

---

# Improved Convergence Rates for Sparse Approximation Methods in Kernel-Based Learning

---

Sattar Vakili<sup>1</sup> Jonathan Scarlett<sup>2</sup> Da-shan Shiu<sup>1</sup> Alberto Bernacchia<sup>1</sup>

## Abstract

Kernel-based models such as kernel ridge regression and Gaussian processes are ubiquitous in machine learning applications for regression and optimization. It is well known that a major downside for kernel-based models is the high computational cost; given a dataset of  $n$  samples, the cost grows as  $\mathcal{O}(n^3)$ . Existing sparse approximation methods can yield a significant reduction in the computational cost, effectively reducing the actual cost down to as low as  $\mathcal{O}(n)$  in certain cases. Despite this remarkable empirical success, significant gaps remain in the existing results for the analytical bounds on the error due to approximation. In this work, we provide novel confidence intervals for the Nyström method and the sparse variational Gaussian process approximation method, which we establish using novel interpretations of the approximate (surrogate) posterior variance of the models. Our confidence intervals lead to improved performance bounds in both regression and optimization problems.

## 1. Introduction

Kernel-based modeling is an elegant and natural technique to extend linear models to highly general nonlinear ones with great representation capacity. Powerful predictors and uncertainty estimates have made this approach very common across several application domains (Shahriari et al., 2016). It is, however, well known that regression using standard kernel-based methods incurs a high  $\mathcal{O}(n^3)$  computational cost in the number of data points  $n$ , which is the main limiting factor in many applications. Fortunately, sparse approximation methods have been developed for kernel-based models, which significantly reduce the computational cost. In particular, the computational cost typically reduces to

$\mathcal{O}(nm^2)$ , with  $m \ll n$  being the number of the *inducing points* that serve to summarize the model. Despite their great popularity in practice, significant gaps remain in the mathematical understanding of such methods' bounds on the error due to approximation, and the implications for regression and optimization problems. This challenging issue is the main focus of this work.

The literature on kernel-based models has developed under two mostly separate approaches with different terminologies (e.g., see, Kanagawa et al., 2018): kernel ridge regression (KRR) and Gaussian process (GP) models. These two approaches are based on frequentist and Bayesian interpretations, respectively. Nonetheless, they lead to the same solution for regression. In this work, we cover both terminologies and provide insights into their connections.

In the GP approach, a prior (surrogate) GP distribution is assumed for the unknown function  $f$ , and Gaussian noise is assumed. The regressor is then the posterior mean (the maximum likelihood estimate) of  $f$ , conditioned on the data points. In contrast, in KRR approach, the regressor is the regularized least squares estimator of  $f$ , within a reproducing kernel Hilbert space (RKHS). For further details, see (Rasmussen & Williams, 2006) on the former, and (Schölkopf et al., 2002) on the latter.

For GPs, sparse approximation methods find a Bayesian *variational* interpretation and are referred to as sparse variational Gaussian processes (SVGPs) in the literature. The SVGP method was introduced in (Titsias, 2009) and has been studied in several works including (Hensman et al., 2013; Matthews et al., 2016; Burt et al., 2019; Rossi et al., 2021). See also (Liu et al., 2020) for a recent overview. The SVGP models enjoy great popularity in practice, with implementations in TensorFlow (Abadi et al., 2015) and GPflow (Matthews et al., 2017) libraries.

For KRR, sparse approximation methods are interpreted as regularized least squares estimators of  $f$  within a reduced rank RKHS. This enables a Nyström low rank matrix approximation; the corresponding sparse approximation method is thus often referred to as the Nyström method. Representative works include (Williams & Seeger, 2001; Seeger et al., 2003; Rudi et al., 2015; Lin & Cevher, 2018).

---

<sup>1</sup>MediaTek Research <sup>2</sup>National University of Singapore. Correspondence to: Sattar Vakili <sattr.vakili@mtkresearch.com>.

**Contributions and paper structure.** In Section 2, we introduce the notations and preliminaries on kernel-based models and their sparse approximations. In Section 3, we establish a novel representation of the approximate posterior variance of a GP model in terms of *i*) the error due to the low rank projection of the approximation method, *ii*) the prediction error from noise-free observations in the low rank RKHS, and *iii*) the effect of noise (see Theorem 3.2). We then use this result to derive novel confidence intervals for approximate kernel-based models (see Theorem 3.5). These results/tools are broadly applicable in kernel-based models, and may be of general interest.

In Section 4, we use our confidence intervals to provide explicit bounds on the prediction error of a regression problem using sparse approximation methods (see, Theorem 4.1). In particular, we prove that the sparse approximate prediction  $\bar{\mu}_n$  of  $f$ , using  $n$  well distributed data points, converges uniformly to  $f$ , with a number of inducing points no larger than  $\tilde{O}(\gamma_k(n))$ ; here,  $\gamma_k(n)$  denotes the maximal information gain of the kernel, and can be accurately bounded for specific kernels (e.g., see Srinivas et al., 2010; Vakili et al., 2021b;c). For example, in the case of the squared exponential (SE) kernel, our results imply that with  $m = \mathcal{O}((\log(n))^{d+2})$  inducing points, where  $d$  is the dimension of the input,  $\bar{\mu}_n$  uniformly converges to  $f$  at rate  $\tilde{O}(\frac{1}{\sqrt{n}})$ . In the case of the Matérn family of kernels, with  $m = \tilde{O}(n^{\frac{d}{2\nu+d}})$ ,  $\bar{\mu}_n$  uniformly converges to  $f$  at a rate  $\tilde{O}(n^{-\frac{\nu}{2\nu+d}})$ , where  $\nu$  is the smoothness parameter of the kernel. Here and subsequently, the notations  $\mathcal{O}$  and  $\tilde{O}$  are used for standard mathematical order, and that up to hiding logarithmic factors, respectively.

In Section 5, we consider a kernel-based black-box optimization (kernel-based bandit) problem using sparse approximation methods. Leveraging our novel confidence intervals, we prove  $\tilde{O}(\sqrt{\gamma_k(N)N})$  regret bounds on a time horizon  $N$ , for an optimization algorithm based on batch observations (see Theorem 5.2). This is optimal up to logarithmic terms in the cases where lower bounds on the regret are known. Our algorithm is an adaptation of Batched Pure Exploration (BPE) algorithm of (Li & Scarlett, 2022), in which we replace the exact GP statistics with their sparse approximations. We thus refer to this algorithm as Sparse BPE (S-BPE). To our knowledge, this is the first algorithm for kernel-based bandits using sparse approximation methods while achieving near-optimal regret bounds.

### 1.1. Related Work

The celebrated work of (Burt et al., 2019) proved a fundamental result on bounding the Kullback-Leibler (KL) divergence between the approximate and exact GP posteriors in SVGPs (see also Burt et al., 2020). In particular, they showed how large the number  $m$  of inducing points should

be chosen in relation to the sample size  $n$  in order to ensure that the expected KL divergence vanishes for large  $n$ . As a consequence, they showed that the approximate posterior mean and variance converge to the exact ones when  $m$  is large enough. These results hold in expectation, where the expectation is taken with respect to a prior on the dataset, as well as the prior GP distribution of  $f$ . (Nieman et al., 2021) proved similar bounds on the KL divergence of the approximate and exact (surrogate) GP posteriors, when  $f$  is a fixed function in the RKHS. Their result also holds in expectation, where the expectation is taken only with respect to a prior on the dataset (and not a prior on  $f$ ).

(Nieman et al., 2021) also proved a result similar to our Theorem 4.1 on the convergence rate of the approximate prediction  $\bar{\mu}_n$  to the fixed  $f$  in the RKHS. Neither this result nor our Theorem 4.1 imply the other, and the comparison can be summarized as follows. Both results find the same  $m = \tilde{O}(n^{\frac{d}{2\nu+d}})$  rates for the number of inducing points in the case of the Matérn kernel. Our convergence measure is stronger than theirs in the sense that our convergence guarantees are uniform (given in terms of the  $L^\infty$  norm), while their convergence guarantees are given in terms of Hellinger distance. Thus, their convergence result is in an average sense rather than in a uniform sense, where the average is taken with respect to a prior distribution on the data. In contrast, in the setting of Theorem 4.1, the data is collected based on uncertainty reduction. While these settings are different, they serve the same purpose of ensuring a nearly uniform dataset across the domain. In addition, we prove novel general confidence intervals based on approximate statistics (Theorem 3.5) and near-optimal regret bounds for black-box optimization (Theorem 5.2), which are new compared to existing works.

It is also noteworthy that the bound on  $m$  required for vanishing KL divergence in (Burt et al., 2019; 2020) is larger than the bound on  $m$  required for the convergence of  $\bar{\mu}_n$  to  $f$ . For example, in the case of Matérn kernels, the results of (Burt et al., 2019; 2020) require  $m = \mathcal{O}(n^{\frac{2d}{2\nu+d}})$  inducing points (see Burt et al., 2020, Table 1), which is a looser bound compared to  $m = \mathcal{O}(n^{\frac{d}{2\nu+d}})$  in (Nieman et al., 2021) and in our results.

**Selection of inducing points.** Several methods for selecting the inducing points have been employed in the existing work. (Burt et al., 2019) considered selecting the inducing points based on a  $k$  determinantal point process ( $k$ -DPP) on the data. However, the cost of sampling from an exact  $k$ -DPP may defy the computational gain from the sparse approximation. They thus suggested to use a *Markov Chain Monte Carlo* based approximate sampling from a  $k$ -DPP. In practice, the inducing points are sometimes selected based on  $k$ -means clustering, without theoretical guarantees (Hensman et al., 2013). Selecting the inducing points

based on ridge leverage score (RLS) sampling, which captures the importance of an individual data point, comes with strong theoretical guarantees (Alaoui & Mahoney, 2014). (Musco & Musco, 2016) introduced a practical recursive RLS with guarantees for low computational complexity (see Lemma 3.8). (Chen & Yang, 2021) improved the bounds on computational complexity of recursive RLS sampling. (Calandriello et al., 2019) showed that selecting the inducing points based on approximate posterior variance enjoys the same performance guarantees as RLS. In our analysis, we use the theoretical guarantees on inducing points given in (Musco & Musco, 2016). However, we also establish general sufficient conditions that may apply to other methods of selecting the inducing points.

**Kernel-based black-box optimization.** Black box optimization of kernel-based models is often referred to as kernel-based bandits, GP bandits, or Bayesian optimization in the literature. The terminology “bandit” signifies zeroth-order observations from  $f$ , in contrast to, e.g., first-order observations in gradient descent. Classical algorithms such as GP upper confidence bound (GP-UCB) (Srinivas et al., 2010; Chowdhury & Gopalan, 2017), expected improvement (EI) (Mockus et al., 1978), and Thompson sampling (GP-TS) (Chowdhury & Gopalan, 2017) sequentially select the observation points based on a score referred to as *acquisition function*. When the objective function  $f$  is a fixed function in the RKHS, the best known regret bounds for these algorithms scales as  $\mathcal{O}(\gamma_k(N)\sqrt{N})$  over a time horizon  $N$  (see the references above). The  $\mathcal{O}(\gamma_k(N)\sqrt{N})$  scaling is not tight in general, and may even fail to be sublinear in many cases of interest, since  $\gamma_k(N)$ , although sublinear, may grow faster than  $\sqrt{N}$ . It remains an open problem whether the suboptimal regret bounds of these acquisition based algorithms is a fundamental limitation or a shortcoming of their proof (see Vakili et al., 2021e, for the details). On discrete domains, the SupKernelUCB algorithm was shown to have an  $\tilde{\mathcal{O}}(\sqrt{\gamma_k(N)N})$  regret bound (Valko et al., 2013); moreover, this result extends to compact domains under mild Lipschitz-like assumptions (Janz et al., 2020). This nearly matches the lower bounds proven in (Scarlett et al., 2017) for the SE and Matérn kernels.

SupKernelUCB is not considered to be practical. Several more practical algorithms with  $\tilde{\mathcal{O}}(\sqrt{\gamma_k(N)N})$  regret have been proposed recently: a tree-based domain-shrinking algorithm known as GP-ThreDS (Salgia et al., 2021), Robust Inverse Propensity Score (RIPS) for experimental design (Camilleri et al., 2021), and BPE (Li & Scarlett, 2022). The latter two use batches of observations, with RIPS using  $\mathcal{O}(\log(N))$  batches and BPE using  $\mathcal{O}(\log \log(N))$  batches.

There have been several works considering sparse approximation methods in kernel-based optimization algorithms. (Calandriello et al., 2019) introduced budgeted kernelized

bandit (BKB), which is an adaptation of GP-UCB by replacing the exact statistics with sparse approximate ones. They proved that BKB achieves the same  $\tilde{\mathcal{O}}(\gamma_k(N)\sqrt{N})$  regret bound as GP-UCB, where the number of inducing points scales as  $m = \mathcal{O}(\gamma_k(n))$  at time  $n$ . (Vakili et al., 2021d) introduced S-GP-TS, which is an adaptation of GP-TS by replacing the exact statistics with sparse approximate ones. They also proved that S-GP-TS achieves the same  $\tilde{\mathcal{O}}(\gamma_k(N)\sqrt{N})$  regret bound as GP-TS, where the number of inducing points scales as  $m = \mathcal{O}(n^{\frac{2d}{2\nu-d}})$ . The analysis in (Vakili et al., 2021d) relies on KL divergence bounds from (Burt et al., 2019), and their weaker bound on  $m$  is due to the weaker scaling of the number of inducing points in (Burt et al., 2019).

We also highlight that a recent concurrent work again attained  $\tilde{\mathcal{O}}(\gamma_k(N)\sqrt{N})$  regret with a low time complexity via a distinct approach of resampling the same points multiple times (Calandriello et al., 2022). In addition, scalable algorithms for the contextual kernel bandit setting were studied in (Zenati et al., 2022).

Our confidence intervals do not apply to the analysis of the BKB or S-GP-TS algorithms, but we show that they apply to an adaptation of BPE and offer a near-optimal  $\tilde{\mathcal{O}}(\sqrt{\gamma_k(N)N})$  regret bound, which is an  $\mathcal{O}(\sqrt{\gamma_k(N)})$  improvement over BKB and S-GP-TS.

## 2. Kernel-Based Models and Sparse Approximations

In this section, we give an overview of the relevant background on regression using kernel-based models, as well as the sparse approximation methods.

### 2.1. Kernel-Based Regression

Consider a noisy dataset  $\mathcal{D}_n = \{(x_i, y_i)_{i=1}^n\}$ , where  $y_i = f(x_i) + \epsilon_i$ , and  $\epsilon_i$  is the observation noise. We use the vectorized notations  $X_n = [x_1, \dots, x_n]^\top$ ,  $Y_n = [y_1, \dots, y_n]^\top$ . Consider a positive definite kernel  $k(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ ,  $\mathcal{X} \subset \mathbb{R}^d$ ,  $d \in \mathbb{N}$ . We have the following closed form expression for the kernel-based regressor:

$$\mu_{k, \tau, X_n, Y_n}(\cdot) = k_{X_n}^\top(\cdot) (k_{X_n, X_n} + \tau^2 I_n)^{-1} Y_n, \quad (1)$$

where  $k_{X_n}(\cdot) = [k(\cdot, x_1), \dots, k(\cdot, x_n)]^\top$ ,  $k_{X_n, X_n} = [k(x_i, x_j)]_{i, j=1}^n$  is the  $n \times n$  kernel matrix,  $\tau$  is a free parameter, and  $I_n$  is the  $n \times n$  identity matrix. We use the shorthand notation  $\mu_n = \mu_{k, \tau, X_n, Y_n}$  when the values of  $k, \tau, X_n, Y_n$  are clear from the context.

In the KRR approach,  $\mu_n$  is interpreted as the regularized least squares estimator with regularization parameter  $\tau^2$ :

$$\mu_n = \arg \min_{g \in \mathcal{H}_k} \sum_{i=1}^n (y_i - g(x_i))^2 + \tau^2 \|g\|_{\mathcal{H}_k}^2, \quad (2)$$

where  $\mathcal{H}_k$  and  $\|\cdot\|_{\mathcal{H}_k}$  denote, respectively, the RKHS corresponding to a kernel  $k$  and its norm.

In the GP approach,  $\mu_n$  is interpreted as the posterior mean of a centered surrogate GP model with kernel  $k$ , conditioned on  $\mathcal{D}_n$ , assuming zero-mean Gaussian observation noise with variance  $\tau^2$ . The KRR and GP interpretations are known to be equivalent; e.g., see, (Kanagawa et al., 2018).

Under the GP approach, the posterior covariance of the (surrogate) GP model is often leveraged as a measure of uncertainty in the regression model. The posterior covariance is given in closed form as

$$k_{\tau, X_n}(\cdot, \cdot') = k(\cdot, \cdot') - k_{X_n}^\top(\cdot) (k_{X_n, X_n} + \tau^2 I_n)^{-1} k_{X_n}(\cdot'). \quad (3)$$

We define the posterior variance as  $\sigma_{k, \tau, X_n}^2(\cdot) = k_{\tau, X_n}(\cdot, \cdot)$ . We use the simplified notations of  $k_n = k_{\tau, X_n}$  and  $\sigma_n = \sigma_{k, \tau, X_n}$ , when the values of the omitted variables are clear from the context.

## 2.2. Sparse Approximation Methods

Kernel-based models provide powerful prediction and uncertainty estimation (posterior covariance) tools for very general nonlinear regression. However, the computation quickly becomes prohibitive due to the high  $\mathcal{O}(n^3)$  cost of matrix inversion. Sparse approximation methods are thus essential tools for applications of kernel-based models to large-scale learning problems. A relatively small number  $m \ll n$  of *inducing points* can be used to obtain approximate prediction and uncertainty estimates with an efficient computation time of  $\mathcal{O}(nm^2)$ . In the following, we give an overview of the relevant sparse approximation methods.

Let  $Z_m = [z_1, z_2, \dots, z_m]^\top \in \mathcal{X}^m$  be the set of inducing points. The approximate regressor and posterior covariance of the (surrogate) GP model are given as (Wild et al., 2021):

$$\begin{aligned} \bar{\mu}_{k, \tau, Z_m, X_n, Y_n}(\cdot) &= V_n^\top(\cdot) Y_n, \\ \bar{k}_{\tau, Z_m, X_n}(\cdot, \cdot') &= k(\cdot, \cdot') - k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} k_{Z_m}(\cdot') \\ &+ k_{Z_m}^\top(\cdot) \left( k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{X_n, Z_m}(\cdot'), \end{aligned} \quad (4)$$

where we define

$$V_n^\top(\cdot) = k_{Z_m}^\top(\cdot) (\tau^2 k_{Z_m, Z_m} + k_{X_n, Z_m}^\top k_{X_n, Z_m})^{-1} \times k_{Z_m, X_n}, \quad (5)$$

and where  $k_{Z_m}(\cdot) = [k(\cdot, z_1), k(\cdot, z_2), \dots, k(\cdot, z_m)]^\top$ ,  $k_{Z_m, Z_m} = [k(z_i, z_j)]_{i, j=1}^m$  is the  $m \times m$  kernel matrix over the inducing points, and  $k_{X_n, Z_m} = [k(x_i, z_j)]_{i=1, \dots, n, j=1, \dots, m}$ . The approximate posterior variance is

defined as  $\bar{\sigma}_{k, \tau, Z_m, X_n}^2(\cdot) = \bar{k}_{\tau, Z_m, X_n}(\cdot, \cdot)$ . We use the simplified notations  $\bar{\mu}_n = \bar{\mu}_{k, \tau, Z_m, X_n, Y_n}$ ,  $\bar{k}_n = \bar{k}_{\tau, Z_m, X_n}$ , and  $\bar{\sigma}_n = \bar{\sigma}_{k, \tau, Z_m, X_n}$ , when the omitted variables are clear from the context.

In the KRR approach,  $\bar{\mu}_n$  is interpreted as the regularized least squares estimator within a reduced rank RKHS induced by  $Z_m$  (e.g., see, Wild et al., 2021). In particular, define

$$\kappa(\cdot, \cdot') = k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} k_{Z_m}(\cdot'), \quad (6)$$

and let  $\mathcal{H}_\kappa$  be the associated RKHS. We then have the following (e.g., see Wild et al., 2021):

$$\bar{\mu}_n = \arg \min_{g \in \mathcal{H}_\kappa} \sum_{i=1}^n (y_i - g(x_i))^2 + \tau^2 \|g\|_{\mathcal{H}_\kappa}^2. \quad (7)$$

In the GP approach, the approximate expressions are obtained through a Bayesian variational method. In particular, a variational family of GP distributions is introduced with mean  $k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} \mu^{(v)}$  and covariance  $k(\cdot, \cdot') - k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} (k_{Z_m, Z_m} - \Sigma^{(v)}) k_{Z_m, Z_m}^{-1} k_{Z_m}(\cdot')$ , where  $\mu^{(v)}$  and  $\Sigma^{(v)}$  (the mean vector and the covariance matrix of a postulated multivariate Gaussian distribution on  $f_{Z_m}$ ) are the variational parameters. Following the standard steps in Bayesian variational methods, the variational parameters are chosen to minimize the KL divergence between the approximate and exact posteriors, or equivalently, maximize the evidence lower bound (ELBO). (Titsias, 2009) showed that this convex optimization problem can be explicitly solved, leading to the optimal variational parameters:

$$\begin{aligned} \mu_*^{(v)} &= k_{Z_m, Z_m} (\tau^2 k_{Z_m, Z_m} + k_{X_n, Z_m}^\top k_{X_n, Z_m})^{-1} k_{Z_m, X_n} Y_n \\ \Sigma_*^{(v)} &= k_{Z_m, Z_m} \left( k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{Z_m, Z_m}. \end{aligned}$$

Substituting  $\mu_*^{(v)}$  and  $\Sigma_*^{(v)}$  in the mean and covariance of the variational distribution results in the expressions for the approximate statistics given in (4).

**Other notations.** For any function  $g : \mathcal{X} \rightarrow \mathbb{R}$ , we write  $g_{X_n} = [g(x_1), g(x_2), \dots, g(x_n)]^\top$ . We also define the noise vector  $E_n = [\epsilon_1, \epsilon_2, \dots, \epsilon_n]^\top$ , and let  $k_{\max} = \sup_{x \in \mathcal{X}} k(x, x)$ , which is assumed to be finite.

For convenience, the most commonly-used notation is briefly summarized in Appendix A.

## 3. Confidence Intervals for Approximate Kernel-Based Regression

Consider the regression problem and its exact solution given in Section 2.1. The posterior variance can be leveraged to quantify the uncertainty in the prediction. In particular, various results exist stating that with probability at least

$1 - \delta$ , the prediction function satisfies  $|f(x) - \mu_n(x)| \leq \beta(\delta)\sigma_n(x)$  (either for fixed  $x$  or simultaneously for all  $x$ ), where the confidence interval width multiplier  $\beta(\delta)$  depends on the properties of the observation noise and the complexity of  $f$  (e.g., see, [Srinivas et al., 2010](#); [Chowdhury & Gopalan, 2017](#); [Vakili et al., 2021a](#)). In this section, we prove similar confidence intervals for the approximate statistics given in [Section 2.2](#). As an intermediate step, we first provide a novel interpretation of the approximate posterior variance, which may be of independent interest.

### 3.1. Approximate Posterior Variance

In this section, we prove a novel expression for the approximate posterior variance as the sum of three terms: *i*) the error due to the projection of  $f \in \mathcal{H}_k$  onto  $\mathcal{H}_\kappa$ , *ii*) the prediction error from noise-free observations within  $\mathcal{H}_\kappa$ , and *iii*) a term capturing the effect of the noise. Given a function  $f \in \mathcal{H}_k$ , let

$$\Pi_\kappa[f] = \arg \min_{g \in \mathcal{H}_\kappa} \|f - g\|_{\mathcal{H}_k} \quad (8)$$

denote the projection of  $f \in \mathcal{H}_k$  onto  $\mathcal{H}_\kappa$ . This projection is the minimum norm interpolator of  $f$  over  $Z_m$  ( $\mu_{k,0,Z_m,fZ_m}(\cdot) = k_{Z_m}^\top(\cdot)k_{Z_m}^{-1}f_{Z_m}$ ), as formalized in the following lemma.

**Lemma 3.1** ([\(Wild et al., 2021\)](#)). *For any  $f \in \mathcal{H}_k$ , the projection of  $f$  onto  $\mathcal{H}_\kappa$  is equivalent to its minimum norm interpolation over  $Z_m$ :  $\Pi_\kappa[f] = \mu_{k,0,Z_m,fZ_m}$ .*

We are now ready to state our first main result on the decomposition of the approximate posterior variance. Recall that  $V_n(\cdot)$  is defined in [\(5\)](#).

**Theorem 3.2.** *Given a dataset  $\mathcal{D}_n$ , inducing points  $Z_m$ , and  $(\bar{\mu}_n, \bar{\sigma}_n)$  given in [Section 2.2](#), we have*

$$\begin{aligned} \bar{\sigma}_n^2(\cdot) &= \sup_{g: \|g\|_{\mathcal{H}_k} \leq 1} (g(\cdot) - \Pi_\kappa[g](\cdot))^2 \\ &+ \sup_{g: \|g\|_{\mathcal{H}_\kappa} \leq 1} (g(\cdot) - V_n^\top(\cdot)g_{X_n})^2 + \tau^2 \|V_n(\cdot)\|_{l^2}^2. \end{aligned}$$

The first term captures the maximum error between any  $g$  in the unit ball of  $\mathcal{H}_k$  and its projection onto  $\mathcal{H}_\kappa$ . The second term captures the maximum prediction error from noise-free observations within the unit ball of  $\mathcal{H}_\kappa$ . Lastly, the third term measures the impact of the observation noise ( $l^2$  norm of the coefficients vector), and helps with bounding the regression error due to observation noise, with high probability. The proof is provided in [Appendix C](#).

### 3.2. Confidence Intervals for Kernel-Based Sparse Approximate Regression

We use the representation of the approximate posterior variance given in [Theorem 3.2](#) to derive novel confidence inter-

vals for approximate regression methods. We first formalize our assumptions.

**Assumption 3.3.** The RKHS norm of the objective function  $f$  is bounded, i.e.,  $\|f\|_{\mathcal{H}_k} \leq C_k$ , for some  $C_k \in \mathbb{R}_{>0}$ .

**Assumption 3.4.** The observation noise terms  $\epsilon_i$  are independent sub-Gaussian random variables. That is, for all  $i \in \mathbb{N}$ , for all  $\eta \in \mathbb{R}$ , and for some  $R > 0$ , the moment generating function of  $\epsilon_i$  satisfies  $\mathbb{E}[\exp(\eta\epsilon_i)] \leq \exp(\frac{\eta^2 R^2}{2})$ .

[Assumption 3.4](#) can be relaxed from sub-Gaussian noise to light-tailed noise (only requiring that the moment generating function exists), at the price of slightly increasing the confidence width. This variation is provided in [Appendix E](#).

The representation of the sparse approximate posterior variance given in [Theorem 3.2](#) holds for any set of inducing points. In order for  $\bar{\mu}_n$  to be a good prediction of  $f$ , however, the set of inducing points should be selected appropriately. Interestingly, we will see that the requirements for the set of inducing points can be concisely captured by the spectral norm of the approximation error in the kernel matrix. Let  $\lambda_{\max}$  denote the maximum eigenvalue of  $k_{X_n, X_n} - \kappa_{X_n, X_n}$ . This quantity appears in our confidence intervals, and turns out to be important for our analysis in the subsequent sections on regression and optimization. A key finding from our results is the following: *A good method for choosing inducing points is one that ensures an upper bound on  $\lambda_{\max}$ .*

We now present our confidence intervals based on sparse approximate statistics. The proof is given in [Appendix D](#).

**Theorem 3.5.** *Recall the definitions of  $\bar{\mu}_n$  and  $\bar{\sigma}_n$  given in [Section 2.2](#). Assume that  $X_n$  and  $Z_m$  are chosen independent of the observation noise  $E_n$ . Under [Assumptions 3.3](#) and [3.4](#), the following equations each hold with probability at least  $1 - \delta$  for any fixed  $x \in \mathcal{X}$ :*

$$f(x) \leq \bar{\mu}_n(x) + \beta(\delta)\bar{\sigma}_n(x), \quad f(x) \geq \bar{\mu}_n(x) - \beta(\delta)\bar{\sigma}_n(x),$$

where  $\beta(\delta) = \left( \left( 2 + \frac{\sqrt{\lambda_{\max}}}{\tau} \right) C_k + \frac{R}{\tau} \sqrt{2 \log\left(\frac{1}{\delta}\right)} \right)$ , and where  $C_k$  and  $R$  are the parameters specified in [Assumptions 3.3](#) and [3.4](#).

If the domain  $\mathcal{X}$  is finite, then uniform confidence bounds readily follow from this result via a union bound, and  $\frac{\delta}{|\mathcal{X}|}$  can be substituted for  $\delta$ . For continuous domains, more effort is required, and we use a discretization argument. First, consider the following continuity assumption.

**Assumption 3.6.** For each  $n \in \mathbb{N}$ , there exists a discretization  $\mathbb{X}$  of  $\mathcal{X}$  such that, for any  $f \in \mathcal{H}_k$  with  $\|f\|_{\mathcal{H}_k} \leq C_k$ , we have  $f(x) - f([x]) \leq \frac{1}{n}$ , where  $[x] = \operatorname{argmin}_{x' \in \mathbb{X}} \|x' - x\|_{l^2}$  is the closest point in  $\mathbb{X}$  to  $x$ , and  $|\mathbb{X}| \leq cC_k^d n^d$ , where  $c$  is a constant independent of  $n$  and  $C_k$ .

[Assumption 3.6](#) is a mild assumption that holds for typical

kernels such as SE and Matérn with  $\nu > 1$  (Srinivas et al., 2010; Chowdhury & Gopalan, 2017; Vakili et al., 2021a).

**Corollary 3.7.** *Under the setting of Theorem 3.5, and under Assumption 3.6, the following equations each hold uniformly in  $x \in \mathcal{X}$  with probability at least  $1 - \delta$ :*

$$\begin{aligned} f(x) &\leq \bar{\mu}_n(x) + \frac{2}{n} + \tilde{\beta}_n(\delta) \left( \bar{\sigma}_n(x) + \frac{2}{\sqrt{n}} \right), \\ f(x) &\geq \bar{\mu}_n(x) - \frac{2}{n} - \tilde{\beta}_n(\delta) \left( \bar{\sigma}_n(x) + \frac{2}{\sqrt{n}} \right), \end{aligned}$$

where  $\tilde{\beta}_n(\delta) = \beta(\frac{\delta}{2\Gamma_n})$ ,  $\Gamma_n = c(C_k^{\bar{\mu}_n}(\frac{\delta}{2}))^d n^d$ , and  $C_k^{\bar{\mu}_n}(\delta) = C_k(1 + \frac{\sqrt{n}k_{\max}}{\tau}) + \frac{\sqrt{n}R}{\tau} \sqrt{2 \log(\frac{2n}{\delta})}$ ;  $C_k$ ,  $R$  and  $c$  are the constants specified in Assumptions 3.3, 3.4, and 3.6, respectively, and  $k_{\max} = \sup_{x \in \mathcal{X}} k(x, x)$ .

The proof is given in Appendix F, and involves using Assumption 3.6 to characterize the effect of discretizing on  $f$ ,  $\bar{\mu}_n$ , and  $\bar{\sigma}_n$ . For this purpose, we derive a high probability upper bound  $C_k^{\bar{\mu}_n}(\delta)$  on  $\|\bar{\mu}_n(\cdot)\|_{\mathcal{H}_k}$  (see Lemma F.1), which appears in the expression of  $\Gamma_n$ . A separate argument is also used to characterize the Lipschitz behavior of  $\bar{\sigma}_n(\cdot)$ . We then set  $\Gamma_n$  to be the size of discretization required to ensure  $\mathcal{O}(\frac{1}{\sqrt{n}})$  discretization errors; note that  $\Gamma_n$  depends on  $n$  according to  $\mathcal{O}(C_k^d n^{3d/2})$ . Effectively, the uniform confidence intervals scale as  $\tilde{\beta}_n(\delta) = \mathcal{O}((1 + \frac{\sqrt{\lambda_{\max}}}{\tau})C_k + \frac{R}{\tau} \sqrt{d \log(\frac{nC_k}{\delta})})$ .

In comparison with the exact models, Theorem 3.5 is analogous to Theorem 1 of (Vakili et al., 2021a) which proves  $1 - \delta$  confidence bounds of the form  $\mu_n(x) \pm \beta(\delta)\sigma_n(x)$  where  $\beta(\delta) = C_k + \sqrt{2 \log(\frac{1}{\delta})}$ . The confidence interval width multipliers are similar except for the scaling of  $C_k$  with a  $(2 + \frac{\sqrt{\lambda_{\max}}}{\tau})$  factor. We now investigate  $\lambda_{\max}$ .

### 3.3. Inducing Points and $\lambda_{\max}$

As discussed above, our results show that the effect of the choice of inducing points is concisely captured by  $\lambda_{\max}$ . In the literature, there exist several methods of selecting inducing points that ensure bounds on  $\lambda_{\max}$ . As discussed in the introduction, these methods typically choose the inducing points based on the ridge leverage score (RLS). Here we highlight a theoretical guarantee for recursive RLS sampling from (Musco & Musco, 2016). We note, however, that our results apply to any rule for selecting inducing points with a guaranteed bound on  $\lambda_{\max}$ .

The result depends on a quantity denoted by  $\tilde{d}_k(n)$ , which is the so called *effective dimension* of the kernel. The rough idea is that while the feature space for typical kernels is infinite-dimensional, for a finite dataset, the number of features with a significant impact on the model is finite. The

effective dimension for a given kernel and dataset is often defined as (Zhang, 2005; Valko et al., 2013)

$$\tilde{d}_k(n) = \text{tr} (K_{X_n, X_n} (K_{X_n, X_n} + \tau^2 I_n)^{-1}). \quad (9)$$

**Lemma 3.8** ((Musco & Musco, 2016), Theorem 7). *Given  $\delta \in (0, 1/32)$ , there exists a method for choosing the inducing points based on a recursive RLS (Musco & Musco, 2016, Algorithm 2), with  $\mathcal{O}(nm^2)$  time complexity, guaranteeing that with probability  $1 - \delta$ , it holds that  $m \leq 384\tilde{d}_k(n) \log(3\tilde{d}_k(n)/\delta)$  and  $\lambda_{\max} \leq \tau^2$ .*

Note that when  $\lambda_{\max} \leq \tau^2$ , the term  $(2 + \frac{\sqrt{\lambda_{\max}}}{\tau})C_k$  of  $\beta(\delta)$  in Theorem 3.5 simply becomes  $3C_k$ .

It is also useful to define a related notion of information gain. To do so, let  $F$  be a centered GP on the domain  $\mathcal{X}$  with kernel  $k$ . Information gain then refers to the mutual information  $\mathcal{I}(Y_n; F)$  between the data values  $Y_n = [y_i]_{i=1}^n$  and  $F$ . From the closed form expression of mutual information between two multivariate Gaussian distributions (Cover, 1999), it follows that  $\mathcal{I}(Y_n; F) = \frac{1}{2} \log \det (I_n + \frac{1}{\tau^2} K_{X_n, X_n})$ . We then define the data-independent and kernel-specific *maximal information gain* as follows (Srinivas et al., 2010):

$$\gamma_k(n) = \sup_{X_n \subset \mathcal{X}} \mathcal{I}(Y_n; F). \quad (10)$$

It is known that the information gain and the effective dimension are the same up to logarithmic factors. Specifically, we have  $\tilde{d}_k(n) \leq \mathcal{I}(Y_n; F)$ , and  $\mathcal{I}(Y_n; F) = \mathcal{O}(\tilde{d}_k(n) \log(n))$  (Calandriello et al., 2019).

We will state our results on the number of inducing points, regression error and optimization regret in terms of  $\gamma_k$ . For specific kernels, explicit bounds on  $\gamma_k$  are given in (Srinivas et al., 2010; Vakili et al., 2021c).

## 4. Sparse Approximate Kernel-Based Regression with Uncertainty Reduction

In this section, we study the application of our confidence intervals with sparse approximate statistics to kernel-based regression. We are interested in the uniform concentration of  $\bar{\mu}_n$  around the true data generating function  $f$  in the form of explicit bounds on  $\|f - \bar{\mu}_n\|_{L^\infty}$ , in terms of  $n$  and  $m$ . We show that, with  $m = \tilde{\mathcal{O}}(\gamma_k(n))$  and suitably-chosen data points,  $\|f - \bar{\mu}_n\|_{L^\infty}$  converges to zero with high probability. In addition, we characterize the precise convergence rate. For a more clear presentation of our analysis, we first consider the case of a finite domain  $|\mathcal{X}| < \infty$ , and then show the extension of the results to compact subsets of  $\mathbb{R}^d$ . Next, we describe the collection of a dataset based on uncertainty reduction, and then state our result on the error rate of the regression problem.

#### 4.1. Dataset

We recursively collect a dataset based on a maximal uncertainty reduction design. At each discrete time  $j = 1, 2, \dots$ , a point  $x_j$  is chosen based on the rule  $x_j = \arg \max_{x \in \mathcal{X}} \bar{\sigma}_{j-1}(x)$ , where  $\bar{\sigma}_{j-1}^2 = \bar{\sigma}_{k, \tau, Z_{m_{j-1}}^{(j-1)}, X_{j-1}}^2$  is the approximate posterior variance based on previous observation points  $X_{j-1}$ , using a set of inducing points  $Z_{m_{j-1}}^{(j-1)}$ .

For correctness, we assume that the set of inducing points at each time  $j$  is selected based on the recursive RLS method (Musco & Musco, 2016, Algorithm 2) formalized in Lemma 3.8. Let  $\delta_j$  denote the confidence parameter of Lemma 3.8 in selecting  $Z_{m_j}^{(j)}$ . Then, with probability at least  $1 - \delta_j$ ,  $m_j = \mathcal{O}(\tilde{d}_k(j) \log(\tilde{d}_k(j))/\delta_j)$  and  $\lambda_{\max}^{(j)} \leq \tau^2$ , where  $\lambda_{\max}^{(j)}$  denotes the maximum eigenvalue of  $k_{X_j, X_j} - \kappa_{X_j, X_j}$ . The set of inducing points may be selected based on any other method guaranteeing similar bounds on  $\lambda_{\max}^{(j)}$ .

Intuitively, this dataset  $\mathcal{D}_n = \{(x_j, y_j)_{j=1}^n\}$  is roughly evenly spread over the entire domain, which allows us to prove the convergence of  $\bar{\mu}_n$  to  $f$ . Note that without assumptions on the dataset, we cannot expect such a result. For example, if all data points are collected from a small region in the domain, it is not expected for the prediction error to be small at the points that are far from this region.

#### 4.2. Uniform Convergence of $\bar{\mu}_n$ to $f$

We prove that for the dataset described above,  $\bar{\mu}_n$  uniformly converges to  $f$  with high probability.

**Theorem 4.1.** *For  $\mathcal{D}_n$  described above with  $\delta_j = \frac{3\delta}{\pi^2 j^2}$ , under Assumptions 3.3 and 3.4, we have with probability at least  $1 - \delta$  that*

$$\|f - \bar{\mu}_n\|_{L^\infty} = \mathcal{O}\left(\sqrt{\frac{\gamma_k(n)}{n} \log\left(\frac{|\mathcal{X}|}{\delta}\right)}\right).$$

We emphasize that by Lemma 3.8, the number of inducing points used in computation of  $\bar{\mu}_n$  is bounded as  $m_n = \tilde{\mathcal{O}}(\gamma_k(n) \log(n^2 \gamma_k(n)/\delta))$ .

The proof of Theorem 4.1 relies on the confidence intervals proven in Theorem 3.5, and a bound on the ratio  $\bar{\sigma}_n^2(x)/\sigma_n^2(x)$  between the approximate and exact posterior variances of the surrogate GP model, which we show to be bounded by a constant depending on  $\lambda_{\max}$  and  $\tau$ . The proof, and a detailed expression of the bound including the hidden terms in  $\mathcal{O}$  notation, are provided in Appendix G.

#### 4.3. Extension to Continuous Domains

The following theorem extends the above convergence result to compact domains  $\mathcal{X} \subset \mathbb{R}^d$ .

**Theorem 4.2.** *Consider  $\mathcal{D}_n$  described in Section 4.1 with  $\delta_j = \frac{3\delta}{\pi^2 j^2}$ . Under Assumptions 3.3, 3.4 and 3.6, we have with probability at least  $1 - \delta$  that*

$$\|f - \bar{\mu}_n\|_{L^\infty} = \mathcal{O}\left(\sqrt{\frac{d\gamma_k(n)}{n} \log\left(\frac{n}{\delta}\right)}\right). \quad (11)$$

The proof of Theorem 4.2 follows the same steps as the proof of Theorem 3.5, using the uniform confidence intervals from Corollary 3.7 instead of the confidence intervals from Theorem 3.5. The proof and a detailed expression of the bound including the hidden terms in  $\mathcal{O}$  notation is provided in Appendix G.

*Remark 4.3.* Our results can be specified for the Matérn and SE kernels using the bounds on  $\gamma_k(n)$  from (Srinivas et al., 2010) and (Vakili et al., 2021c). For the SE kernel, we have  $\gamma_k(n) = \mathcal{O}((\log(n))^{d+1})$ ; thus, setting  $m = \mathcal{O}((\log(n))^{d+2})$ , we have with probability at least  $1 - \delta$  that  $\|f - \bar{\mu}_n\|_{L^\infty} = \tilde{\mathcal{O}}\left(\sqrt{\frac{d}{n} \log\left(\frac{n}{\delta}\right)}\right)$ . In the case of Matérn kernel, we have  $\gamma_k(n) = \tilde{\mathcal{O}}(n^{\frac{d}{2\nu+d}})$ ; thus, setting  $m = \tilde{\mathcal{O}}(n^{\frac{d}{2\nu+d}})$ , we have with probability at least  $1 - \delta$  that  $\|f - \bar{\mu}_n\|_{L^\infty} = \tilde{\mathcal{O}}\left(n^{\frac{\nu}{2\nu+d}} \sqrt{d \log\left(\frac{n}{\delta}\right)}\right)$ .

## 5. Kernel-Based Black-Box Optimization Using Sparse Approximation Methods

In this section, we consider the problem of optimizing a black-box function  $f$  over  $\mathcal{X}$ . In this setting, a learning algorithm sequentially selects observation points  $x_n$  indexed by a discrete time  $n = 1, 2, \dots$ , and receives the corresponding real-valued noisy observations  $f(x_n) + \epsilon_n$ , where  $\epsilon_n$  is the observation noise. The goal is to minimize the (cumulative) regret, defined as the sum of the losses compared to the maximum attainable objective, over a time horizon  $N$ :

$$\mathcal{R}(N) = \sum_{n=1}^N (f(x^*) - f(x_n)), \quad (12)$$

where  $x^* \in \arg \max_{x \in \mathcal{X}} f(x)$  is a global maximum of  $f$ .

### 5.1. Algorithm

Our S-BPE algorithm is an adaptation of BPE algorithm proposed in (Li & Scarlett, 2022), where the exact GP statistics are replaced by their sparse approximations. The algorithm proceeds in batches. Specifically, the time horizon  $N$  is partitioned into  $B$  batches indexed by  $i = 1, \dots, B$ , and the length of the  $i$ -th batch is denoted by  $N_i$ .

Similar to the previous section, we first consider the case that  $|\mathcal{X}| < \infty$ , and then show the extension of the results to compact subsets of  $\mathbb{R}^d$ . The S-BPE algorithm maintains a set  $\mathcal{X}_i$  of the potential maximizers of  $f$  (initialized to be the

full domain  $\mathcal{X}$ ). During each batch  $i$  the algorithm explores the points in  $\mathcal{X}_i$ , and at the end, reduces the size of  $\mathcal{X}_i$  by eliminating some of its elements.

For exploration, the algorithm picks the points in  $\mathcal{X}_i$  with the highest uncertainty, which is measured by the sparse approximate posterior variance. To ensure the validity of the confidence intervals, we must ignore the previous batches and compute the variance only based on the observations in the current batch. We thus introduce the notations

$$\begin{aligned}\bar{\mu}_{j-1,i}(\cdot) &= \bar{\mu}_{k,\tau,Z_{m_{j-1,i}}^{(j-1,i)},X_{j-1,i},Y_{j-1,i}}(\cdot), \\ \bar{\sigma}_{j-1,i}^2(\cdot) &= \bar{\sigma}_{k,\tau,Z_{m_{j-1,i}}^{(j-1,i)},X_{j-1,i}}^2(\cdot),\end{aligned}$$

for the approximate posterior mean and variance based on the first  $j-1$  observations  $\{X_{j-1,i}, Y_{j-1,i}\}$  in batch  $i$ .

During each batch  $i$ , the set of inducing points at each time  $j$  is selected based on the recursive RLS method (Musco & Musco, 2016, Algorithm 2) formalized in Lemma 3.8. Let  $1 - \delta_j$  denote the confidence level of Lemma 3.8 in selecting  $Z_{m_{j,i}}^{(j,i)}$ . Then, with probability at least  $1 - \delta_j$ ,  $m_{j,i} = \mathcal{O}(\bar{d}_k(j) \log(\bar{d}_k(j))/\delta_j)$  and  $\lambda_{\max}^{(j,i)} \leq \tau^2$ , where  $\lambda_{\max}^{(j,i)}$  denotes the maximum eigenvalue of  $k_{X_{j,i}, X_{j,i}} - \kappa_{X_{j,i}, X_{j,i}}$ .

The  $j$ -th point selected in batch  $i$  is specified as  $x_j = \arg \max_{x \in \mathcal{X}_i} \bar{\sigma}_{j-1,i}(x)$ .

*Remark 5.1.* The selection of points in each batch  $i$  (as well as the design of the dataset in Section 4) is based on uncertainty reduction, and only uses the sparse approximate variance. This ensures the validity of the confidence interval in Theorem 3.5, which requires the observation points and the inducing points to be selected independently from the observation noise (i.e., without using  $Y_{j-1,i}$ ).

At the end of batch  $i$ , the elements of  $\mathcal{X}_i$  that are unlikely to be the maximizers of  $f$  are removed. To do this, Theorem 3.5 is used to create the following upper and lower confidence bounds:

$$\begin{aligned}u_i(\cdot) &= \bar{\mu}_{N_i,i}(\cdot) + \beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_i,i}(\cdot), \\ l_i(\cdot) &= \bar{\mu}_{N_i,i}(\cdot) - \beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_i,i}(\cdot).\end{aligned}\quad (13)$$

Using  $l_i$  and  $u_i$ , S-BPE sets  $\mathcal{X}_{i+1} = \{x \in \mathcal{X}_i : u_i(x) \geq \max_{x' \in \mathcal{X}_i} l_i(x')\}$ . The rationale here is that conditioned on the validity of the confidence bounds, if  $u_i(x) < l_i(z)$  for some  $z \in \mathcal{X}_i$ , then  $x$  cannot be a maximizer of  $f$ . Pseudocode for S-BPE is provided in Algorithm 1 in Appendix B.

**Theorem 5.2.** *Consider the black-box optimization problem described above. In S-BPE, set  $\delta_j = \frac{3\delta}{\pi^2 B j^2}$ , and  $l_i, u_i$  according to (13). Under Assumptions 3.3 and 3.4, the regret performance of S-BPE satisfies the following with*

probability at least  $1 - \delta$ :

$$\mathcal{R}(N) = \tilde{\mathcal{O}} \left( \sqrt{N \gamma_k(N) \log \left( \frac{|\mathcal{X}|}{\delta} \right)} \right). \quad (14)$$

The proof is given in Appendix H, along with a more precise expression for the bound showing constants and logarithmic terms.

## 5.2. Extension to Continuous Domains

The extension to the continuous domains is similar to Section 4.3. For continuous domains, we replace the confidence bounds  $l_i$  and  $u_i$  in (13) with the following:

$$\begin{aligned}u_i(\cdot) &= \bar{\mu}_{N_i,i}(\cdot) + \frac{2}{N_i} + \tilde{\beta}_{N_i} \left( \frac{\delta}{4B} \right) \left( \bar{\sigma}_{N_i,i}(\cdot) + \frac{2}{\sqrt{N_i}} \right), \\ l_i(\cdot) &= \bar{\mu}_{N_i,i}(\cdot) - \frac{2}{N_i} - \tilde{\beta}_{N_i} \left( \frac{\delta}{4B} \right) \left( \bar{\sigma}_{N_i,i}(\cdot) + \frac{2}{\sqrt{N_i}} \right),\end{aligned}\quad (15)$$

which are uniformly valid as per Corollary 3.7.

**Theorem 5.3.** *Consider the black-box optimization problem described above. In S-BPE, set  $\delta_j = \frac{3\delta}{\pi^2 B j^2}$ , and  $l_i, u_i$  according to (15). Under Assumptions 3.3, 3.4 and 3.6, the regret of S-BPE satisfies the following with probability at least  $1 - \delta$ :*

$$\mathcal{R}(N) = \tilde{\mathcal{O}} \left( \sqrt{N d \gamma_k(N) \log \left( \frac{N}{\delta} \right)} \right).$$

The proof and a more precise expression for the bound are given in Appendix H.

Using Lemma 3.8, the total time complexity of S-BPE over the time horizon  $N$  is bounded by  $\tilde{\mathcal{O}}(N^2(\gamma_k(N))^2)$ , which is the same as that of BKB (Calandriello et al., 2019) and lower than that of S-GP-TS (Vakili et al., 2021d). All are significant improvements over the time complexity of standard kernel-based bandit algorithms using exact GP statistics.

We note that the time complexity  $\tilde{\mathcal{O}}(N^2(\gamma_k(N))^2)$  excludes the complexity of optimizing the acquisition function (here, sparse approximate variance). This aspect is largely orthogonal to our focus on the approximate posterior calculation, and it is similarly not considered in the most relevant works on optimization using sparse GPs (Calandriello et al., 2019; Vakili et al., 2021d).

*Remark 5.4.* In the case of the SE kernel, Theorem 5.2 implies the following regret bound:

$$\mathcal{R}(N) = \tilde{\mathcal{O}} \left( \sqrt{N d (\log(N))^{d+1} \log \left( \frac{N}{\delta} \right)} \right), \quad (16)$$

and in the case of Matérn kernels, Theorem 5.2 implies

$$\mathcal{R}(N) = \tilde{\mathcal{O}} \left( N^{\frac{\nu+d}{2\nu+d}} (\log(N))^{\frac{\nu}{2\nu+d}} \sqrt{d \log \left( \frac{N}{\delta} \right)} \right). \quad (17)$$



Comparing to the lower bounds proven in (Scarlett et al., 2017), both of these bounds are tight up to logarithmic factors in  $N$ . To the best of our knowledge, this is the first kernel-based bandit optimization algorithm based on sparse approximation methods enjoying such near-optimality guarantees. In comparison, for Matérn kernels, the regret bounds for both BKB and S-GP-TS scale at a rate  $\tilde{O}(N^{\frac{2\nu+3d}{4\nu+2d}})$ , which become super-linear (trivial) when  $\nu \leq d/2$ .

## 6. Conclusion

We have proved confidence intervals for kernel-based regression with sparse approximate models, using a novel representation of the sparse approximate posterior variance. We demonstrated the applications of the confidence intervals for regression problems, as well as kernel-based bandit optimization problems. For the latter, our regret bounds are optimal up to logarithmic factors for the cases where lower bounds on the regret are known.

**Acknowledgment.** J. Scarlett was supported by the Singapore National Research Foundation (NRF) under grant number R-252-000-A74-281.

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>.
- Alaoui, A. E. and Mahoney, M. W. Fast randomized kernel methods with statistical guarantees. *arXiv preprint arXiv:1411.0306*, 2014.
- Antonini, R. G., Kozachenko, Y., and Volodin, A. Convergence of series of dependent  $\varphi$ -subgaussian random variables. *Journal of Mathematical Analysis and Applications*, 338(2):1188–1203, 2008.
- Burt, D., Rasmussen, C. E., and Van Der Wilk, M. Rates of convergence for sparse variational Gaussian process regression. In *International Conference on Machine Learning*, pp. 862–871, 2019.
- Burt, D. R., Rasmussen, C. E., and van der Wilk, M. Convergence of sparse variational inference in Gaussian processes regression. *Journal of Machine Learning Research*, 21(131):1–63, 2020.
- Calandriello, D., Carratino, L., Lazaric, A., Valko, M., and Rosasco, L. Gaussian process optimization with adaptive sketching: Scalable and no regret. In *Conference on Learning Theory*, 2019.
- Calandriello, D., Carratino, L., Lazaric, A., Valko, M., and Rosasco, L. Scaling Gaussian process optimization by evaluating a few unique candidates multiple times. In *International Conference on Machine Learning*, 2022.
- Camilleri, R., Jamieson, K., and Katz-Samuels, J. High-dimensional experimental design and kernel bandits. In *International Conference on Machine Learning*, pp. 1227–1237. PMLR, 2021.
- Chareka, P., Chareka, O., and Kennendy, S. Locally sub-Gaussian random variable and the strong law of large numbers. *Atlantic Electronic Journal of Mathematics*, 1(1):75–81, 2006.
- Chen, Y. and Yang, Y. Fast statistical leverage score approximation in kernel ridge regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 2935–2943. PMLR, 2021.
- Chowdhury, S. R. and Gopalan, A. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pp. 844–853, 2017.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.
- Janz, D., Burt, D., and Gonzalez, J. Bandit optimisation of functions in the matern kernel RKHS. In *Proceedings of Machine Learning Research*, volume 108, pp. 2486–2495. PMLR, 26–28 Aug 2020.
- Kanagawa, M., Hennig, P., Sejdinovic, D., and Sriperumbudur, B. K. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*, 2018.
- Li, Z. and Scarlett, J. Gaussian process bandit optimization with few batches. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- Lin, J. and Cevher, V. Optimal rates of sketched-regularized algorithms for least-squares regression over Hilbert spaces. In *International Conference on Machine Learning*, pp. 3102–3111. PMLR, 2018.
- Liu, H., Ong, Y.-S., Shen, X., and Cai, J. When Gaussian process meets big data: A review of scalable GPs. *IEEE Transactions on Neural Networks and Learning Systems*, 31(11):4405–4423, 2020.

- Matthews, A. G. d. G., Hensman, J., Turner, R., and Ghahramani, Z. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. In *International Conference on Artificial Intelligence and Statistics*, pp. 231–239, 2016.
- Matthews, A. G. d. G., van der Wilk, M., Nickson, T., Fujii, K., Boukouvalas, A., León-Villagrà, P., Ghahramani, Z., and Hensman, J. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, apr 2017.
- Mockus, J., Tiesis, V., and Zilinskas, A. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2(117-129):2, 1978.
- Musco, C. and Musco, C. Recursive sampling for the Nyström method. *arXiv preprint arXiv:1605.07583*, 2016.
- Nieman, D., Szabo, B., and van Zanten, H. Contraction rates for sparse variational approximations in Gaussian process regression. *arXiv preprint arXiv:2109.10755*, 2021.
- Rasmussen, C. E. and Williams, C. K. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Rossi, S., Heinonen, M., Bonilla, E., Shen, Z., and Filippone, M. Sparse Gaussian processes revisited: Bayesian approaches to inducing-variable approximations. In *International Conference on Artificial Intelligence and Statistics*, pp. 1837–1845. PMLR, 2021.
- Rudi, A., Camoriano, R., and Rosasco, L. Less is more: Nyström computational regularization. *Advances in Neural Information Processing Systems*, 28, 2015.
- Salgia, S., Vakili, S., and Zhao, Q. A domain-shrinking based Bayesian optimization algorithm with order-optimal regret performance. *Conference on Neural Information Processing Systems*, 34, 2021.
- Scarlett, J., Bogunovic, I., and Cevher, V. Lower bounds on regret for noisy Gaussian process bandit optimization. In *Conference on Learning Theory*, pp. 1723–1742, 2017.
- Schölkopf, B., Smola, A. J., Bach, F., et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Seeger, M. W., Williams, C. K., and Lawrence, N. D. Fast forward selection to speed up sparse Gaussian process regression. In *International Workshop on Artificial Intelligence and Statistics*, pp. 254–261. PMLR, 2003.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning*, pp. 1015–1022, 2010.
- Titsias, M. K. Variational learning of inducing variables in sparse Gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Vakili, S., Bouziani, N., Jalali, S., Bernacchia, A., and Shiu, D.-s. Optimal order simple regret for Gaussian process bandits. In *Conference on Neural Information Processing Systems*, 2021a.
- Vakili, S., Bromberg, M., Garcia, J., Shiu, D.-s., and Bernacchia, A. Uniform generalization bounds for overparameterized neural networks. *arXiv preprint arXiv:2109.06099*, 2021b.
- Vakili, S., Khezeli, K., and Picheny, V. On information gain and regret bounds in Gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 82–90, 2021c.
- Vakili, S., Picheny, V., and Artemev, A. Scalable Thompson sampling using sparse Gaussian process models. In *Conference on Neural Information Processing Systems*, 2021d.
- Vakili, S., Scarlett, J., and Javidi, T. Open problem: Tight online confidence intervals for RKHS elements. In *Conference on Learning Theory*, pp. 4647–4652. PMLR, 2021e.
- Valko, M., Korda, N., Munos, R., Flaounas, I., and Cristianini, N. Finite-time analysis of kernelised contextual bandits. In *Conference on Uncertainty in Artificial Intelligence*, pp. 654–663, 2013.
- Wild, V., Kanagawa, M., and Sejdinovic, D. Connections and equivalences between the Nyström method and sparse variational Gaussian processes. *arXiv preprint arXiv:2106.01121*, 2021.
- Williams, C. and Seeger, M. Using the Nyström method to speed up kernel machines. In *Conference on Neural Information Processing Systems*, pp. 682–688, 2001.
- Zenati, H., Bietti, A., Diemert, E., Mairal, J., Martin, M., and Gaillard, P. Efficient kernelized UCB for contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 5689–5720. PMLR, 2022.
- Zhang, T. Learning bounds for kernel regression using effective data dimensionality. *Neural Computation*, 17(9):2077–2098, 2005.

In the following appendices, we provide a summary of some common notation, pseudo-code for S-BPE, the proofs of our theorems, as well as a variation of Theorem 3.5 relaxing the sub-Gaussian noise assumption. The proofs of auxiliary lemmas are provided in Appendix I.

## A. Summary of Main Notation

For the reader's convenience, some of the main notation used throughout the paper is summarized as follows:

- Inputs  $X_n = [x_1, \dots, x_n]^\top$ , observations  $Y_n = [y_1, \dots, y_n]^\top$ , noise terms  $E_n = [\epsilon_1, \epsilon_2, \dots, \epsilon_n]^\top$ , inducing points  $Z_m = [z_1, z_2, \dots, z_m]^\top$ , regularization parameter  $\tau$ .
- Exact kernel  $k(\cdot, \cdot)$ , reduced rank kernel  $\kappa(\cdot, \cdot)$  (see (6)).
- Exact posterior mean (or kernel regressor)  $\mu_{k, \tau, X_n, Y_n}(\cdot)$ , covariance  $k_{\tau, X_n}(\cdot, \cdot)$ , and variance  $\sigma_{k, \tau, X_n}^2(\cdot)$  (see (1) and (3)), or in shorthand notation,  $\mu_n$ ,  $k_n$ , and  $\sigma_n$ .
- Approximate posterior mean  $\bar{\mu}_{k, \tau, Z_m, X_n, Y_n}(\cdot)$ , covariance  $\bar{k}_{\tau, Z_m, X_n}(\cdot, \cdot)$ , and variance  $\bar{\sigma}_{k, \tau, Z_m, X_n}^2(\cdot)$  (see (4)), or in shorthand notation,  $\bar{\mu}_n$ ,  $\bar{k}_n$ , and  $\bar{\sigma}_n$ .
- Auxiliary vector-valued function  $V_n$  (see (5)) such that  $\bar{\mu}_n(\cdot) = V_n^\top(\cdot)Y_n$ .
- (Section 3) Projected function  $\Pi_\kappa[f]$  (see (8)), RKHS norm bound  $C_k$ , sub-Gaussianity parameter  $R$ , maximum eigenvalue  $\lambda_{\max}$  (of  $k_{X_n, X_n} - \kappa_{X_n, X_n}$ ), confidence interval multiplier  $\beta(\delta)$ , discretized input  $[x]$ .

## B. Pseudo-Code for S-BPE

Pseudo-code for S-BPE is provided in Algorithm 1 below.

---

### Algorithm 1 Sparse Batched Pure Exploration (S-BPE)

---

```

1: Initialization:  $k, \mathcal{X}, f, N, N_0 \leftarrow 1$ 
2: for  $i \leftarrow 1, 2, \dots$  do
3:    $N_i \leftarrow \lceil \sqrt{N \cdot N_{i-1}} \rceil$ 
4:    $\mathcal{S}_i \leftarrow \emptyset$ 
5:   for  $j \leftarrow 1, 2, \dots, N_i$  do
6:     Update  $Z_{m_{j-1}, i}^{(j-1, i)}$ 
7:     Update  $\bar{\sigma}_{j-1, i}(\cdot)$ 
8:      $x_n \leftarrow \arg \max_{x \in \mathcal{X}_i} \bar{\sigma}_{j-1, i}(x)$ 
9:      $\mathcal{S}_i \leftarrow \mathcal{S}_i \cup \{x_n\}$ 
10:     $n \leftarrow n + 1$ 
11:    if  $n > N$  then
12:      Terminate
13:    end if
14:    Collect observations for all points in  $\mathcal{S}_i$ 
15:  end for
16:  Compute  $\bar{\mu}_{N_i, i}$  and  $\bar{\sigma}_{N_i, i}$ 
17:   $\mathcal{X}_{i+1} \leftarrow \{x \in \mathcal{X}_i : u_i(x) \geq \max_{x' \in \mathcal{X}_i} l_i(x')\}$ 
18: end for

```

---

For discrete and continuous domains  $l_i$  and  $u_i$  are set according to (13) and (15), respectively.

## C. Proof of Theorem 3.2 (Decomposition of the Approximate Posterior Variance)

From the expression for the approximate posterior covariance in (4), we have

$$\bar{\sigma}_n^2(\cdot, \cdot) = \underbrace{k(\cdot, \cdot) - k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} k_{Z_m, Z_m}(\cdot)}_{\text{Term 1}} + \underbrace{k_{Z_m}^\top(\cdot) \left( k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{Z_m}(\cdot)}_{\text{Term 2}}.$$

Term 1 is equivalent to the posterior variance of a noise-free GP model with kernel  $k$  conditioned on  $Z_m$ :

$$\text{Term 1} = \sigma_{k,0,Z_m}^2(\cdot). \quad (18)$$

The following lemma shows that Term 2 is equivalent to the posterior variance of a GP model with kernel  $\kappa$  conditioned on  $X_n$ .

**Lemma C.1.** *For any kernel  $k$  and any  $Z_m$ ,  $X_n$ , and  $Y_n$ , we have*

$$\text{Term 2} = \sigma_{\kappa,\tau,X_n}^2(\cdot),$$

where  $\kappa(\cdot, \cdot') = k_{Z_m}(\cdot)k_{Z_m,Z_m}^{-1}k_{Z_m}(\cdot')$ .

In addition, we have the following expression for  $V_n(\cdot)$  based on  $\kappa$ .

**Lemma C.2.** *For any kernel  $k$  and any  $Z_m$ ,  $X_n$ , and  $Y_n$ , we have*

$$V_n^\top(\cdot) = \kappa_{X_n}^\top(\cdot)(\tau^2 I_n + \kappa_{X_n,X_n})^{-1}, \quad (19)$$

where  $\kappa(\cdot, \cdot') = k_{Z_m}(\cdot)k_{Z_m,Z_m}^{-1}k_{Z_m}(\cdot')$ .

Combining (1) and (4) with Lemma C.2, we find that  $\bar{\mu}_{k,\tau,Z_m,X_n,\Pi_\kappa[f]_{X_n}} = \mu_{\kappa,\tau,X_n,\Pi_\kappa[f]_{X_n}}$ , where  $\Pi_\kappa[f]_{X_n} = [\Pi_\kappa[f](x_1), \dots, \Pi_\kappa[f](x_n)]^\top$  denotes the vector of values of the projection of  $f$  onto  $\mathcal{H}_\kappa$ , at the data points.

Then, Theorem 3.2 follows from the following two lemmas on the posterior variance of the exact GP models.

**Lemma C.3** (e.g., see, (Kanagawa et al., 2018), Proposition 3.10). *For any positive definite kernel  $k$  on  $\mathcal{X}$  and  $Z_m \in \mathcal{X}^m$ , we have*

$$\sigma_{k,0,Z_m}(\cdot) = \sup_{g:\|g\|_k \leq 1} |g(\cdot) - \mu_{k,0,Z_m,g_{Z_m}}(\cdot)|.$$

**Lemma C.4** ((Vakili et al., 2021a), Proposition 1). *For any positive definite kernel  $k$  on  $\mathcal{X}$  and  $X_n \in \mathcal{X}^n$ , we have*

$$\sigma_{k,\tau,X_n}^2(\cdot) = \sup_{g:\|g\|_{\mathcal{H}_k} \leq 1} (g(\cdot) - \mu_{k,\tau,X_n,g_{X_n}}(\cdot))^2 + \tau^2 \left\| \kappa_{X_n}(\cdot)(\tau^2 I_n + \kappa_{X_n,X_n})^{-1} \right\|_{l^2}^2.$$

Combining Lemmas C.1 and C.3, we have

$$\begin{aligned} \text{Term 1} &= \sigma_{k,0,Z_m}^2(\cdot) \\ &= \sup_{g:\|g\|_{\mathcal{H}_k} \leq 1} (g(\cdot) - \mu_{k,0,Z_m,g_{Z_m}}(\cdot))^2 \\ &= \sup_{g:\|g\|_{\mathcal{H}_k} \leq 1} (g(\cdot) - \Pi_\kappa[g](\cdot))^2, \end{aligned} \quad (20)$$

where the last line holds by Lemma 3.1.

In addition, using Lemmas C.1, C.4 and C.2 in that order, we obtain the following:

$$\begin{aligned} \text{Term 2} &= \sigma_{\kappa,\tau,X_n}^2(\cdot) \\ &= \sup_{g:\|g\|_{\mathcal{H}_\kappa} \leq 1} (g(\cdot) - \mu_{\kappa,\tau,X_n,g_{X_n}}(\cdot))^2 + \tau^2 \left\| \kappa_{X_n}(\cdot)(\tau^2 I_n + \kappa_{X_n,X_n})^{-1} \right\|_{l^2}^2 \\ &= \sup_{g:\|g\|_{\mathcal{H}_\kappa} \leq 1} (g(\cdot) - V_n^\top(\cdot)g_{X_n})^2 + \tau^2 \left\| V_n(\cdot) \right\|_{l^2}^2. \end{aligned} \quad (21)$$

Putting (20) and (21) together, the proof is complete.

## D. Proof of Theorem 3.5 (Confidence Interval)

Recall the definitions of  $V_n^\top(\cdot)$  and  $\Pi_\kappa[f]$  in (5) and (8). We can decompose the prediction error  $f(x) - \bar{\mu}_n(x)$  as follows:

$$\begin{aligned}
 f(x) - \bar{\mu}_n(x) &= f(x) - V_n^\top(x)Y_n \\
 &= f(x) - V_n^\top(x)f_{X_n} - V_n^\top(x)E_n \\
 &= \underbrace{f(x) - \Pi_\kappa[f](x)}_{\text{Term 1}} + \underbrace{\Pi_\kappa[f](x) - V_n^\top(x)\Pi_\kappa[f]_{X_n}}_{\text{Term 2}} - \underbrace{V_n^\top(x)(f_{X_n} - \Pi_\kappa[f]_{X_n})}_{\text{Term 3}} - \underbrace{V_n^\top(x)E_n}_{\text{Term 4}}. \quad (22)
 \end{aligned}$$

The first term corresponds to the error due to projection of  $f$  onto the approximate RKHS  $\mathcal{H}_\kappa$ . The second term corresponds to the prediction error from noise-free observations within  $\mathcal{H}_\kappa$ . The third term corresponds to the error due to projection of observations onto  $\mathcal{H}_\kappa$ , which can be interpreted as an observation noise. Finally, the fourth term corresponds to the effect of noise. We bound each term based on the expression for the approximate posterior variance given in Theorem 3.2.

**Term 1:** The first term is bounded as follows.

$$\begin{aligned}
 |f(x) - \Pi_\kappa[f](x)| &\leq \|f\|_{\mathcal{H}_\kappa} \bar{\sigma}_n(x) \\
 &\leq C_k \bar{\sigma}_n(x),
 \end{aligned}$$

where the first line holds by Theorem 3.2 (after lower bounding the second and third terms therein by zero), and by the fact that rescaling  $f$  by its RKHS norm creates a function with RKHS norm exactly one. By Lemma 3.1, this rescaling amounts to the same rescaling of  $\Pi_\kappa[f]$ . The second line follows directly from the definition of  $C_k$ .

**Term 2:** The second term is bounded as follows:

$$\begin{aligned}
 |\Pi_\kappa[f](x) - V_n^\top(x)\Pi_\kappa[f]_{X_n}| &\leq \|\Pi_\kappa[f]\|_{\mathcal{H}_\kappa} \bar{\sigma}_n(x) \\
 &\leq C_k \bar{\sigma}_n(x),
 \end{aligned}$$

where the first line holds by Theorem 3.2 (again, rescaling  $\Pi_\kappa[f]$  by its RKHS norm creates a function with RKHS norm exactly one, and this rescaling amounts to the same rescaling of  $V_n^\top(x)\Pi_\kappa[f]_{X_n}$ ). The second line comes from the definition of  $C_k$ , and the fact that  $\|\Pi_\kappa[f]\|_{\mathcal{H}_\kappa} = \|\Pi_\kappa[f]\|_{\mathcal{H}_k} \leq \|f\|_{\mathcal{H}_k}$ , by definition of  $\Pi_\kappa[f]$ .

**Term 3:** To bound the third term, we use the following known result.

**Lemma D.1** (e.g., see (Kanagawa et al., 2018), Lemma 3.9). *For any kernel  $k$  on  $\mathcal{X}$ ,  $x_1, x_2, \dots, x_n \in \mathcal{X}$  and  $c_1, c_2, \dots, c_n \in \mathbb{R}$ , the following holds:*

$$\sup_{f: \|f\|_{\mathcal{H}_k} \leq 1} \sum_{i=1}^n c_i f(x_i) = \left\| \sum_{i=1}^n c_i k(\cdot, x_i) \right\|_{\mathcal{H}_k}. \quad (23)$$

Define  $\kappa^\perp = k - \kappa$ . Using Lemma D.1, we can write

$$\begin{aligned}
 \sup_{\|g\|_{\mathcal{H}_{\kappa^\perp}} \leq 1} (V_n^\top(x)g_{X_n})^2 &= \|V_n^\top(x)\kappa_{X_n}^\perp(\cdot)\|_{\mathcal{H}_{\kappa^\perp}}^2 \\
 &= \langle V_n^\top(x)\kappa_{X_n}^\perp(\cdot), V_n^\top(x)\kappa_{X_n}^\perp(\cdot) \rangle_{\mathcal{H}_{\kappa^\perp}} \\
 &= V_n^\top(x)\kappa_{X_n, X_n}^\perp V_n(x) \\
 &= V_n^\top(x)(k_{X_n, X_n} - \kappa_{X_n, X_n})V_n(x) \\
 &\leq \lambda_{\max} \|V_n(x)\|_l^2 \\
 &\leq \frac{\lambda_{\max} \bar{\sigma}_n^2(x)}{\tau^2}, \quad (24)
 \end{aligned}$$

where the third equality uses the reproducing property, the first inequality comes from the definition of  $\lambda_{\max}$ , and the second inequality holds by Theorem 3.2.

With  $\Pi_{\kappa}^{\perp}[f] = f - \Pi_{\kappa}[f]$ , and rescaling by  $\|\Pi_{\kappa}^{\perp}[f]\|_{\mathcal{H}_{\kappa}^{\perp}}$ , Equation (24) implies

$$\begin{aligned} V_n^{\top}(x)(f_{X_n} - \Pi_{\kappa}[f]_{X_n}) &\leq \frac{\|\Pi_{\kappa}^{\perp}[f]\|_{\mathcal{H}_{\kappa}^{\perp}} \sqrt{\lambda_{\max}} \bar{\sigma}_n(x)}{\tau} \\ &\leq \frac{C_k \sqrt{\lambda_{\max}} \bar{\sigma}_n(x)}{\tau}, \end{aligned}$$

where the second inequality uses  $\|\Pi_{\kappa}^{\perp}[f]\|_{\mathcal{H}_{\kappa}^{\perp}} = \|\Pi_{\kappa}^{\perp}[f]\|_{\mathcal{H}_{\kappa}} \leq \|f\|_{\mathcal{H}_{\kappa}}$ , i.e., when decomposing  $f$  into the orthogonal components  $\Pi_{\kappa}[f]$  and  $\Pi_{\kappa}^{\perp}[f]$ , their norms cannot exceed that of  $f$ .

**Term 4:** The fourth term can be bounded with high probability using Chernoff-Hoeffding inequality, noticing that  $V_n^{\top}(x)E_n$  is a sub-Gaussian random variable with parameter  $R\|V_n(x)\|_{l^2}$ , which is bounded by  $\frac{R}{\tau}\bar{\sigma}_n(x)$  based on Theorem 3.2. This is formalized in the following lemma.

**Lemma D.2.** *Under Assumption 3.4, when the vector  $V_n(x)$  is independent of  $E_n$ , the following equations each hold with probability at least  $1 - \delta$  for any fixed  $x \in \mathcal{X}$ :*

$$\begin{aligned} V_n^{\top}(x)E_n &\leq \frac{R\bar{\sigma}_n(x)}{\tau} \sqrt{2 \log\left(\frac{1}{\delta}\right)}, \\ V_n^{\top}(x)E_n &\geq -\frac{R\bar{\sigma}_n(x)}{\tau} \sqrt{2 \log\left(\frac{1}{\delta}\right)}. \end{aligned}$$

Combining the bounds on the four terms completes the proof of Theorem 3.5.

## E. Relaxing the Sub-Gaussian Noise Assumption

Assumption 3.4, required in the analysis of Theorem 3.5, can be relaxed from sub-Gaussian noise to light-tailed noise, at the price of slightly increasing the confidence interval width. Recall the definition of light-tailed distributions: A random variable  $X$  is called light-tailed if its moment generating function  $M(h)$  exists; specifically, there exists  $h_0 > 0$  such that for all  $|h| \leq h_0$ ,  $M(h) < \infty$ .

For a zero mean light-tailed random variable  $X$ , we have (Chareka et al., 2006)

$$M(h) \leq \exp(\xi_0 h^2/2), \quad \forall |h| \leq h_0, \quad \xi_0 = \sup\{M^{(2)}(h), |h| \leq h_0\}, \quad (25)$$

where  $M^{(2)}(\cdot)$  denotes the second derivative of  $M(\cdot)$ , and  $h_0$  is the parameter specified above. We observe that the upper bound in (25) is the moment generating function of a zero mean Gaussian random variable with variance  $\xi_0$ . Thus, light-tailed distributions can also be considered locally sub-Gaussian distributions. We formalize the light-tailed noise assumption as follows.

**Assumption E.1 (Light-Tailed Noise).** The noise terms  $\epsilon_i$  are zero mean independent random variables over  $i$ , and there is a constant  $\xi_0 > 0$  such that  $\forall h \leq h_0, \forall n \in \mathbb{N}, \mathbb{E}[e^{h\epsilon_n}] \leq \exp(\frac{h^2 \xi_0}{2})$ .

The following is a variation of Theorem 3.5 with Assumption E.1 in place of Assumption 3.4.

**Theorem E.2.** *Recall the definitions of  $\bar{\mu}_n$  and  $\bar{\sigma}_n$  given in Section 2.2. Assume  $X_n$  and  $Z_m$  are chosen independent of the observation noise  $E_n$ . Under Assumptions 3.3 and E.1, the following equations each hold with probability at least  $1 - \delta$  for any fixed  $x \in \mathcal{X}$ :*

$$\begin{aligned} f(x) &\leq \bar{\mu}_n(x) + \beta(\delta)\bar{\sigma}_n(x), \\ f(x) &\geq \bar{\mu}_n(x) - \beta(\delta)\bar{\sigma}_n(x), \end{aligned}$$

where

$$\beta(\delta) = \left( \left( 2 + \frac{\sqrt{\lambda_{\max}}}{\tau} \right) C_k + \frac{1}{\tau} \sqrt{2 \max\left\{ \xi_0, \frac{2 \log(\frac{1}{\delta})}{h_0^2} \right\} \log\left(\frac{1}{\delta}\right)} \right),$$

and  $C_k, \xi_0$  and  $h_0$  are the parameters specified in Assumption 3.3 and Assumption E.1.

*Proof.* The proof follows the same steps as in the proof of Theorem 3.5. The regression error  $f(x) - \mu_n(x)$  is expanded as Terms 1 to 4 in Equation (22). The first three terms are bounded exactly as before. The only difference is in bounding Term 4,  $V_n^\top(x)E_n$ , which captures the observation noise. The modified is given in the following lemma.

**Lemma E.3.** *Under Assumption E.1, when the vector  $V_n(x)$  is independent of  $E_n$ , the following equations each hold with probability at least  $1 - \delta$  for any fixed  $x \in \mathcal{X}$ :*

$$\begin{aligned} V_n^\top(x)E_n &\leq \frac{\bar{\sigma}_n(x)}{\tau} \sqrt{2 \max \left\{ \xi_0, \frac{2 \log(\frac{1}{\delta})}{h_0^2} \right\} \log \left( \frac{1}{\delta} \right)}, \\ V_n^\top(x)E_n &\geq -\frac{\bar{\sigma}_n(x)}{\tau} \sqrt{2 \max \left\{ \xi_0, \frac{2 \log(\frac{1}{\delta})}{h_0^2} \right\} \log \left( \frac{1}{\delta} \right)}. \end{aligned}$$

Combining Lemma E.3 with the bound on the first three terms in the expression of  $f(x) - \bar{\mu}_n(x)$ , the proof of Theorem E.2 is complete.  $\square$

## F. Proof of Corollary 3.7 (Uniform Confidence Intervals on Continuous Domains)

To extend the confidence interval given in Theorem 3.5 to hold uniformly on  $\mathcal{X}$ , we use a discretization argument. For this purpose, we apply Assumption 3.6 to  $f$  and  $\bar{\mu}_n$ , and also use Assumption 3.6 to bound the discretization error in  $\bar{\sigma}_n$ .

In the following lemma, we establish a high-probability bound on  $\|\bar{\mu}_n\|_{\mathcal{H}_k}$ .

**Lemma F.1.** *Under Assumptions 3.3 and 3.4, the RKHS norm of  $\bar{\mu}_n = V_k^\top(x)Y_n$  satisfies the following with probability at least  $1 - \delta$ :*

$$\|\bar{\mu}_n\|_{\mathcal{H}_k} \leq C_k \left( 1 + \frac{\sqrt{n}k_{\max}}{\tau} \right) + \frac{\sqrt{n}R}{\tau} \sqrt{2 \log \left( \frac{2n}{\delta} \right)}.$$

Let  $C_k^{\bar{\mu}_n}(\delta) = C_k \left( 1 + \frac{\sqrt{n}k_{\max}}{\tau} \right) + \frac{\sqrt{n}R}{\tau} \sqrt{2 \log \left( \frac{2n}{\delta} \right)}$  denote the  $1 - \delta$  upper confidence bound on  $\|\bar{\mu}_n\|_{\mathcal{H}_k}$ , and define  $\mathcal{E} = \{\|\bar{\mu}_n\|_{\mathcal{H}_k} \leq C_k^{\bar{\mu}_n}(\frac{\delta}{2})\}$ . By Lemma F.1, we have  $\Pr[\mathcal{E}] \geq 1 - \frac{\delta}{2}$ . Let  $\mathbb{X}$  be the discretization of  $\mathcal{X}$  specified in Assumption 3.6 with RKHS norm bound  $C_k^{\bar{\mu}_n}(\frac{\delta}{2})$ . That is, for any  $g \in \mathcal{H}_k$  with  $\|g\|_{\mathcal{H}_k} \leq C_k^{\bar{\mu}_n}(\frac{\delta}{2})$ , we have  $g(x) - g([x]) \leq \frac{1}{n}$ , where  $[x] = \operatorname{argmin}_{x' \in \mathbb{X}} \|x' - x\|_{l_2}$  is the closest point in  $\mathbb{X}$  to  $x$ , and  $|\mathbb{X}| \leq \Gamma_n$ , where  $\Gamma_n = c \left( C_k^{\bar{\mu}_n}(\frac{\delta}{2}) \right)^d n^d$ .

Applying Assumption 3.6 to  $f$  and  $\bar{\mu}_n$  with this discretization, it holds for all  $x \in \mathcal{X}$  that

$$|f(x) - f([x])| \leq \frac{1}{n}.$$

In addition, under  $\mathcal{E}$ , for all  $x \in \mathcal{X}$

$$|\bar{\mu}_n(x) - \bar{\mu}_n([x])| \leq \frac{1}{n}.$$

Furthermore, we have the following lemma, which can roughly be viewed as a Lipschitz continuity property for  $\bar{\sigma}_n(\cdot)$ .

**Lemma F.2.** *Under Assumption 3.6, with the discretization  $\mathbb{X}$  described above, it holds for all  $x \in \mathcal{X}$  that*

$$\bar{\sigma}_n(x) - \bar{\sigma}_n([x]) \leq \frac{2}{\sqrt{n}}.$$

Using this discretization, we also define the event  $\mathcal{E}_2 = \{\bar{\mu}_n(x) - \beta(\frac{\delta}{2\Gamma_n})\bar{\sigma}_n(x) \leq f(x), \forall x \in \mathbb{X}\}$ . By Theorem 3.5 and a probability union bound over  $x \in \mathbb{X}$ , we have  $\Pr[\mathcal{E}_2] \geq 1 - \frac{\delta}{2}$ .

Accounting for the discretization errors given above, we have under  $\mathcal{E}_1 \cap \mathcal{E}_2$  (which holds with probability at least  $1 - \delta$ ) that the following lower bound holds for all  $x \in \mathcal{X}$ :

$$f(x) \geq \bar{\mu}_n(x) - \frac{2}{n} - \tilde{\beta}_n(\delta) \left( \bar{\sigma}_n(x) + \frac{2}{\sqrt{n}} \right),$$

where  $\tilde{\beta}_n(\delta) = \beta(\frac{\delta}{2\Gamma_n})$ . The uniform upper bound holds similarly.

## G. Proofs of Theorems 4.1 and 4.2 (Error Bounds for Regression)

We will prove these theorems using Theorem 3.5, a bound on the ratio between the approximate and true posterior variances, and a bound on the sum of conditional exact variances based on  $\gamma_k(n)$ . The following lemma states that the ratio between posterior variances can be bounded by a constant depending on  $\lambda_{\max}$ .

**Lemma G.1.** *For  $\sigma_n^2$  and  $\bar{\sigma}_n^2$  defined in Section 2, we have for all  $x \in \mathcal{X}$  that*

$$\begin{aligned} \sigma_n^2(x) &\leq \bar{\sigma}_n^2(x), \\ \bar{\sigma}_n^2(x) &\leq \left(1 + \frac{\lambda_{\max}}{\tau^2}\right) \sigma_n^2(x). \end{aligned}$$

We have for all  $j \leq n$  and  $x \in \mathcal{X}$  that

$$\begin{aligned} \bar{\sigma}_n(x) &\leq \left(\sqrt{1 + \frac{\lambda_{\max}}{\tau^2}}\right) \sigma_n(x) \\ &\leq \left(\sqrt{1 + \frac{\lambda_{\max}}{\tau^2}}\right) \sigma_{j-1}(x) \\ &\leq \left(\sqrt{1 + \frac{\lambda_{\max}}{\tau^2}}\right) \bar{\sigma}_{j-1}(x) \\ &\leq \left(\sqrt{1 + \frac{\lambda_{\max}}{\tau^2}}\right) \bar{\sigma}_{j-1}(x_j) \\ &\leq \left(1 + \frac{\lambda_{\max}}{\tau^2}\right) \sigma_{j-1}(x_j), \end{aligned}$$

where the first, third, and final inequalities all follow from Lemma G.1, the second inequality holds since conditioning on a larger set reduces the posterior variance of an exact GP model (which follows from positive-definiteness of the kernel matrix), and the fourth inequality follows from the data selection rule  $x_j = \arg \max_{x \in \mathcal{X}} \bar{\sigma}_{j-1}(x)$ .

Averaging both sides over  $j = 1, \dots, n$ , we have

$$\bar{\sigma}_n^2(x) \leq \frac{1}{n} \left(1 + \frac{\lambda_{\max}}{\tau^2}\right)^2 \sum_{j=1}^n \sigma_{j-1}^2(x_j). \quad (26)$$

It is well known from (Srinivas et al., 2010) that the quantity  $\sum_{j=1}^n \sigma_{j-1}^2(x_j)$  can be bounded in terms of the mutual information given in (3.3) for exact GP models:

$$\sum_{j=1}^n \sigma_{j-1}^2(x_j) \leq \frac{2}{\log(1 + 1/\tau^2)} \mathcal{I}(Y_n, F). \quad (27)$$

Using this bound and the definition of  $\gamma_k(n)$  in (10), we obtain

$$\bar{\sigma}_n^2(x) \leq \frac{2(1 + \frac{\lambda_{\max}}{\tau^2})^2 \gamma_k(n)}{\log(1 + 1/\tau^2) n}. \quad (28)$$

Define the event  $\mathcal{E}_1 = \{\bar{\mu}_n(x) - \beta(\frac{\delta}{4|\mathcal{X}|}) \bar{\sigma}_n(x) \leq f(x) \leq \bar{\mu}_n(x) + \beta(\frac{\delta}{4|\mathcal{X}|}) \bar{\sigma}_n(x), \forall x \in \mathcal{X}\}$ . By Theorem 3.5 and a probability union bound over  $x$ , we have  $\Pr[\mathcal{E}_1] \geq 1 - \frac{\delta}{2}$ . In addition, let  $\mathcal{E}_2 = \{\lambda_{\max}^{(j)} \leq \tau^2, \forall j \leq n\}$ . By Lemma 3.8, and



a probability union bound over  $j$ , we have  $\Pr[\mathcal{E}_2] \geq 1 - \frac{\delta}{2}$ . We thus condition the rest of the proof on  $\mathcal{E}_1 \cap \mathcal{E}_2$ , which holds true with probability at least  $1 - \delta$ .

We have for all  $x \in \mathcal{X}$  that

$$\begin{aligned} |f(x) - \bar{\mu}_n(x)| &\leq \beta\left(\frac{\delta}{4|\mathcal{X}|}\right)\bar{\sigma}_n(x) \\ &\leq \left(3C_k + \frac{R}{\tau}\sqrt{2\log\left(\frac{4|\mathcal{X}|}{\delta}\right)}\right)\bar{\sigma}_n(x) \\ &\leq \left(3C_k + \frac{R}{\tau}\sqrt{2\log\left(\frac{4|\mathcal{X}|}{\delta}\right)}\right)\sqrt{\frac{8\gamma_k(n)}{\log(1+1/\tau^2)n}}, \end{aligned} \quad (29)$$

where the first and second lines hold by  $\mathcal{E}_1$  and  $\mathcal{E}_2$ , and the third line comes from substituting the bound on  $\bar{\sigma}_n(x)$ . Thus, we have with probability at least  $1 - \delta$  that

$$\|f - \bar{\mu}_n\|_{L^\infty} \leq \left(3C_k + \frac{R}{\tau}\sqrt{2\log\left(\frac{4|\mathcal{X}|}{\delta}\right)}\right)\sqrt{\frac{8\gamma_k(n)}{\log(1+1/\tau^2)n}}, \quad (30)$$

which completes the proof of Theorem 4.1.

**Extension to continuous domains.** Theorem 4.2 is proved following the same steps as in the proof of Theorem 3.5. Let us define the event

$$\mathcal{E}'_1 = \left\{ \bar{\mu}_n(x) - \frac{2}{n} - \tilde{\beta}_n\left(\frac{\delta}{4}\right)\left(\bar{\sigma}_n(x) + \frac{2}{\sqrt{n}}\right) \leq f(x) \leq \bar{\mu}_n(x) + \frac{2}{n} + \tilde{\beta}_n\left(\frac{\delta}{4}\right)\left(\bar{\sigma}_n(x) + \frac{2}{\sqrt{n}}\right), \forall x \in \mathcal{X} \right\}.$$

By Corollary 3.7,  $\Pr[\mathcal{E}'_1] \geq 1 - \frac{\delta}{2}$ . Recalling also that  $\mathcal{E}_2 = \{\lambda_{\max}^{(j)} \leq \tau^2, \forall j \leq n\}$  and  $\Pr[\mathcal{E}_2] \geq 1 - \frac{\delta}{2}$ , we can condition on  $\mathcal{E}'_1 \cap \mathcal{E}_2$ , which holds with probability at least  $1 - \delta$ .

We have for all  $x \in \mathcal{X}$  that

$$\begin{aligned} |f(x) - \bar{\mu}_n(x)| &\leq \tilde{\beta}_n\left(\frac{\delta}{4}\right)\left(\bar{\sigma}_n(x) + \frac{2}{\sqrt{n}}\right) + \frac{2}{n} \\ &\leq \left(3C_k + \frac{R}{\tau}\sqrt{2\log\left(\frac{8\Gamma_n}{\delta}\right)}\right)\left(\bar{\sigma}_n(x) + \frac{2}{\sqrt{n}}\right) + \frac{2}{n} \\ &\leq \left(3C_k + \frac{R}{\tau}\sqrt{2\log\left(\frac{8\Gamma_n}{\delta}\right)}\right)\left(\sqrt{\frac{8\gamma_k(n)}{\log(1+1/\tau^2)n}} + \frac{2}{\sqrt{n}}\right) + \frac{2}{n}, \end{aligned} \quad (31)$$

where the first and second lines hold by  $\mathcal{E}_1$  and  $\mathcal{E}_2$  (as well as  $\tilde{\beta}_n(\delta) = \beta(\frac{\delta}{2\Gamma_n})$ ), and the third line comes from substituting the bound on  $\bar{\sigma}_n(x)$  from (28). Recall that  $\Gamma_n = c(C_k^{\bar{\mu}_n}(\frac{\delta}{2}))^d n^d$  and  $C_k^{\bar{\mu}_n}(\delta) = C_k(1 + \frac{\sqrt{nk_{\max}}}{\tau}) + \frac{\sqrt{nR}}{\tau}\sqrt{2\log(\frac{2n}{\delta})}$ . This gives  $\Gamma_n = \tilde{O}(C_k^d n^{3d/2})$ , and we arrive at

$$\|f - \bar{\mu}_n\|_{L^\infty} = \mathcal{O}\left(\sqrt{\frac{d\gamma_k(n)}{n}\log\left(\frac{n}{\delta}\right)}\right), \quad (32)$$

which completes the proof.

## H. Proofs of Theorems 5.2 and 5.3 (Regret Bounds for Optimization)

Let  $t_i = \sum_{j=1}^i N_j$  denote the time at the end of batch  $i$ , and  $\mathcal{R}_i = \sum_{j=t_{i-1}+1}^{t_i} (f(x^*) - f(x_j))$  denote the regret incurred in batch  $i$ . We thus have  $\mathcal{R}(N) = \sum_{i=1}^B \mathcal{R}_i$ .

Let  $\mathcal{E}_1 = \{\forall i, \forall x \in \mathcal{X}, l_i(x) \leq f(x) \leq u_i(x)\}$ , for  $l_i, u_i$  given in (13), be the event that all the confidence intervals at the end of all batches are satisfied. By Theorem 3.5 and a probability union bound, we have  $\Pr[\mathcal{E}_1] \geq 1 - \frac{\delta}{2}$ . In addition, let  $\mathcal{E}_2 = \{\forall i, j, \lambda_{\max}^{(i,j)} \leq \tau^2\}$ . By Lemma 3.8, and a probability union bound over  $i$  and  $j$ , we have  $\Pr[\mathcal{E}_2] \geq 1 - \frac{\delta}{2}$ . We thus condition the rest of the proof on  $\mathcal{E}_1 \cap \mathcal{E}_2$ , which holds with probability at least  $1 - \delta$ .

Under  $\mathcal{E}_1$ , we have for all  $x \in \mathcal{X}$  that

$$\begin{aligned} u_i(x^*) &\geq f(x^*) \\ &\geq f(x) \\ &\geq l_i(x). \end{aligned}$$

Thus,  $u_i(x^*) \geq \max_{x \in \mathcal{X}_i} l_i(x)$  for all  $i$ , and  $x^*$  will not be eliminated by S-BPE; i.e.,  $x^* \in \mathcal{X}_i$  for all  $i$ .

The maximum regret in the first batch is bounded as  $\mathcal{R}_1 \leq \sqrt{N}(\max_{x \in \mathcal{X}}(f(x^*) - f(x)))$ . Moreover, by the reproducing property,  $f(x) = \langle f(\cdot), k(x, \cdot) \rangle \leq \|f\|_{\mathcal{H}_k} k(x, x) \leq C_k k_{\max}$ . Thus, we have  $\mathcal{R}_1 \leq C_k k_{\max} \sqrt{N}$ .

For  $n > t_{i-1}$  and  $i \geq 2$ , the regret can be bounded using a similar approach as in the analysis of Theorem 4.1 (see also (Li & Scarlett, 2022)):

$$\begin{aligned} f(x^*) - f(x_n) &\leq u_{i-1}(x^*) - l_{i-1}(x_n) \\ &\leq l_{i-1}(x^*) + 2\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_{i-1}, i-1}(x^*) - u_{i-1}(x_n) + 2\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_{i-1}, i-1}(x_n) \\ &\leq 2\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_{i-1}, i-1}(x^*) + 2\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \bar{\sigma}_{N_{i-1}, i-1}(x_n) \\ &\leq 4\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \sqrt{\frac{8\gamma_k(N_{i-1})}{\log(1 + 1/\tau^2)N_{i-1}}}, \end{aligned}$$

where the third line follows from the elimination rule of S-BPE, and the last line is proven following the same steps as in the proof of Equation (28).

Thus,

$$\begin{aligned} \mathcal{R}_i &= \sum_{n=t_{i-1}+1}^{t_i} (f(x^*) - f(x_n)) \\ &\leq 4N_i \beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) \sqrt{\frac{8\gamma_k(N_{i-1})}{\log(1 + 1/\tau^2)N_{i-1}}} \\ &\leq 4\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) (\sqrt{N} + 1) \sqrt{\frac{8\gamma_k(N_{i-1})}{\log(1 + 1/\tau^2)}} \\ &\leq 4\beta \left( \frac{\delta}{4B|\mathcal{X}|} \right) (\sqrt{N} + 1) \sqrt{\frac{8\gamma_k(N)}{\log(1 + 1/\tau^2)}}, \end{aligned}$$

where the second inequality follows from the choice of  $N_i$  in Line 1 of Algorithm 1.

The definition of  $N_i$  ensures that the number of batches satisfies  $B \leq \lceil \log \log(N) \rceil + 1$  (Li & Scarlett, 2022, Proposition 1).

As a result, we have

$$\begin{aligned}
 \mathcal{R}(N) &= \sum_{i=1}^B \mathcal{R}_i \\
 &\leq 4\beta\left(\frac{\delta}{4B|\mathcal{X}|}\right)(\sqrt{N}+1)(\log\log(N)+2)\sqrt{\frac{8\gamma_k(N)}{\log(1+1/\tau^2)}} \\
 &\leq \left(12C_k + \frac{4R}{\tau}\sqrt{2\log\left(\frac{4|\mathcal{X}|(\log\log(N)+2)}{\delta}\right)}\right)(\sqrt{N}+1)(\log\log(N)+2)\sqrt{\frac{8\gamma_k(N)}{\log(1+1/\tau^2)}} \\
 &= \mathcal{O}\left(\sqrt{N\gamma_k(N)\log\log(N)\log\left(\frac{|\mathcal{X}|\log\log(N)}{\delta}\right)}\right),
 \end{aligned}$$

where the third line uses the definition of  $\beta(\cdot)$  in Theorem 3.5. Using  $\tilde{\mathcal{O}}(\cdot)$  notation, this simplifies to  $\mathcal{R}(N) = \tilde{\mathcal{O}}\left(\sqrt{N\gamma_k(N)\log\left(\frac{|\mathcal{X}|}{\delta}\right)}\right)$ , which establishes Theorem 5.2.

**Extension to continuous domains.** The proof of Theorem 5.3 is similar to that of Theorem 5.2. Define the event  $\mathcal{E}'_1 = \{\forall i, \forall x \in \mathcal{X}, l_i(x) \leq f(x) \leq u_i(x)\}$ , for  $l_i, u_i$  given in (15). By Corollary 3.7, we have  $\Pr[\mathcal{E}'_1] \geq 1 - \frac{\delta}{2}$ . Recalling that  $\mathcal{E}_2 = \{\forall i, j, \lambda_{\max}^{(i,j)} \leq \tau^2\}$  and  $\Pr[\mathcal{E}_2] \geq 1 - \frac{\delta}{2}$ , we can condition on  $\mathcal{E}'_1 \cap \mathcal{E}_2$ , which holds with probability at least  $1 - \delta$ . Then, following similar steps as in the proof of Theorem 5.2, we have

$$f(x^*) - f(x_n) \leq 4\tilde{\beta}_{N_{i-1}}\left(\frac{\delta}{4B}\right)\left(\sqrt{\frac{8\gamma_k(N_{i-1})}{\log(1+1/\tau^2)N_{i-1}}} + \frac{2}{\sqrt{N_{i-1}}}\right) + \frac{8}{N_{i-1}},$$

and consequently, for  $i \geq 2$ ,

$$\mathcal{R}_i \leq 4\tilde{\beta}_N\left(\frac{\delta}{4B}\right)\left((\sqrt{N}+1)\sqrt{\frac{8\gamma_k(N)}{\log(1+1/\tau^2)}} + 2(\sqrt{N}+1)\right) + 8\sqrt{N}. \quad (33)$$

Note that  $\tilde{\beta}_n(\delta)$  is non-decreasing in  $n$ . Summing the regret over batches, we have

$$\begin{aligned}
 \mathcal{R}(N) &= \sum_{i=1}^B \mathcal{R}_i \\
 &\leq 4\tilde{\beta}_N\left(\frac{\delta}{4B}\right)(\sqrt{N}+1)(\log\log(N)+2)\left(\sqrt{\frac{8\gamma_k(N)}{\log(1+1/\tau^2)}} + 2\right) + 8\sqrt{N}(\log\log(N)+2) \\
 &\leq \left(12C_k + \frac{4R}{\tau}\sqrt{2\log\left(\frac{8\Gamma_N(\log\log(N)+2)}{\delta}\right)}\right)(\sqrt{N}+1)(\log\log(N)+2)\left(\sqrt{\frac{8\gamma_k(N)}{\log(1+1/\tau^2)}} + 2\right) \\
 &\quad + 8\sqrt{N}(\log\log(N)+2) \\
 &= \mathcal{O}\left(\sqrt{N\gamma_k(N)\log\log(N)\log\left(\frac{\Gamma_N\log\log(N)}{\delta}\right)}\right).
 \end{aligned}$$

Recall the notations  $\tilde{\beta}_n(\delta) = \beta\left(\frac{\delta}{2\Gamma_n}\right)$ ,  $\Gamma_n = c(C_k^{\bar{\mu}_n}(\frac{\delta}{2}))^d n^d$  and  $C_k^{\bar{\mu}_n}(\delta) = C_k(1 + \frac{\sqrt{nk_{\max}}}{\tau}) + \frac{\sqrt{nR}}{\tau}\sqrt{2\log(\frac{2n}{\delta})}$ . This yields  $\Gamma_n = \tilde{\mathcal{O}}(C_k^d n^{3d/2})$ , and we arrive at

$$\mathcal{R}(N) = \tilde{\mathcal{O}}\left(\sqrt{Nd\gamma_k(N)\log\left(\frac{N}{\delta}\right)}\right),$$

which completes the proof of Theorem 5.3.

## I. Proofs of Auxiliary Lemmas

In this section, we prove the auxiliary lemmas stated in the preceding sections.

### I.1. Proof of Lemma 3.1 (Expression for the Projection)

We can express the projection as  $\Pi_\kappa[f] = \sum_{j=1}^m \alpha_j^* k(\cdot, z_j)$ , where  $\alpha^* = [\alpha_1, \dots, \alpha_m]^\top$  is the solution of

$$\min_{\alpha \in \mathbb{R}^m} \left\| f - \sum_{j=1}^m \alpha_j k(\cdot, z_j) \right\|_{\mathcal{H}_k}. \quad (34)$$

By the reproducing property, we have

$$\left\| f(\cdot) - \sum_{j=1}^m \alpha_j k(\cdot, z_j) \right\|_{\mathcal{H}_k}^2 = \|f\|_{\mathcal{H}_k}^2 - 2\alpha^\top f_{Z_m} + \alpha^\top k_{Z_m, Z_m} \alpha. \quad (35)$$

Explicitly solving for the minimum of this quadratic expression gives  $\alpha^* = k_{Z_m, Z_m}^{-1} f_{Z_m}$ , and thus  $\Pi_\kappa[f] = k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} f_{Z_m}$ , which completes the proof.

### I.2. Proof of Lemma C.1 (Expression for Term 2)

We show that Term 2 =  $k_{Z_m}^\top(\cdot) (k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m})^{-1} k_{Z_m}(\cdot)$  is equal to  $\sigma_{\kappa, \tau, X_n}^2(\cdot)$ . We have

$$\begin{aligned} \text{Term 2} &= k_{Z_m}^\top(\cdot) \left( k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{Z_m}(\cdot) \\ &= k_{Z_m}^\top(\cdot) \left( k_{Z_m, Z_m}^{-1} - k_{Z_m, Z_m}^{-1} k_{X_n, Z_m}^\top \left( k_{X_n, Z_m} k_{Z_m, Z_m}^{-1} k_{X_n, Z_m}^\top + \tau^2 I_n \right)^{-1} k_{X_n, Z_m} k_{Z_m, Z_m}^{-1} \right) k_{Z_m}(\cdot) \\ &= k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} k_{Z_m}(\cdot) \\ &\quad - k_{Z_m}^\top(\cdot) k_{Z_m, Z_m}^{-1} k_{Z_m, X_n} \left( k_{Z_m, X_n} k_{Z_m, Z_m}^{-1} k_{Z_m, X_n} + \tau^2 I_n \right)^{-1} k_{Z_m, X_n} k_{Z_m, Z_m}^{-1} k_{Z_m}(\cdot) \\ &= \kappa(\cdot, \cdot) - \kappa_{X_n}^\top(\cdot) (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot) \\ &= \sigma_{\kappa, \tau, X_n}^2(\cdot), \end{aligned} \quad (36)$$

where the second line follows from Woodbury matrix identity, and the fourth line follows from the definition of  $\kappa(\cdot, \cdot)$  given in Equation (6).

### I.3. Proof of Lemma C.2 (Expression for $V_n$ )

We have

$$\begin{aligned} \kappa_{X_n}^\top(\cdot) (\tau^2 I_n + \kappa_{X_n, X_n})^{-1} &= k_{Z_m}(\cdot) k_{Z_m, Z_m}^{-1} k_{X_n, Z_m}^\top \left( k_{X_n, Z_m} k_{Z_m, Z_m}^{-1} k_{X_n, Z_m}^\top + \tau^2 I_n \right)^{-1} \\ &= k_{Z_m}(\cdot) k_{Z_m, Z_m}^{-1} \left( k_{X_n, Z_m}^\top k_{X_n, Z_m} k_{Z_m, Z_m}^{-1} + \tau^2 I_n \right)^{-1} k_{X_n, Z_m}^\top \\ &= k_{Z_m}^\top(\cdot) \left( \tau^2 k_{Z_m, Z_m} + k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{X_n, Z_m}^\top \\ &= V_n, \end{aligned}$$

where the first line comes from the definition of  $\kappa$ , the second line uses the push-through matrix identity  $A(I + BA)^{-1} = (I + AB)^{-1}A$  with  $A = k_{X_n, Z_m}^\top$  and  $B = k_{X_n, Z_m} k_{Z_m, Z_m}^{-1}$ , and the third line uses  $C^{-1}D^{-1} = (DC)^{-1}$  with  $C = k_{Z_m, Z_m}$  and  $D = \tau^2 I_n + k_{X_n, Z_m}^\top k_{X_n, Z_m}$ .

#### I.4. Proof of Lemma D.2 (High-Probability Bounds on Term 4)

For a sub-Gaussian random variable  $X$  with parameter  $R$ , Chernoff-Hoeffding inequality (e.g., see Antonini et al., 2008) implies that the following inequalities hold with probability at least  $1 - \delta$  each (or  $1 - 2\delta$  jointly):

$$\begin{aligned} X &\leq R\sqrt{2\log\left(\frac{1}{\delta}\right)}, \\ X &\geq -R\sqrt{2\log\left(\frac{1}{\delta}\right)}. \end{aligned}$$

When the vector  $V_n(x)$  is constructed independent of  $E_n$ , and under Assumption 3.4,  $V_n^\top(x)E_n$  is sub-Gaussian with parameter  $R\|V_n(x)\|_{l^2}$ . Using Theorem 3.2, we have  $R\|V_n(x)\|_{l^2} \leq \frac{R\bar{\sigma}_n(x)}{\tau}$ . Thus, applying Chernoff-Hoeffding inequality to  $V_n^\top(x)E_n$ , we have the following with probability  $1 - \delta$  each:

$$\begin{aligned} V_n^\top(x)E_n &\leq \frac{R\bar{\sigma}_n(x)}{\tau}\sqrt{2\log\left(\frac{1}{\delta}\right)}, \\ V_n^\top(x)E_n &\geq -\frac{R\bar{\sigma}_n(x)}{\tau}\sqrt{2\log\left(\frac{1}{\delta}\right)}. \end{aligned}$$

#### I.5. Proof of Lemma E.3 (High-Probability Bounds on the Noise Term)

When  $V_n(x)$  is independent of  $E_n$ , and when the noise vector  $E_n$  satisfies Assumption E.1, following a standard Chernoff-like confidence bound, the following equations each hold with probability at least  $1 - \delta$  (e.g., see Vakili et al., 2021a):

$$\begin{aligned} V_n^\top(x)E_n &\leq \|V_n(x)\|_{l^2}\sqrt{2\max\left\{\xi_0, \frac{2\log\left(\frac{1}{\delta}\right)}{h_0^2}\right\}\log\left(\frac{1}{\delta}\right)}, \\ V_n^\top(x)E_n &\geq -\|V_n(x)\|_{l^2}\sqrt{2\max\left\{\xi_0, \frac{2\log\left(\frac{1}{\delta}\right)}{h_0^2}\right\}\log\left(\frac{1}{\delta}\right)}. \end{aligned}$$

Substituting  $\|V_n(x)\|_{l^2} \leq \frac{\bar{\sigma}_n(x)}{\tau}$  from Theorem 3.2, we arrive at the lemma.

#### I.6. Proof of Lemma F.1 (RKHS Norm of the Approximate Posterior Mean)

Recalling that  $\bar{\mu}_n(\cdot) = V_n^\top(\cdot)Y_n$ , we have

$$\begin{aligned} \|\bar{\mu}_n(\cdot)\|_{\mathcal{H}_k} &= \|V_n^\top(\cdot)Y_n\|_{\mathcal{H}_k} \\ &= \|V_n^\top(\cdot)(f_{X_n} + E_n)\|_{\mathcal{H}_k} \\ &\leq \underbrace{\|V_n^\top(\cdot)\Pi_\kappa[f]_{X_n}\|_{\mathcal{H}_k}}_{\text{Term 1}} + \underbrace{\|V_n^\top(\cdot)(f_{X_n} - \Pi_\kappa[f]_{X_n})\|_{\mathcal{H}_k}}_{\text{Term 2}} + \underbrace{\|V_n^\top(\cdot)E_n\|_{\mathcal{H}_k}}_{\text{Term 3}}. \end{aligned}$$

We proceed by bounding these three terms.

**Term 1.** Recall the RKHS-based definition of  $\bar{\mu}_n$ :

$$\bar{\mu}_n = \arg \min_{g \in \mathcal{H}_\kappa} \sum_{i=1}^n (y_i - g(x_i))^2 + \tau^2 \|g\|_{\mathcal{H}_\kappa}^2.$$

Consider a noiseless projected dataset  $\tilde{\mathcal{D}}_n = \{(x_i, \Pi_\kappa[f](x_i))_{i=1}^n\}$ , and note that  $\bar{\mu}_{k,\tau,Z_m,X_n,\Pi_\kappa[f]_{X_n}} = V_n^\top(x)\Pi_\kappa[f]_{X_n}$ . Moreover, using the RKHS-based definition of  $\bar{\mu}_{k,\tau,Z_m,X_n,\Pi_\kappa[f]_{X_n}}$ , we have

$$\begin{aligned} \sum_{i=1}^n (\Pi_\kappa[f](x_i) - V_n^\top(x_i)\Pi_\kappa[f]_{X_n})^2 + \tau^2 \|V_n^\top(\cdot)\Pi_\kappa[f]_{X_n}\|_{\mathcal{H}_k}^2 &\leq \sum_{i=1}^n (\Pi_\kappa[f](x_i) - \Pi_\kappa[f](x_i))^2 + \tau^2 \|\Pi_\kappa[f]\|_{\mathcal{H}_\kappa}^2 \\ &= \tau^2 \|\Pi_\kappa[f]\|_{\mathcal{H}_\kappa}^2. \end{aligned}$$

Thus,  $\|V_n^\top(\cdot)\Pi_\kappa[f]_{X_n}\|_{\mathcal{H}_k} \leq \|\Pi_\kappa[f]\|_{\mathcal{H}_\kappa} = \|\Pi_\kappa[f]\|_{\mathcal{H}_k} \leq \|f\|_{\mathcal{H}_k} \leq C_k$ .

**Term 2.** First note that for any vector  $v \in \mathbb{R}^n$ , we have

$$\begin{aligned}
 \|V_n^\top(\cdot)v\|_{\mathcal{H}_\kappa}^2 &= \langle V_n^\top(\cdot)v, V_n^\top(\cdot)v \rangle_{\mathcal{H}_\kappa} \\
 &= \langle \kappa_{X_n}^\top(\cdot)(\tau^2 I_n + \kappa_{X_n, X_n})^{-1}v, \kappa_{X_n}^\top(\cdot)(\tau^2 I_n + \kappa_{X_n, X_n})^{-1}v \rangle_{\mathcal{H}_\kappa} \\
 &= v^\top (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n, X_n} (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} v \\
 &= v^\top (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} (\kappa_{X_n, X_n} + \tau^2 I_n - \tau^2 I_n) (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} v \\
 &= v^\top (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} v - \tau^2 v^\top (\kappa_{X_n, X_n} + \tau^2 I_n)^{-2} v \\
 &\leq v^\top (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} v \\
 &\leq \frac{\|v\|_{l_2}^2}{\tau^2},
 \end{aligned} \tag{37}$$

where the second line uses  $V_n^\top(\cdot) = \kappa_{X_n}^\top(\cdot)(\tau^2 I_n + \kappa_{X_n, X_n})^{-1}$  from Lemma C.2, and the third line follows from the reproducing property of the RKHS. Thus,

$$\begin{aligned}
 \|V_n^\top(\cdot)(f_{X_n} - \Pi_\kappa[f]_{X_n})\|_{\mathcal{H}_\kappa} &\leq \frac{\|f_{X_n} - \Pi_\kappa[f]_{X_n}\|_{l_2}}{\tau} \\
 &\leq \frac{\sqrt{n} \max_{1 \leq i \leq n} (f(x_i) - \Pi_\kappa[f](x_i))}{\tau}.
 \end{aligned}$$

Recall the notations  $f^\perp = f - \Pi_\kappa[f]$  and  $\kappa^\perp = k - \kappa$ . By the reproducing property of the RKHS, we have

$$\begin{aligned}
 f^\perp(x) &= \langle f^\perp(\cdot), \kappa^\perp(x, \cdot) \rangle_{\mathcal{H}_{\kappa^\perp}} \\
 &\leq \|f^\perp\|_{\mathcal{H}_{\kappa^\perp}} \kappa^\perp(x, x) \\
 &\leq C_k k_{\max},
 \end{aligned}$$

where we used  $\|f^\perp\|_{\mathcal{H}_{\kappa^\perp}} = \|f^\perp\|_{\mathcal{H}_\kappa} \leq \|f\|_{\mathcal{H}_\kappa} \leq C_k$ , as well as  $\kappa^\perp(x, x) \leq k(x, x)$ , which follows from the definition of  $\kappa$ .

Thus, we have

$$\|V_n^\top(\cdot)(f_{X_n} - \Pi_\kappa[f]_{X_n})\|_{\mathcal{H}_\kappa} \leq \frac{\sqrt{n} C_k k_{\max}}{\tau}.$$

**Term 3.** For the third term, using  $\|V_n^\top(\cdot)v\|_{\mathcal{H}_\kappa} \leq \frac{\|v\|_{l_2}}{\tau}$  as established in (37), we have

$$\|V_n^\top(\cdot)E_n\|_{\mathcal{H}_\kappa}^2 \leq \frac{\|E_n\|_{l_2}^2}{\tau^2}.$$

Under Assumption 3.4, as a result of Chernoff-Hoeffding inequality, we have for each  $i \in \{1, \dots, n\}$  that

$$\epsilon_i^2 \leq 2R^2 \log\left(\frac{2}{\delta}\right)$$

with probability at least  $1 - \delta$ . Hence, using a probability union bound, we obtain

$$\frac{\|E_n\|_{l_2}^2}{\tau^2} \leq \frac{2nR^2}{\tau^2} \log\left(\frac{2n}{\delta}\right)$$

with probability at least  $1 - \delta$ , which in turn gives

$$\|V_n^\top(\cdot)E_n\|_{\mathcal{H}_\kappa} \leq \frac{\sqrt{n}R}{\tau} \sqrt{2 \log\left(\frac{2n}{\delta}\right)}.$$

Putting the bounds on Terms 1 to 3 together, we arrive at the lemma.

### I.7. Proof of Lemma F.2 (Discretization Error in the Posterior Standard Deviation)

Recall the expression

$$\bar{\sigma}_n^2(\cdot, \cdot) = \underbrace{k(\cdot, \cdot) - k_{Z_m}^\top(\cdot)k_{Z_m, Z_m}^{-1}k_{Z_m}(\cdot)}_{\text{Term 1}} + \underbrace{k_{Z_m}^\top(\cdot) \left( k_{Z_m, Z_m} + \frac{1}{\tau^2} k_{X_n, Z_m}^\top k_{X_n, Z_m} \right)^{-1} k_{Z_m}(\cdot)}_{\text{Term 2}},$$

from the proof of Theorem 3.2, and recall that Term 2 =  $\sigma_{\kappa, \tau, X_n}^2(\cdot)$  from Lemma C.1. By the definition  $\kappa(\cdot, \cdot) = k_{Z_m}^\top(\cdot)k_{Z_m, Z_m}^{-1}k_{Z_m}(\cdot)$ , it follows that

$$\bar{\sigma}_n^2(\cdot, \cdot) = k(\cdot, \cdot) - \kappa_{X_n}^\top(\cdot)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot).$$

By reproducing property, we have for all  $x \in \mathcal{X}$  that  $\|k(x, \cdot)\|_{\mathcal{H}_k} = k(x, x) \leq k_{\max}$ , where  $k_{\max} = \sup_{x \in \mathcal{X}} k(x, x)$ . We also have for all  $x \in \mathcal{X}$  that

$$\begin{aligned} \|\kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot)\|_{\mathcal{H}_k}^2 &= \langle \kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot), \kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot) \rangle \\ &= \kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n, X_n} (\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(x) \\ &= \kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(x) - \kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-2} \kappa_{X_n}(x) \\ &\leq \frac{1}{\tau^2} \|\kappa_{X_n}(x)\|^2 \\ &\leq \frac{nk_{\max}^2}{\tau^2}, \end{aligned}$$

by similar steps to those leading up to (37). Thus,  $\|\kappa_{X_n}^\top(x)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot)\|_{\mathcal{H}_k} \leq \frac{\sqrt{nk_{\max}}}{\tau}$ .

In the following, define the function  $q(\cdot, \cdot) = \kappa_{X_n}^\top(\cdot)(\kappa_{X_n, X_n} + \tau^2 I_n)^{-1} \kappa_{X_n}(\cdot)$ . For all  $x \in \mathcal{X}$ , applying Assumption 3.6 to  $k(x, \cdot)$  and  $q(x, \cdot)$ , with  $\mathbb{X}$ , we have  $|k(x, x') - k(x, [x'])| \leq \frac{1}{n}$  and  $|q(x, x') - q(x, [x'])| \leq \frac{1}{n}$ . Thus,

$$\begin{aligned} |\bar{\sigma}_n^2(x) - \bar{\sigma}^2([x])| &= |(k(x, x) - q(x, x)) - (k([x], [x]) - q([x], [x]))| \\ &= |(k(x, x) - q(x, x)) - (k(x, [x]) - q(x, [x])) + (k(x, [x]) - q(x, [x])) - (k([x], [x]) - q([x], [x]))| \\ &\leq |k(x, x) - k(x, [x])| + |k(x, [x]) - k([x], [x])| + |q(x, x) - q(x, [x])| - |q(x, [x]) - q([x], [x])| \\ &\leq \frac{4}{n}. \end{aligned}$$

To obtain a discretization error bound for the standard deviation from that of the variance, we write

$$\begin{aligned} (\bar{\sigma}_n(x) - \bar{\sigma}([x]))^2 &\leq |\bar{\sigma}_n(x) - \bar{\sigma}([x])| (\bar{\sigma}_n(x) + \bar{\sigma}([x])) \\ &= |\bar{\sigma}_n^2(x) - \bar{\sigma}^2([x])| \\ &\leq \frac{4}{n}. \end{aligned}$$

Therefore,

$$|\bar{\sigma}_n(x) - \bar{\sigma}([x])| \leq \frac{2}{\sqrt{n}}.$$

### I.8. Proof of Lemma G.1 (Bounds on the Approximate Posterior Variance)

For this proof, we introduce a feature map  $\phi(\cdot)$  (for example, the Mercer feature vector) such that  $k(x, x') = \phi^\top(x)\phi(x')$ . Similarly, we introduce  $\psi(\cdot)$  such that  $\kappa(x, x') = \psi^\top(x)\psi(x')$ . Let  $\Phi = [\phi^\top(x_1), \dots, \phi^\top(x_n)]^\top$  be the feature matrix at observation points, and similarly for  $\Psi$ . Note that  $k_{X_n, X_n} = \Phi\Phi^\top$  and  $\kappa_{X_n, X_n} = \Psi\Psi^\top$ . For this proof, we use the feature space forms of the posterior variances (e.g., see, Calandriello et al., 2019):  $\sigma_n^2(x) = \phi^\top(x)(\Phi^\top\Phi + \tau^2 I)^{-1}\phi(x)$ , and  $\bar{\sigma}_n^2(x) = \phi^\top(x)(\Psi^\top\Psi + \tau^2 I)^{-1}\phi(x)$ .

With these definitions in place, the definition of  $\lambda_{\max}$  can equivalently be expressed as

$$\Phi^\top \Phi - \Psi^\top \Psi \preceq \lambda_{\max} I,$$

where  $I$  is the identity operator, and the notation  $A \preceq B$  means  $B - A$  is positive semi-definite. Thus,

$$\begin{aligned} \Phi^\top \Phi + \tau^2 I &\preceq \Psi^\top \Psi + \tau^2 I + \lambda_{\max} I \\ &\preceq \Psi^\top \Psi + \tau^2 I + \frac{\lambda_{\max}}{\tau^2} (\Psi^\top \Psi + \tau^2 I) \\ &= \left(1 + \frac{\lambda_{\max}}{\tau^2}\right) (\Psi^\top \Psi + \tau^2 I), \end{aligned}$$

and rearranging gives

$$(\Psi^\top \Psi + \tau^2 I)^{-1} \preceq \left(1 + \frac{\lambda_{\max}}{\tau^2}\right) (\Phi^\top \Phi + \tau^2 I)^{-1}.$$

Thus, by the above equations for posterior variances in the feature space representation, we obtain

$$\bar{\sigma}_n^2(x) \leq \left(1 + \frac{\lambda_{\max}}{\tau^2}\right) \sigma_n^2(x).$$

The remaining inequality in Lemma G.1 follows similarly from  $\Psi^\top \Psi \preceq \Phi^\top \Phi$ .