Lemma 3 above.

Proof. For any (u, v) for which $\mathbf{c}(\mathbf{x}_u) = i$ and $\mathbf{c}(\mathbf{x}_v) = j$ and $i \neq j$, we have:

$$\begin{aligned} \left\| \Theta^{-1/2}(\mathbf{x}_u - \mathbf{x}_v) \right\|_2^2 &\geq \left\| \Theta^{-1/2}(\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 - \left\| \Theta^{-1/2}(\mathbf{x}_u - \mathbf{c}_i) \right\|_2^2 - \left\| \Theta^{-1/2}(\mathbf{x}_v - \mathbf{c}_j) \right\|_2^2 \\ &\geq \left\| \Theta^{-1/2}(\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 - \frac{1}{2} \log \left(\frac{2^{a+i}}{2^{a+i} - 1} \right) - \frac{1}{2} \log \left(\frac{2^{a+j}}{2^{a+j} - 1} \right) \\ &\geq \left\| \Theta^{-1/2}(\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 - \log \left(\frac{2^a}{2^a - 1} \right) \\ \Rightarrow \mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) &= \exp \left(-\frac{1}{2} \left\| \Theta^{-1/2}(\mathbf{x}_u - \mathbf{x}_v) \right\|_2^2 \right) \\ &\leq \exp \left(-\frac{1}{2} \cdot \left\| \Theta^{-1/2}(\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 + \frac{1}{2} \log \left(\frac{2^a}{2^a - 1} \right) \right) < \left(1 - \frac{1}{2^a}\right)^{\frac{1}{4}} \end{aligned}$$

since for all (i, j) , by Condition 3:

$$\left\| \Theta^{-1/2}(\mathbf{c}_i - \mathbf{c}_j) \right\|_2^2 > \frac{3}{2} \log \left(\frac{2^a}{2^a - 1} \right). \quad (17)$$

This completes our proof for the stated result of Lemma 4. \square

Corollary 1. With probability at least $1 - \delta$, there are exactly n entries that are greater than $1 - 2^{-(a+b)}$ where $a = \frac{1}{\log 2} \log \left(\frac{n^4}{n^4 - \lambda^4}\right)$. These are the diagonal entries $\mathbf{K}(\mathbf{x}_u, \mathbf{x}_u)$ with $1 \leq u \leq n$.

Proof. Lemma 3 asserts that with probability $1 - \delta$, all kernel entries $\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v)$, where $\mathbf{c}(\mathbf{x}_u) = \mathbf{c}(\mathbf{x}_v) = i$, belong to their respective band $\kappa_i = \{(u, v) \mid 1 - 1/2^{a+i-1} \leq \mathbf{K}^4(\mathbf{x}_u, \mathbf{x}_v) \leq 1 - 1/2^{a+i}\}$. When this happens, all in-cluster entries (except the diagonal entries) will have values between $1 - 1/2^a$ and $1 - 1/2^{a+b}$ (since there are b bands) and as such, off-cluster entries will either be smaller than $1 - 1/2^a$ or larger than $1 - 1/2^{a+b}$. But then Lemma 4 further guarantees that all off-cluster entries are smaller than $1 - 1/2^a$, following Condition 3. Thus, it follows that the only entries that are larger than $1 - 1/2^{a+b}$ are the diagonal items and there are exactly n of them. \square

Lemma 5. Let $\kappa_i = \{(u, v) \mid 1 - 1/2^{a+i-1} \leq \mathbf{K}^4(\mathbf{x}_u, \mathbf{x}_v) < 1 - 1/2^{a+i}\}$. It follows that for each $i \in [1 \dots b]$, with probability at least $1 - \delta/b$:

$$\sum_{(u,v) \in \mathcal{G}_i} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{\lambda^2}{b}, \quad (18)$$

if the kernel approximation $\mathbf{K}'(\mathbf{x}_u, \mathbf{x}_v) \triangleq \frac{1}{p} \sum_{t=1}^p \mathbf{K}_{\epsilon_t}(\mathbf{x}_u, \mathbf{x}_v)$ is formed using at least $p = \frac{b|\kappa_i|}{\lambda^2 \cdot 2^{a+i}} \log \left(\frac{b|\kappa_i|}{\delta}\right) = \mathbf{O} \left(\frac{\log^2 n}{\lambda^2} \log \left(\frac{\log n}{\delta}\right)\right)$ samples.

Proof. For all (u, v) , we have $\mathbf{K}_{\epsilon_t}(\mathbf{x}_u, \mathbf{x}_v) = \cos(\epsilon_t^\top \Theta^{-1/2} (\mathbf{x}_u - \mathbf{x}_v))$ where $\epsilon_t \sim \mathbf{N}(0, \mathbf{I})$ and,

$$\mathbf{K}_{\epsilon_t}(\mathbf{x}_u, \mathbf{x}_v) = \cos\left(\sum_{\ell=1}^d \epsilon_t^\ell \cdot \left(\frac{\mathbf{x}_u^\ell - \mathbf{x}_v^\ell}{\theta_\ell}\right)\right) \triangleq \cos(\mathbf{z}_{uv}^t). \quad (19)$$

Since $\epsilon_t^\ell \sim \mathbf{N}(0, 1)$, \mathbf{z}_{uv}^t is then a weighted sum of Gaussian random variables and $\mathbf{z}_{uv}^t \sim \mathbf{N}(0, \Sigma_{uv}^t)$, where $\Sigma_{uv}^t \triangleq (\mathbf{x}_u - \mathbf{x}_v)^\top \Theta^{-1} (\mathbf{x}_u - \mathbf{x}_v)$, which in turn implies:

$$\begin{aligned} \mathbb{E}[\cos(\mathbf{z}_{uv}^t)] &= \exp(-0.5 \Sigma_{uv}^t) = \mathbf{K}(\mathbf{x}_u, \mathbf{x}_v), \\ \mathbb{V}[\cos(\mathbf{z}_{uv}^t)] &= \frac{1}{2} [1 - \mathbb{E}[\cos(\mathbf{z}_{uv}^t)]]^2 = \frac{1}{2} (1 - \mathbf{K}^2(\mathbf{x}_u, \mathbf{x}_v))^2 \leq 2 \times \frac{1}{2^{a+i}}, \end{aligned} \quad (20)$$

where the last inequality follows from the choice of $(u, v) \in \kappa_i$ and the definition of the κ_i above. Next, applying the Chernoff-Hoeffding inequality and union bounding over the κ_i , we have:

$$\begin{aligned} \Pr\left(\forall (u, v) \in \kappa_i : \Delta(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{\epsilon}{p}\right) &\geq 1 - 2|\kappa_i| \exp\left(-\frac{\epsilon^2}{4 \sum_{t=1}^p \mathbb{V}[\cos(\mathbf{z}_{uv}^t)]}\right) \\ \Rightarrow \Pr\left(\sum_{(u,v) \in \kappa_i} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{|\kappa_i| \epsilon^2}{p^2}\right) &\geq 1 - 2|\kappa_i| \exp\left(-\frac{\epsilon^2 \cdot 2^{a+i}}{8p}\right). \end{aligned} \quad (21)$$

Thus, setting $\epsilon^2 = \frac{\lambda^2 p^2}{4b|\kappa_i|}$ and $p \geq \frac{32b|\kappa_i|}{\lambda^2 \cdot 2^{a+i}} \log\left(\frac{2b|\kappa_i|}{\delta}\right)$ yields:

$$\Pr\left(\sum_{(u,v) \in \mathcal{G}_i} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{\lambda^2}{4b}\right) \geq 1 - 2|\kappa_i| \exp\left(-\frac{\lambda^2 p \cdot 2^{a+i}}{32b|\kappa_i|}\right) \geq 1 - \frac{\delta}{b}. \quad (22)$$

where the last inequality follows from the above choice of p . Since $|\kappa_i| = 2^i$ by Condition 2, we further have $p \geq \frac{32b}{\lambda^2 \cdot 2^a} \log\left(\frac{b \cdot 2^{b+1}}{\delta}\right) = \mathbf{O}\left(\frac{\log^2 n}{\lambda^2} \log\left(\frac{\log n}{\delta}\right)\right)$. \square

Lemma 5 thus establishes a very strong sample complexity of $\mathbf{O}(\log^2 n \log \log n)$ for approximating all kernel entries within a narrow band of values, which is significantly cheaper than the sample complexity of $\mathbf{O}(n^2 \log n)$ we would get if we were to ignore the distribution of kernel values in different bands. This is made clear in Corollary 2 below, which combines Lemmas 3, 4 and 5 to establish an overall sample complexity resulting in only a small approximation loss accumulated over all bands.

Corollary 2. *If a kernel approximation \mathbf{K}' of \mathbf{K} is formed such that $\mathbf{K}'(\mathbf{x}_u, \mathbf{x}_v) \triangleq \frac{1}{p} \sum_{t=1}^p \mathbf{K}_{\epsilon_t}(\mathbf{x}_u, \mathbf{x}_v)$ for all in-cluster entries $(u, v) \in \mathcal{C}$ using $p = \mathbf{O}((\log^2 n / \lambda^2) \log(\log n / \delta))$ samples and $\mathbf{K}'(\mathbf{x}_{u'}, \mathbf{x}_{v'}) \triangleq 0$ for all off-cluster entries $(u', v') \in \mathcal{C}'$, then,*

$$\|\mathbf{K} - \mathbf{K}'\|_2^2 \leq \|\mathbf{K} - \mathbf{K}'\|_F^2 \leq \lambda^2,$$

with probability at least $1 - 2\delta$ with $\delta \geq \mathbf{O}(\exp(\log n - \sqrt{d}))$. This immediately guarantees that \mathbf{K}' is spectrally close to \mathbf{K} using the notion of λ -closeness (see Definition 1).

Proof. By Lemma 3, with probability $1 - \delta$, $|\kappa_i| = |\mathcal{C}_i|$ simultaneously for all i . Thus, applying a union bound over this event and the results obtained in Lemma 5 for all clusters, we have the following bound on the total approximation loss over in-cluster entries in \mathcal{C} with probability $1 - 2\delta$:

$$\sum_{(u,v) \in \mathcal{C}} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{\lambda^2}{4}. \quad (23)$$

Furthermore, by Lemma 4, we also have the following bound on the total approximation loss over off-cluster entries in \mathcal{C}' (which were approximated uniformly by zero):

$$\sum_{(u,v) \in \mathcal{C}'} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{3n^2}{4} (\mathbf{K}(\mathbf{x}_u, \mathbf{x}_v) - 0)^2 \leq \frac{3n^2}{4} \sqrt{1 - \frac{1}{2^a}} = \frac{3\lambda^2}{4}, \quad (24)$$

when the last inequality is due to the facts (established in Lemma 4) that $\mathbf{K}^4(\mathbf{x}_u, \mathbf{x}_v) \leq 1 - 1/2^a$ and that $a = \frac{1}{\log 2} \log \left(\frac{n^4}{n^4 - \lambda^4} \right)$. Finally, combining these yields:

$$\|\mathbf{K} - \mathbf{K}'\|_2^2 \leq \|\mathbf{K} - \mathbf{K}'\|_F^2 = \sum_{(u,v) \in \mathcal{C}} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) + \sum_{(u,v) \in \mathcal{C}'} \Delta^2(\mathbf{x}_u, \mathbf{x}_v) \leq \frac{1}{4}\lambda^2 + \frac{3}{4}\lambda^2 = \lambda^2 \quad (25)$$

□

B Intermediate Results for Theorem 3

Lemma 6. *Let \mathbf{K} and \mathbf{K}' be positive semidefinite matrices in $\mathbb{R}^{n \times n}$ such that $-\lambda \mathbf{I} \preceq \mathbf{K} - \mathbf{K}' \preceq \lambda \mathbf{I}$, $\mathbf{Q} \triangleq \mathbf{K} + \sigma^2 \mathbf{I}$ and $\mathbf{Q}' \triangleq \mathbf{K}' + \sigma^2 \mathbf{I}$ for some $\lambda, \sigma > 0$, then:*

$$\|\mathbf{Q}'^{-1}\|_2 = \left(1 \pm \frac{\lambda}{\sigma^2}\right) \|\mathbf{Q}^{-1}\|_2. \quad (26)$$

Proof. By definition of the spectral norm, we have $\forall \mathbf{x} \in \mathbb{R}^n$:

$$\mathbf{K} - \mathbf{K}' \preceq \lambda \mathbf{I}, \quad (27)$$

which implies

$$\begin{aligned} \mathbf{Q} &\preceq \mathbf{K}' + (\sigma^2 + \lambda)\mathbf{I} \\ &\preceq (\sigma^2 + \lambda)\mathbf{K}' + (\sigma^2 + \lambda)\mathbf{I} \\ &= \left(1 + \frac{\lambda}{\sigma^2}\right) \mathbf{Q}'. \end{aligned} \quad (28)$$

where \preceq and \succeq denote the Loewner inequality operators. Likewise, by symmetry, we also have:

$$\mathbf{Q}' \preceq \left(1 + \frac{\lambda}{\sigma^2}\right) \mathbf{Q}. \quad (29)$$

Let $\mathbf{A} \triangleq (1 + \lambda/\sigma^2)\mathbf{Q}'$ and $\mathbf{B} \triangleq \mathbf{Q}$. Since \mathbf{A} and \mathbf{B} are symmetric and positive semidefinite, there exist \mathbf{U}, \mathbf{V} with orthogonal rows and columns and diagonal matrices Σ, Σ' for which $\mathbf{A} = \mathbf{U}\Sigma\mathbf{U}^\top$ and $\mathbf{B} = \mathbf{V}\Sigma'\mathbf{V}^\top$. We further let $\mathbf{A}^{-1/2} \triangleq \mathbf{U}\Sigma^{-1/2}$ and $\mathbf{B}^{-1/2} \triangleq \mathbf{V}\Sigma'^{-1/2}$.

Then, we can rewrite Eq. (28) as:

$$\begin{aligned} \mathbf{A} - \mathbf{B} &\succeq 0 \\ \Rightarrow \mathbf{B}^{-1/2}(\mathbf{A} - \mathbf{B})\mathbf{B}^{-1/2} &\succeq 0 \\ \Rightarrow \mathbf{B}^{-1/2}\mathbf{A}\mathbf{B}^{-1/2} - \mathbf{I} &\succeq 0 \\ \Rightarrow \mathbf{A}^{-1/2}\mathbf{B}^{1/2}(\mathbf{B}^{-1/2}\mathbf{A}\mathbf{B}^{-1/2})\mathbf{B}^{-1/2}\mathbf{A}^{1/2} &\succeq \mathbf{A}^{-1/2}\mathbf{B}^{1/2}\mathbf{B}^{-1/2}\mathbf{A}^{1/2} \\ &\Rightarrow \mathbf{A}^{1/2}\mathbf{B}^{-1}\mathbf{A}^{1/2} &\succeq \mathbf{I} \\ \Rightarrow \mathbf{A}^{-1/2}(\mathbf{A}^{1/2}\mathbf{B}^{-1}\mathbf{A}^{1/2})\mathbf{A}^{-1/2} &\succeq \mathbf{A}^{-1/2}\mathbf{A}^{-1/2} \\ &\Rightarrow \mathbf{B}^{-1} &\succeq \mathbf{A}^{-1} \\ &\Rightarrow \mathbf{Q}^{-1} &\succeq \frac{\sigma^2}{\sigma^2 + \lambda}\mathbf{Q}'^{-1} \\ &\Rightarrow \left(1 + \frac{\lambda}{\sigma^2}\right) \mathbf{Q}^{-1} &\succeq \mathbf{Q}'^{-1}. \end{aligned} \quad (30)$$

Again, by symmetry, we can rewrite Eq. 29 as:

$$\begin{aligned} \mathbf{Q}'^{-1} &\succeq \frac{\sigma^2}{\sigma^2 + \lambda}\mathbf{Q}^{-1} \succeq \left(1 - \frac{\lambda}{\sigma^2 + \lambda}\right) \mathbf{Q}^{-1} \\ &\succeq \left(1 - \frac{\lambda}{\sigma^2}\right) \mathbf{Q}^{-1}. \end{aligned} \quad (31)$$

Therefore, we have $\|\mathbf{Q}'^{-1}\|_2 = (1 \pm \lambda/\sigma^2)\|\mathbf{Q}^{-1}\|_2$. □

Let $g(\mathbf{x}_*)$ and $g'(\mathbf{x}_*)$ respectively denote the predictive distributions of full GP and the approximated GP pertaining to an arbitrary test input \mathbf{x}_* . We then state the following lemmas:

Lemma 7. *Let \mathbf{K}' denote an approximation that is λ -close to the original kernel \mathbf{K} . The induced predictive mean of \mathbf{K}' is bounded by a factor of $1 \pm \lambda/\sigma^2$ times the original predictive mean.*

$$\mathbb{E}[g(\mathbf{x}_*)] = \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{E}[g'(\mathbf{x}_*)]. \quad (32)$$

Proof. Let $\mathbf{k}_* \triangleq [k(\mathbf{x}_*, \mathbf{x}_i)]_{i=1}^n$ where \mathbf{x}_i denotes the i -th training data point. We have:

$$\begin{aligned} \mathbb{E}[g(\mathbf{x}_*)] &= \frac{1}{2} ((\mathbf{k}_* + \mathbf{y})^\top \mathbf{Q}^{-1} (\mathbf{k}_* + \mathbf{y}) - \mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{k}_* - \mathbf{y}^\top \mathbf{Q}^{-1} \mathbf{y}) \\ &= \frac{1}{2} \left(1 \pm \frac{\lambda}{\sigma^2}\right) \left((\mathbf{k}_* + \mathbf{y})^\top \mathbf{Q}'^{-1} (\mathbf{k}_* + \mathbf{y}) - \mathbf{k}_*^\top \mathbf{Q}'^{-1} \mathbf{k}_* - \mathbf{y}^\top \mathbf{Q}'^{-1} \mathbf{y} \right) \\ &= \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{E}[g'(\mathbf{x}_*)], \end{aligned} \quad (33)$$

where the first and third equations follow from adding and subtracting the same terms to the expression of $g(\mathbf{x}_*)$ – see Eq. (2) – while the second equation follows from applying Lemma 6 above. \square

Lemma 8. *Let \mathbf{K}' denote an approximation that is λ -close to the original kernel \mathbf{K} . The induced predictive variance of \mathbf{K}' is bounded by a factor of $1 \pm \lambda/\sigma^2$ of the original predictive variance up to a constant bias of λ/σ^2 ,*

$$\mathbb{V}[g(\mathbf{x}_*)] = \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{V}[g'(\mathbf{x}_*)] \pm \frac{\lambda}{\sigma^2}. \quad (34)$$

Proof. Following the definition of the Gaussian kernel, we assume that the signal of the SE (Squared Exponential) kernel is unitary⁸. As such,

$$\begin{aligned} \mathbb{V}[g(\mathbf{x}_*)] &= 1 - \mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{k}_* \\ &= 1 - \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbf{k}_*^\top \mathbf{Q}'^{-1} \mathbf{k}_* \\ &= \left(1 \pm \frac{\lambda}{\sigma^2}\right) \left(1 - \mathbf{k}_*^\top \mathbf{Q}'^{-1} \mathbf{k}_*\right) \pm \frac{\lambda}{\sigma^2} \\ &= \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbb{V}[g'(\mathbf{x}_*)] \pm \frac{\lambda}{\sigma^2}, \end{aligned} \quad (35)$$

where (again) the above equation follows straightforwardly from applying Lemma 6 and standard algebraic manipulation. Lemma 7 and Lemma 8 thus provide an explicit bound on the difference between the original and approximated predictive distributions. We will now establish another bound on the difference between the original and approximated negative log likelihoods (i.e., the training objectives) in Lemma 9 and Lemma 10 below. \square

Lemma 9. *Let \mathbf{K}' denote an approximation that is λ -close to the original kernel \mathbf{K} . Let $\mathbf{Q} = \mathbf{K} + \sigma^2 \mathbf{I}$ and $\mathbf{Q}' = \mathbf{K}' + \sigma^2 \mathbf{I}$. We have:*

$$\log |\mathbf{Q}'| = \left(1 \pm \tau_{\lambda, \sigma}(\mathbf{K})\right) \log |\mathbf{Q}|. \quad (36)$$

where the spectral constant $\tau_{\lambda, \sigma}(\mathbf{K})$ of \mathbf{K} is defined below:

$$\tau_{\lambda, \sigma}(\mathbf{K}) \triangleq \frac{\max\left(\left|\log\left(1 + \frac{\lambda}{\sigma^2}\right)\right|, \left|\log\left(1 - \frac{\lambda}{\sigma^2}\right)\right|\right)}{\min\left(\left|\log(\lambda_{\min}(\mathbf{K}) + \sigma^2)\right|, \left|\log(\lambda_{\max}(\mathbf{K}) + \sigma^2)\right|\right)}. \quad (37)$$

⁸This simplifies the analysis and does not restrict the expressiveness of the kernel since we can either normalize the output or absorb it into the length-scales (i.e., the θ_i).

Proof. Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ and $\lambda'_1 \leq \lambda'_2 \leq \dots \leq \lambda'_n$ be the eigenvalues of \mathbf{K} and \mathbf{K}' respectively. Applying the Courant-Fischer theorem on the result obtained in Lemma 6, we have:

$$\lambda'_i + \sigma^2 = \left(1 \pm \frac{\lambda}{\sigma^2}\right) (\lambda_i + \sigma^2). \quad (38)$$

This implies:

$$\begin{aligned} \log |\mathbf{Q}'| &\leq \sum_{i=1}^n \left| \log(\lambda'_i + \sigma^2) \right| = \sum_{i=1}^n \left| \log(\lambda_i + \sigma^2) + \log \left(1 \pm \frac{\lambda}{\sigma^2}\right) \right| \\ &\leq \sum_{i=1}^n \left| \log(\lambda_i + \sigma^2) \right| + \sum_{i=1}^n \max \left(\left| \log \left(1 + \frac{\lambda}{\sigma^2}\right) \right|, \left| \log \left(1 - \frac{\lambda}{\sigma^2}\right) \right| \right) \\ &\leq \left(1 + \tau_{\lambda, \sigma}(\mathbf{K})\right) \sum_{i=1}^n \left| \log(\lambda_i + \sigma^2) \right| = \left(1 + \tau_{\lambda, \sigma}(\mathbf{K})\right) \log |\mathbf{Q}|. \end{aligned} \quad (39)$$

Similarly, by symmetry, we have:

$$\begin{aligned} \log |\mathbf{Q}'| &\geq \sum_{i=1}^n \left| \log(\lambda_i + \sigma^2) \right| - \sum_{i=1}^n \max \left(\left| \log \left(1 + \frac{\lambda}{\sigma^2}\right) \right|, \left| \log \left(1 - \frac{\lambda}{\sigma^2}\right) \right| \right) \\ &\geq \left(1 - \tau_{\lambda, \sigma}(\mathbf{K})\right) \sum_{i=1}^n \left| \log(\lambda_i + \sigma^2) \right| = \left(1 - \tau_{\lambda, \sigma}(\mathbf{K})\right) \log |\mathbf{Q}|. \end{aligned} \quad (40)$$

Together, Eq. (39) and Eq. (40) imply $\log |\mathbf{Q}'| = \left(1 \pm \tau_{\lambda, \sigma}(\mathbf{K})\right) \log |\mathbf{Q}|$. \square

Lemma 10. Let \mathbf{K}' denote an approximation that is λ -close to the original kernel \mathbf{K} . With $\tau_{\lambda, \sigma}(\mathbf{K})$ previously defined in Lemma 9, we have:

$$\ell'(\Theta) = \left(1 \pm \max \left(\tau_{\lambda, \sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2} \right)\right) \ell(\Theta). \quad (41)$$

where $\ell(\Theta)$ and $\ell'(\Theta)$ respectively denote the negative log likelihood of the full GP and the approximated GP evaluated at the hyper-parameters $\Theta = \text{diag}[\theta_1^2, \theta_2^2, \dots, \theta_d^2]$ as defined previously.

Proof. We have:

$$\begin{aligned} \ell'(\Theta) &= \frac{1}{2} \log |\mathbf{Q}'| + \frac{1}{2} \mathbf{y}^\top \mathbf{Q}'^{-1} \mathbf{y} \\ &= \frac{1}{2} \left(1 \pm \tau_{\lambda, \sigma}(\mathbf{K})\right) \log |\mathbf{Q}| + \frac{1}{2} \left(1 \pm \frac{\lambda}{\sigma^2}\right) \mathbf{y}^\top \mathbf{Q}^{-1} \mathbf{y} \\ &= \left(1 \pm \max \left(\tau_{\lambda, \sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2} \right)\right) \frac{1}{2} \left(\log |\mathbf{Q}| + \mathbf{y}^\top \mathbf{Q}^{-1} \mathbf{y} \right) \\ &= \left(1 \pm \max \left(\tau_{\lambda, \sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2} \right)\right) \ell(\Theta). \end{aligned} \quad \square \quad (42)$$

Using the result of Lemma 10 above, we can further analyze how the quality of the optimized parameter $\Theta'_* = \arg \max_{\Theta} \ell'(\Theta)$ of the approximated training objective compares to the true optimizer of the original objective function $\Theta_* = \arg \max_{\Theta} \ell(\Theta)$ in Lemma 11 below.

Lemma 11. Let Θ_* and Θ'_* denote the optimal hyper-parameters obtained by respectively minimizing the negative log likelihood of the full GP and the approximated GP. We have:

$$\ell'(\Theta'_*) = \left(1 \pm \max \left(\tau_{\lambda, \sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2} \right)\right) \ell(\Theta_*). \quad (43)$$

Proof. By Lemma 10, we have:

$$\begin{aligned} \ell'(\Theta'_*) &\leq \ell'(\Theta_*) \\ &\leq \left(1 + \max \left(\tau_{\lambda, \sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2} \right)\right) \ell(\Theta_*) \end{aligned} \quad (44)$$

and

$$\begin{aligned}\ell'(\Theta'_*) &\geq \left(1 - \max\left(\tau_{\lambda,\sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2}\right)\right) \ell(\Theta'_*) \\ &\geq \left(1 - \max\left(\tau_{\lambda,\sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2}\right)\right) \ell(\Theta_*).\end{aligned}\quad (45)$$

Together, these results imply $\ell'(\Theta'_*) = (1 \pm \max(\tau_{\lambda,\sigma}(\mathbf{K}), \frac{\lambda}{\sigma^2})) \ell(\Theta_*)$. \square

Lemma 12. *Let $\delta \in (0, 1)$ and let \mathbf{K}' denote an approximation of \mathbf{K} for which $\|\mathbf{K} - \mathbf{K}'\|_2^2 \leq \lambda^2$ with probability at least $1 - \delta$ uniformly over the entire parameter space. Let Θ_* and Θ'_* denote the optimal hyper-parameters obtained by respectively minimizing the negative log likelihood of the full GP and the approximated GP. Then, with probability $1 - \delta$, the following holds:*

$$\mathbb{E}[g'(\mathbf{x}_*; \Theta'_*)] = (1 \pm \rho(\lambda, \sigma, \Theta_*, \Theta'_*)) \cdot \mathbb{E}[g(\mathbf{x}_*; \Theta_*)] + \wp(\lambda, \sigma, \Theta_*, \Theta'_*) \quad (46)$$

where $\rho(\lambda, \sigma, \Theta_*, \Theta'_*)$ and $\wp(\lambda, \sigma, \Theta_*, \Theta'_*)$ are constant with respect to $\lambda, \sigma, \Theta_*, \Theta'_*$

Proof. We have:

$$\begin{aligned}\mathbb{E}[g(\mathbf{x}_*; \Theta_*)] &= \mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{y} \Big|_{\Theta_*} \\ &= \frac{1}{2} \left[(\mathbf{k}_* + \mathbf{y})^\top \mathbf{Q}^{-1} (\mathbf{k}_* + \mathbf{y}) - \mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{k}_* + \log |\mathbf{Q}| \right] \Big|_{\Theta_*} - \frac{1}{2} \ell(\Theta_*) \\ &\geq -\frac{1}{2} \left[\ell(\Theta_*) + 1 - \sum_{i=1}^n \log(\lambda_i + \sigma^2) \Big|_{\Theta_*} \right]\end{aligned}\quad (47)$$

On the other hand, we have:

$$\begin{aligned}\mathbb{E}[g(\mathbf{x}_*; \Theta_*)] &= \mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{y} \\ &\leq \frac{1}{2} \left[\mathbf{k}_*^\top \mathbf{Q}^{-1} \mathbf{k}_* + \mathbf{y}^\top \mathbf{Q}^{-1} \mathbf{y} \right] \Big|_{\Theta_*} \\ &\leq \frac{1}{2} \left[\ell(\Theta_*) + 1 - \sum_{i=1}^n \log(\lambda_i + \sigma^2) \Big|_{\Theta_*} \right]\end{aligned}\quad (48)$$

Thus, we have:

$$\mathbb{E}[g(\mathbf{x}_*; \Theta_*)] = \pm \frac{1}{2} \left[\ell(\Theta_*) + 1 - \sum_{i=1}^n \log(\lambda_i + \sigma^2) \Big|_{\Theta_*} \right] \quad (49)$$

and by symmetry:

$$\begin{aligned}\mathbb{E}[g'(\mathbf{x}_*; \Theta'_*)] &= \pm \frac{1}{2} \left[\ell'(\Theta'_*) + 1 - \sum_{i=1}^n \log(\lambda'_i + \sigma^2) \Big|_{\Theta'_*} \right] \\ &= (1 \pm \rho(\lambda, \sigma, \Theta_*, \Theta'_*)) \cdot \mathbb{E}[g(\mathbf{x}_*; \Theta_*)] + \wp(\lambda, \sigma, \Theta_*, \Theta'_*)\end{aligned}\quad (50)$$

where $\wp(\Theta_*, \Theta'_*)$ is a constant as defined below:

$$\begin{aligned}\rho(\lambda, \sigma, \Theta_*, \Theta'_*) &\triangleq \max\left(\tau_{\lambda,\sigma}(\mathbf{K}), \tau_{\lambda,\sigma}(\mathbf{K}'), \frac{\lambda}{\sigma^2}\right) \\ \wp(\lambda, \sigma, \Theta_*, \Theta'_*) &\triangleq \left(\sum_{i=1}^n \log \frac{\lambda_i + \sigma^2}{\lambda'_i + \sigma^2} \Big|_{\Theta_*} \right) \pm \rho(\lambda, \sigma, \Theta_*, \Theta'_*) \cdot \left(1 - \sum_{i=1}^n \log(\lambda_i + \sigma^2) \Big|_{\Theta_*} \right)\end{aligned}$$

\square

C Model Parameterization and Practical Implementation

Our embedding algorithm is based on a VAE implementation where the latent prior, posterior and likelihood of the data generation process are represented via separate mixtures of k Gaussian distributions over a 4-dimensional space. For the latent prior, we set (and fixed) the means of each Gaussian component (i.e., the prior cluster means) at k equidistant points on a 4-dimensional sphere centered at zero with an optimizable radius. For the latent posterior and likelihood, the mean and covariance entries of each component in the mixture are parameterized as outputs of their respective neural networks, which we refer to as Gaussian nets.

In turn, the Gaussian nets are parameterized separately. Each starts with a linear layer comprising of 10 neurons whose outputs are fed simultaneously to two separate hidden (linear) layers with 10 hidden neurons each. Their outputs are then used to form the mean and covariance entries of the corresponding Gaussian component. All neurons are activated by a ReLU unit, and in addition, the (batch) outputs of the first linear layer are also standardized via a learnable 1D batch-norm layer to ensure the stability of batch optimization. The mixing weights that combine such Gaussian nets in the mixtures are also parameterized as the outputs of a linear layer with $k = 8$ neurons where $k = 8$ is also the number of components in our mixture.

The above parameterized latent prior, posterior, and likelihood are then connected in the variational lower-bound (ELBO) as expressed in the first two terms of Eq. (9). This ELBO objective is then combined with two regularization terms weighted with (manually tuned) parameters $\alpha = 8.0$ and $\beta = 1.2$ as detailed in Eq. (10). The entire function is optimized via gradient descent using the standard Adam optimizer with the default setting implemented in PyTorch.

Once learned, the outputs of the latent posterior were used as the encoded data which were fed as input to our revisited SSGP. For a practical implementation, we also found that additionally passing the encoded data to the latent likelihood generates a reconfigured version of the original data which helps to marginally improve the performance. All of our reported results below are generated with respect to this version of reconfiguration. All of our implementations of GP, SSGP and revisited SSGP that makes use of the output of this reconfiguration process, are also in PyTorch. Our experimental code is released at https://github.com/hqminh/gp_sketch_nips.

D Additional Empirical Results and Visualizations

This section provides additional empirical results and visualizations that complement and corroborate the reported results in the main text. In particular, we provide: (a) a more refined and comprehensive visualization of how our embedding algorithm (Section 3.3) re-configures data across different settings; and (b) an extended comparison with SSGP at different levels of sample complexity when evaluated on middle (10K data points) and large (500K data points) data sets. All data samples used in this section were extracted from the GAS SENSOR dataset [4]⁹.

D.1 The Effect of Data Re-configuration: A Visual Demonstration

This section describes an ablation study to demonstrate the effectiveness of our data re-configuration component (i.e., to approximately meet the practical Conditions 1-3 of our refined analysis). Specifically, we demonstrate this by contrasting the scatter plots of data embeddings (see Fig. 3) before and after reconfiguration using our algorithm in Section 3.3 below. The visualizations are shown for 3 different samples of data, each of which has 10K data points.

For each data sample, its embedding was clustered and re-clustered before and after its reconfiguration. Both clustering processes were generated independently using K-Means to provide an objective visual measurement of the reconfiguration effects of our algorithm.

Observing the above visual excerpts, it appears that after reconfiguration, the clusters across different data samples all became significantly more disengtangled with a visibly increased distance between their cluster centers. This provides conclusive evidence to the data disengtangling effect of our

⁹The entire GAS SENSOR dataset contains approximately 4M data points. However, in the body of this paper we only used a sample of 500K points to conduct our experiments.

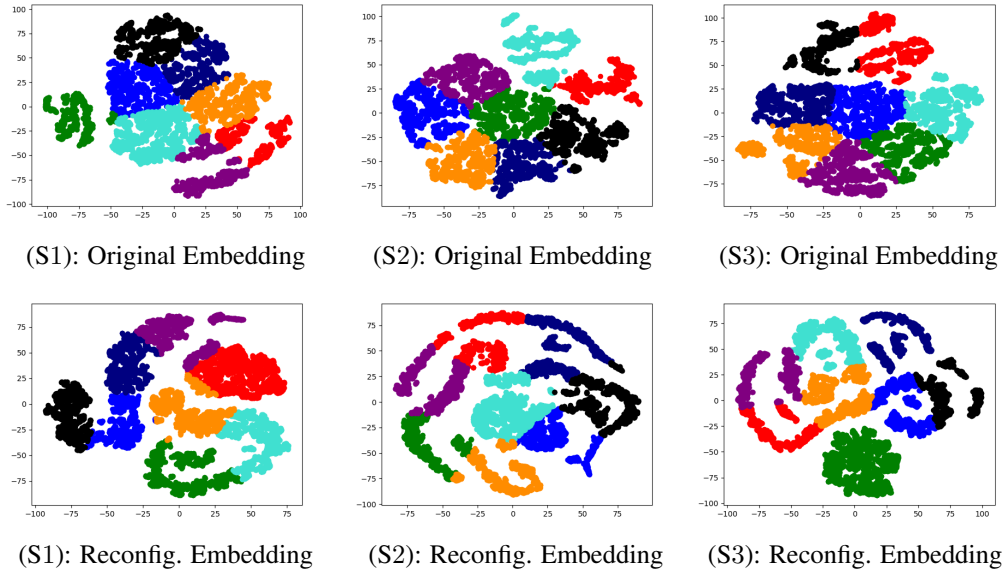


Figure 3: Visualizations of original (top) and reconfigured (bottom) data embeddings for 3 different (randomly selected) data samples annotated with S1, S2 and S3, respectively. Each visual excerpt is annotated with different colors corresponding to the different clusters that the data belong to. All visualizations are generated using T-SNE [41].

embedding algorithm. More importantly, this demonstration further reveals a practical aspect of data that has not been investigated before in the existing literature of GP:

Data (especially experimental data) is often the manifestation of how latent concepts that underlie them were observed and depending on specific parameters of the observation process, these concepts might manifest differently in either more or less useful forms for learning. This raises the question of whether one can reorient the observation process to increase the utility of such data.

In this vein of thought, to address the above question, our data reconfiguration algorithm can be considered to be one potential solution which uses a parameterized construction of a latent space to provide a handle on how to reorient the latent concepts that underlie our data. For an intuitive example, imagine how we would look at the outside world via a narrowed pigeonhole. With different viewing angles, we would perceive the same scene outside differently and apparently, some angles provide a much better perception of that scene (thus, allowing us to interpret the scene more accurately).

In technical terms, such a reorientation is implemented in our algorithm via the regularization of the mixture composition of the latent prior while constraining the entire embedding process to have it reflected on the latent posterior – see Eq. (10) – which was used to encode data into a latent space that exhibits the desired separation effect. Such separation/disentanglement is then shown (empirically) to be richer in information and can be leveraged to improve the sample complexity of SSGP (see Section D.2), thus supporting our theoretical analysis in Appendix A.

D.2 Comparison with SSGP on Large Data

To demonstrate the effectiveness of the data disentanglement in reducing the sample complexity of SSGP, we compare the performance of SSGP and our revisited SSGP (which was instead applied on the reconfigured space of data) at different levels of sample complexity. All results were generated for two different data samples extracted from GAS-SENSOR [4]. One of these (containing 500K data points) is in fact on the same scale of the most extensive datasets used in the GP literature. All performance plots were visualized in Fig. 4 below. For each experiment, the data sample is divided into a train/test partition with an 8-2 ratio. All results were averaged over 5 independent runs.

We see that our revised SSGP consistently achieves better performance than its SSGP counterpart at all complexity levels. In particular, in all cases of the 10K setting, the performance of our revised

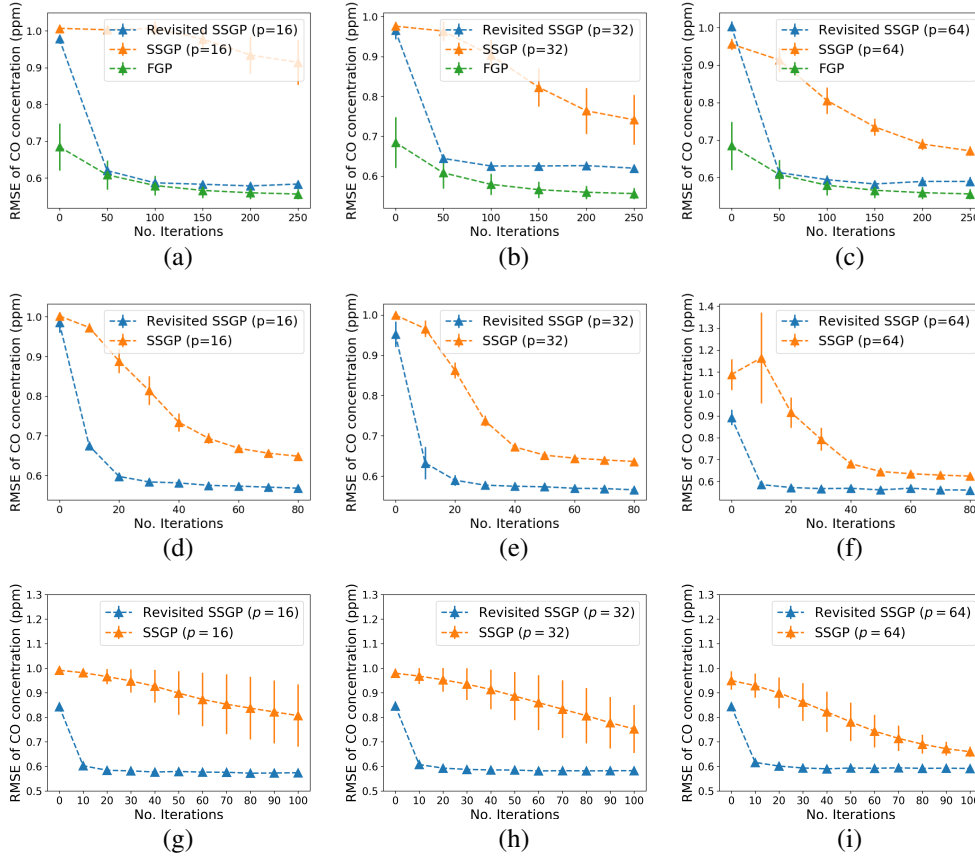


Figure 4: Graphs of performance comparisons between the full GP, our revised SSGP and the traditional SSGP on a 10K sample (a-c); 500K sample (d-f) and the entire GAS SENSOR dataset [4] totalling approximately 4M data points (g-i). In both settings, the performance differences were plotted at $p = 16, 32$ and 64 . Note that for the 500K sized sample and the entire dataset (which contains 4M data points), the full GP model is not applicable due to its inability (memory- and computation-wise) to store and invert the corresponding large covariance matrix.

SSGP is also shown to approach closely that of the full GP, which serves as a gold-standard lower-bound on the achievable prediction error. This concludes our empirical demonstration which (we believe) has shown that with a proper reconfiguration of data, the predictive performance of a GP can be well-preserved at a much cheaper sample complexity as compared to the previous conservative estimate yielded by SSGP. In fact, the performance trend of SSGP as depicted in the above graphs shows that with more samples, it also slowly converges towards the performance level of GP and our revised SSGP but at a much greater sample complexity – see the shrinking performance gap between revisited SSGP and SSGP from Fig. 4d to Fig. 4e; and similarly, from Fig. 4g to Fig. 4h.