

No-Regret Bayesian Optimization with Unknown Hyperparameters

Felix Berkenkamp

*Department of Computer Science
ETH Zurich
Zurich, Switzerland*

BEFELIX@INF.ETHZ.CH

Angela P. Schoellig

*Institute for Aerospace Studies
University of Toronto
Toronto, Canada*

SCHOELLIG@UTIAS.UTORONTO.CA

Andreas Krause

*Department of Computer Science
ETH Zurich
Zurich, Switzerland*

KRAUSEA@ETHZ.CH

Editor: Bayesian Optimization Special Issue

Abstract

Bayesian optimization (BO) based on Gaussian process models is a powerful paradigm to optimize black-box functions that are expensive to evaluate. While several BO algorithms provably converge to the global optimum of the unknown function, they assume that the hyperparameters of the kernel are known in advance. This is not the case in practice and misspecification often causes these algorithms to converge to poor local optima. In this paper, we present the first BO algorithm that is provably no-regret and converges to the optimum without knowledge of the hyperparameters. During optimization we slowly adapt the hyperparameters of stationary kernels and thereby expand the associated function class over time, so that the BO algorithm considers more complex function candidates. Based on the theoretical insights, we propose several practical algorithms that achieve the empirical sample efficiency of BO with online hyperparameter estimation, but retain theoretical convergence guarantees. We evaluate our method on several benchmark problems.

Keywords: Bayesian optimization, Unknown hyperparameters, Reproducing kernel Hilbert space (RKHS), Bandits, No regret

1. Introduction

The performance of machine learning algorithms often critically depends on the choice of tuning inputs, e.g., learning rates or regularization constants. Picking these correctly is a key challenge. Traditionally, these inputs are optimized using grid or random search (Bergstra and Bengio, 2012). However, as data sets become larger the computation time required to train a single model increases, which renders these approaches less applicable. Bayesian optimization (BO, Mockus (2012)) is an alternative method that provably determines good inputs within few evaluations of the underlying objective function. BO methods construct a statistical model of the underlying objective function and use it to evaluate inputs that

are informative about the optimum. However, the theoretical guarantees, empirical performance, and data efficiency of BO algorithms critically depend on their own choice of hyperparameters and, in particular, on the prior distribution over the function space. Thus, we effectively shift the problem of tuning inputs one level up, to the tuning of hyperparameters of the BO algorithm.

In this paper, we use a Gaussian processes (GP, Rasmussen and Williams (2006)) for the statistical model. We present the first BO algorithm that does not require knowledge about the hyperparameters of the GP’s stationary kernel and provably converges to the global optimum. To this end, we adapt the hyperparameters of the kernel and our BO algorithm, so that the associated function space grows over time. The resulting algorithm provably converges to the global optimum and retains theoretical convergence guarantees, even when combined with online estimation of hyperparameters.

Related work General BO has received a lot of attention in recent years. Typically, BO algorithms suggest inputs to evaluate by maximizing an acquisition function that measures informativeness about the optimum. Classical acquisition functions are the *expected improvement* over the best known function value encountered so far given the GP distribution (Mockus et al., 1978) and the *Upper Confidence Bound* algorithm, GP-UCB, which applies the ‘optimism in the face of uncertainty’ principle. The latter is shown to provably converge by Srinivas et al. (2012). Durand et al. (2018) extend this framework to the case of unknown measurement noise. A related method is *truncated variance reduction* by Bogunovic et al. (2016), which considers the reduction in uncertainty at candidate locations for the optimum. Hennig and Schuler (2012) propose *entropy search*, which approximates the distribution of the optimum of the objective function and uses the reduction of the entropy in this distribution as an acquisition function. Alternative information-theoretic methods are proposed by Hernández-Lobato et al. (2014); Wang and Jegelka (2017); Ru et al. (2018). Other alternatives are the *knowledge gradient* (Frazier et al., 2009), which is one-step Bayes optimal, and *information directed sampling* by Russo and Van Roy (2014), which considers a tradeoff between regret and information gained when evaluating an input. Kirschner and Krause (2018) extend the latter framework to heteroscedastic noise.

These BO methods have also been successful empirically. In machine learning, they are used to optimize the performance of learning methods (Brochu et al., 2010; Snoek et al., 2012). BO is also applicable more broadly; for example, in reinforcement learning to optimize a parametric policy for a robot (Calandra et al., 2014; Lizotte et al., 2007; Berkenkamp et al., 2016) or in control to optimize the energy output of a power plant (Abdelrahman et al., 2016). It also forms the backbone of Google vizier, a service for tuning black-box functions (Golovin et al., 2017).

Some of the previous BO algorithms provide theoretical guarantees about convergence to the optimum. These theoretical guarantees only hold when the kernel hyperparameters are known *a priori*. When this is not the case, hyperparameters are often inferred using either *maximum a posteriori* estimates or sampling-based inference (Snoek et al., 2012). Unfortunately, methods that estimate the hyperparameters online are known to get stuck in local optima (Bull, 2011). Instead, we propose to adapt the hyperparameters online in order to enlarge the function space over time, which allows us to provide guarantees in terms of convergence to the global optimum without knowing the hyperparameters. Wang and de

Freitas (2014) analyze this setting when a lower bound on the kernel lengthscales is known *a priori*. They decrease the lengthscales over time and bound the regret in terms of the known lower-bound on the lengthscales. Empirically, similar heuristics are used by Wang et al. (2016); Wabersich and Toussaint (2016). In contrast, this paper considers the case where the hyperparameters are *not known*. Moreover, the scaling of the hyperparameters in the previous two papers did not depend on the dimensionality of the problem, which can cause the function space to increase too quickly.

Considering larger function classes as more data becomes available is the core idea behind structural risk minimization (Vapnik, 1992) in statistical learning theory. However, there data is assumed to be sampled independently and identically distributed. This is not the case in BO, where new data is generated actively based on past information.

Our contribution In this paper, we present Adaptive GP-UCB (A-GP-UCB), the first algorithm that provably converges to the globally optimal inputs when BO hyperparameters are *unknown*. Our method expands the function class encoded in the model over time, but does so slowly enough to ensure sublinear regret and convergence to the optimum. Based on the theoretical insights, we propose practical variants of the algorithm with guaranteed convergence. Since our method can be used as an add-on module to existing algorithms with hyperparameter estimation, it achieves similar performance empirically, but avoids local optima when hyperparameters are misspecified. In summary, we:

- Provide theoretical convergence guarantees for BO with unknown hyperparameters;
- Propose several practical algorithms based on the theoretical insights;
- Evaluate the performance in practice and show that our method retains the empirical performance of heuristic methods based on online hyperparameter estimation, but leads to significantly improved performance when the model is misspecified initially.

The remainder of the paper is structured as follows. We state the problem in Sec. 2 and provide relevant background material in Sec. 3. We derive our main theoretical result in Sec. 4 and use insights gained from the theory to propose practical algorithms. We evaluate these algorithms experimentally in Sec. 5 and draw conclusions in Sec. 6. The technical details of the proofs are given in the appendix.

2. Problem Statement

In general, BO considers global optimization problems of the form

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{D}} f(\mathbf{x}), \quad (1)$$

where $\mathcal{D} \subset \mathbb{R}^d$ is a compact domain over which we want to optimize inputs \mathbf{x} , and $f: \mathcal{D} \rightarrow \mathbb{R}$ is an objective function that evaluates the reward $f(\mathbf{x})$ associated with a given input configuration \mathbf{x} . For example, in a machine learning application, $f(\mathbf{x})$ may be the validation loss and \mathbf{x} may be the tuning inputs (e.g., regularization parameters) of the training algorithm. We do not have any significant prior knowledge about the structure of f . Specifically, we cannot assume convexity or that we have access to gradient information. Moreover, evaluations of f are corrupted by σ -sub-Gaussian noise, a general class of noise models that includes, for example, bounded or Gaussian noise.

Regret We aim to construct a sequence of input evaluations \mathbf{x}_t , that eventually maximizes the function value $f(\mathbf{x}_t)$. One natural way to prove this convergence is to show that an algorithm has sublinear regret. The instantaneous regret at iteration t is defined as $r_t = \max_{\mathbf{x} \in \mathcal{D}} f(\mathbf{x}) - f(\mathbf{x}_t) \geq 0$, which is the loss incurred by evaluating the function at \mathbf{x}_t instead of at the *a priori unknown* optimal inputs. The cumulative regret is defined as $R_T = \sum_{0 < t \leq T} r_t$, the sum of regrets incurred over T steps. If we can show that the cumulative regret is sublinear for a given algorithm, that is, $\lim_{t \rightarrow \infty} R_t/t = 0$, then eventually the algorithm evaluates the function at inputs that lead to close-to-optimal function values most of the time. We say that such an algorithm has *no-regret*. Intuitively, if the average regret approaches zero then, on average, the instantaneous regret must approach zero too, since r_t is strictly positive. This implies that there exists a $t > 0$ such that $f(\mathbf{x}_t)$ is arbitrarily close to $f(\mathbf{x}^*)$ and the algorithm converges. Thus, we aim to design an optimization algorithm that has sublinear regret.

Regularity assumptions Without further assumptions, it is impossible to achieve sub-linear regret on (1). In the worst case, f could be discontinuous at every input in \mathcal{D} . To make the optimization problem in (1) tractable, we make regularity assumptions about f . In particular, we assume that the function f has low complexity, as measured by the norm in a reproducing kernel Hilbert space (RKHS, Christmann and Steinwart (2008)). An RKHS \mathcal{H}_k contains well-behaved functions of the form $f(\mathbf{x}) = \sum_{i \geq 0} \alpha_i k(\mathbf{x}, \mathbf{x}_i)$, for given representer points $\mathbf{x}_i \in \mathbb{R}^d$ and weights $\alpha_i \in \mathbb{R}$ that decay sufficiently quickly. The kernel $k(\cdot, \cdot)$ determines the roughness and size of the function space and the induced RKHS norm $\|f\|_k = \sqrt{\langle f, f \rangle}$ measures the complexity of a function $f \in \mathcal{H}_k$ with respect to the kernel.

In the following, we assume that f in (1) has bounded RKHS norm $\|f\|_{k_\theta} \leq B$ with respect to a kernel k_θ that is parameterized by hyperparameters θ . We write \mathcal{H}_θ for the corresponding RKHS, \mathcal{H}_{k_θ} . For known B and θ , no-regret BO algorithms for (1) are known, e.g., GP-UCB (Srinivas et al., 2012). In practice, these hyperparameters need to be tuned. In this paper, we consider the case where θ and B are unknown. We focus on stationary kernels, which measure similarity based on the distance of inputs, $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$. The most commonly used hyperparameters for these kernels are the lengthscales $\theta \in \mathbb{R}^d$, which scale the inputs to the kernel in order to account for different magnitudes in the different components of \mathbf{x} and effects on the output value. That is, we scale the difference $\mathbf{x} - \mathbf{x}'$ by the lengthscales θ ,

$$k_\theta(\mathbf{x}, \mathbf{x}') = k \left(\frac{[\mathbf{x}]_1 - [\mathbf{x}']_1}{[\theta]_1}, \dots, \frac{[\mathbf{x}]_d - [\mathbf{x}']_d}{[\theta]_d} \right), \quad (2)$$

where $[\mathbf{x}]_i$ denotes the i th element of \mathbf{x} . Typically, these kernels assign larger similarity scores to inputs when the scaled distance between these two inputs is small. Another common hyperparameter is the prior variance of the kernel, a multiplicative constant that determines the magnitude of the kernel. We assume $k(\mathbf{x}, \mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{D}$ without loss of generality, as any multiplicative scaling can be absorbed by the norm bound B .

In summary, our goal is to efficiently solve (1) via a BO algorithm with sublinear regret, where f lies in some RKHS \mathcal{H}_θ , but neither the hyperparameters θ nor the norm-bound $\|f\|_{k_\theta}$ are known.

3. Background

In this section, we review Gaussian processes (GPs) and Bayesian optimization (BO).

3.1. Gaussian processes (GP)

Based on the assumptions in Sec. 2, we can use GPs to infer confidence intervals on f . The goal of GP inference is to infer a posterior distribution over the nonlinear map $f(\mathbf{x}) : D \rightarrow \mathbb{R}$ from an input vector $\mathbf{x} \in D$ to the function value $f(\mathbf{x})$. This is accomplished by assuming that the function values $f(\mathbf{x})$, associated with different values of \mathbf{x} , are random variables and that any finite number of these random variables have a joint Gaussian distribution (Rasmussen and Williams, 2006). A GP distribution is parameterized by a prior mean function and a covariance function or kernel $k(\mathbf{x}, \mathbf{x}')$, which defines the covariance of any two function values $f(\mathbf{x})$ and $f(\mathbf{x}')$ for $\mathbf{x}, \mathbf{x}' \in D$. In this work, the mean is assumed to be zero without loss of generality. The choice of kernel function is problem-dependent and encodes assumptions about the unknown function.

We can condition a $GP(0, k(\mathbf{x}, \mathbf{x}'))$ on a set of t past observations $\mathbf{y}_t = (y_1, \dots, y_t)$ at inputs $\mathcal{A}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$ in order to obtain a posterior distribution on $f(\mathbf{x})$ for any input $\mathbf{x} \in D$. The GP model assumes that observations are noisy measurements of the true function value, $y_t = f(\mathbf{x}_t) + \omega_t$, where $\omega_t \sim \mathcal{N}(0, \sigma^2)$. The posterior distribution is again a $GP(\mu_t(\mathbf{x}), k_t(\mathbf{x}, \mathbf{x}'))$ with mean μ_t , covariance k_t , and variance σ_t , where

$$\mu_t(\mathbf{x}) = \mathbf{k}_t(\mathbf{x})(\mathbf{K}_t + \mathbf{I}\sigma^2)^{-1}\mathbf{y}_t, \quad (3)$$

$$k_t(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_t(\mathbf{x})(\mathbf{K}_t + \mathbf{I}\sigma^2)^{-1}\mathbf{k}_t^\top(\mathbf{x}'), \quad (4)$$

$$\sigma_t^2(\mathbf{x}) = k_t(\mathbf{x}, \mathbf{x}). \quad (5)$$

The covariance matrix $\mathbf{K}_t \in \mathbb{R}^{t \times t}$ has entries $[\mathbf{K}_t]_{(i,j)} = k(\mathbf{x}_i, \mathbf{x}_j)$, $i, j \in \{1, \dots, t\}$, and the vector $\mathbf{k}_t(\mathbf{x}) = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_t)]$ contains the covariances between the input \mathbf{x} and the observed data points in \mathcal{A}_t . The identity matrix is denoted by $\mathbf{I}_t \in \mathbb{R}^{t \times t}$.

3.2. Learning RKHS functions with GPs

The GP framework uses a statistical model that makes different assumptions from the ones made about f in Sec. 2. In particular, we assume a different noise model, and samples from a $GP(0, k(\mathbf{x}, \mathbf{x}'))$ are rougher than RKHS functions and are not contained in \mathcal{H}_k . However, GPs and RKHS functions are closely related (Kanagawa et al., 2018) and it is possible to use GP models to infer reliable confidence intervals on f in (1).

Lemma 1 (Abbasi-Yadkori (2012); Chowdhury and Gopalan (2017)) *Assume that f has bounded RKHS norm $\|f\|_k \leq B$ and that measurements are corrupted by σ -sub-Gaussian noise. If $\beta_t^{1/2} = B + 4\sigma\sqrt{I(\mathbf{y}_t; f) + 1 + \ln(1/\delta)}$, then for all $\mathbf{x} \in D$ and $t \geq 0$ it holds jointly with probability at least $1 - \delta$ that $|f(\mathbf{x}) - \mu_t(\mathbf{x})| \leq \beta_t^{1/2}\sigma_t(\mathbf{x})$.*

Lemma 1 implies that, with high probability, the true function f is contained in the confidence intervals induced by the posterior GP distribution that uses the kernel k from Lemma 1 as a covariance function, scaled by an appropriate factor β_t . Here, $I(\mathbf{y}_t; f)$ denotes the mutual information between the GP prior on f and the t measurements \mathbf{y}_t . Intriguingly, for

GP models this quantity only depends on the inputs \mathbf{x}_t and not the corresponding measurement y_t . Specifically, for a given set of measurements $\mathbf{y}_{\mathcal{A}}$ at inputs $\mathbf{x} \in \mathcal{A}$, the mutual information is given by

$$I(\mathbf{y}_{\mathcal{A}}; f) = 0.5 \log |\mathbf{I} + \sigma^{-2} \mathbf{K}_{\mathcal{A}}|, \quad (6)$$

where $\mathbf{K}_{\mathcal{A}}$ is the kernel matrix $[k(\mathbf{x}, \mathbf{x}')]_{\mathbf{x}, \mathbf{x}' \in \mathcal{A}}$ and $|\cdot|$ is the determinant. Intuitively, the mutual information measures how informative the collected samples $\mathbf{y}_{\mathcal{A}}$ are about the function f . If the function values are independent of each other under the GP prior, they will provide large amounts of new information. However, if measurements are taken close to each other as measured by the kernel, they are correlated under the GP prior and provide less information.

3.3. Bayesian Optimization (BO)

BO aims to find the global maximum of an unknown function (Mockus, 2012). The framework assumes that evaluating the function is expensive in terms of time required or monetary costs, while other computational resources are comparatively inexpensive. In general, BO methods model the objective function f with a statistical model and use it to determine informative sample locations. A popular approach is to model the underlying function with a GP, see Sec. 3.1. GP-based BO methods use the posterior mean and variance predictions in (3) and (5) to compute the next sample location.

One commonly used algorithm is the GP-UCB algorithm by Srinivas et al. (2012). It uses confidence intervals on the function f , e.g., from Lemma 1, in order to select as next input the point with the largest plausible function value according to the model,

$$\mathbf{x}_{t+1} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{D}} \mu_t(\mathbf{x}) + \beta_t^{1/2} \sigma_t(\mathbf{x}). \quad (7)$$

Intuitively, (7) selects new evaluation points at locations where the upper bound of the confidence interval of the GP estimate is maximal. Repeatedly evaluating the function f at inputs \mathbf{x}_{t+1} given by (7) improves the mean estimate of the underlying function and decreases the uncertainty at candidate locations for the maximum, so that the global maximum is provably found eventually (Srinivas et al., 2012). While (7) is also an optimization problem, it only depends on the GP model of f and solving it therefore does not require any expensive evaluations of f .

Regret bounds Srinivas et al. (2012) show that the GP-UCB algorithm has cumulative regret $R_t = \mathcal{O}(\sqrt{t\beta_t\gamma_t})$ for all $t \geq 1$ with the same $(1 - \delta)$ probability as the confidence intervals, e.g., in Lemma 1, hold. Here γ_t is the largest amount of mutual information that could be obtained by any algorithm from at most t measurements,

$$\gamma_t = \max_{\mathcal{A} \subset \mathcal{D}, |\mathcal{A}| \leq t} I(\mathbf{y}_{\mathcal{A}}; f). \quad (8)$$

We refer to γ_t as the *information capacity*, since it can be interpreted as a measure of complexity of the function class associated with a GP prior. It was shown by Srinivas et al. (2012) that γ_t has a sublinear dependence on t for many commonly used kernels such as the Gaussian kernel. As a result, R_t has a sublinear dependence on t so that $R_t/t \rightarrow 0$ and therefore GP-UCB converges to function evaluations close to $f(\mathbf{x}^*)$. These regret bounds

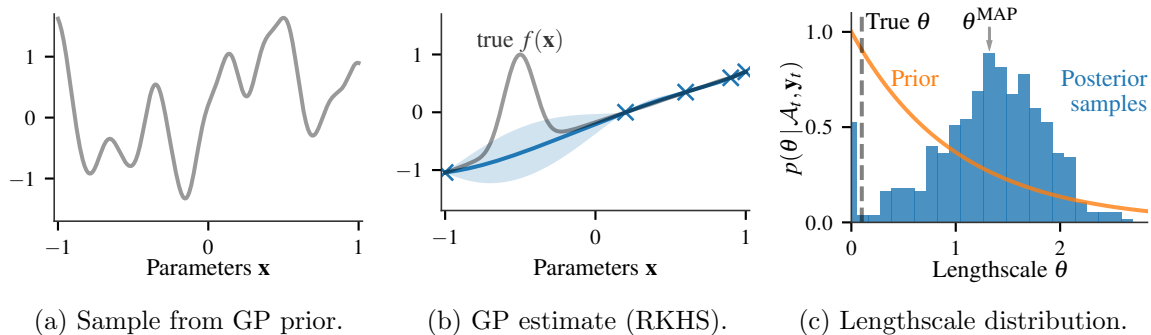


Figure 1: A sample from the GP prior in Fig. 1a typically varies at a consistent rate over the input space. However, RKHS functions with the same kernel may be less consistent and can have bumps, as in Fig. 1b (gray). As a result, inferring the posterior lengthscales based on measurements (blue crosses in Fig. 1b) can lead to erroneous results. In Fig. 1c, most of the probability mass of the posterior lengthscales has concentrated around large lengthscales that encode smooth functions. Consequently, the GP’s 2σ confidence intervals in Fig. 1b (blue shaded) based on the posterior samples do not contain the true function.

were extended to Thompson sampling, an algorithm that uses samples from the posterior GP as the acquisition function, by Chowdhury and Gopalan (2017).

Online hyperparameter estimation In the previous section, we have seen that the GP-UCB algorithm provably converges. However, it requires access to a RKHS norm bound $\|f\|_{\theta} \leq B$ under the correct kernel hyperparameters θ in order to construct reliable confidence intervals using Lemma 1. In practice, these are unknown and have to be estimated online, e.g., based on a prior distribution placed on θ . Unfortunately, it is well-known that online estimation of the inputs, be it via maximum a posteriori (MAP) or sampling methods, does not always converge to the optimum (Bull, 2011). The problem does not primarily lie with the inference scheme, but rather with the assumptions made by the GP. In particular, typical samples drawn from a GP with a stationary kernel tend to have a similar rate of change throughout the input space, see Fig. 1a. In contrast, the functions inside the RKHS, as specified in Sec. 2, can have different rates of change and are thus improbable under the GP prior. For example, the grey function in Fig. 1b is almost linear but has one bump that defines the global maximum, which makes this function an improbable sample under the GP prior even though it belongs to the RKHS induced by the same kernel. This property of GPs with stationary kernels means that, for inference, it is sufficient to estimate the lengthscales in a small part of the state-space in order to make statements about the function space globally. This is illustrated in Fig. 1c, where we show samples from the posterior distribution over the lengthscales based on the measurements obtained from the GP-UCB algorithm in Fig. 1b (blue crosses). Even though the prior distribution on the lengthscales θ is suggestive of short lengthscales, most of the posterior probability mass is concentrated around lengthscales that are significantly larger than the true ones. As a result, even under model averaging over the samples from the posterior distribution of the lengthscales, the GP confidence intervals do not contain the true function in Fig. 1b. This is not a problem of the inference method applied, but rather a direct consequence of




Figure 2: BO algorithms get stuck in local optima when the hyperparameters of the model are misspecified. In Fig. 2a, the true function is not contained within the GP’s confidence intervals (blue shaded), so that GP-UCB only collects data at the local optimum on the right (green arrow), see also Fig. 1. Our method expands the function class over time by scaling the hyperparameters, which encourages additional exploration in Fig. 2b. The function class grows slowly enough, so that the global optimum is provably found in Fig. 2c.

the probabilistic model that we have specified based on the stationary kernel, which does not consider functions with different rates of change to be likely.

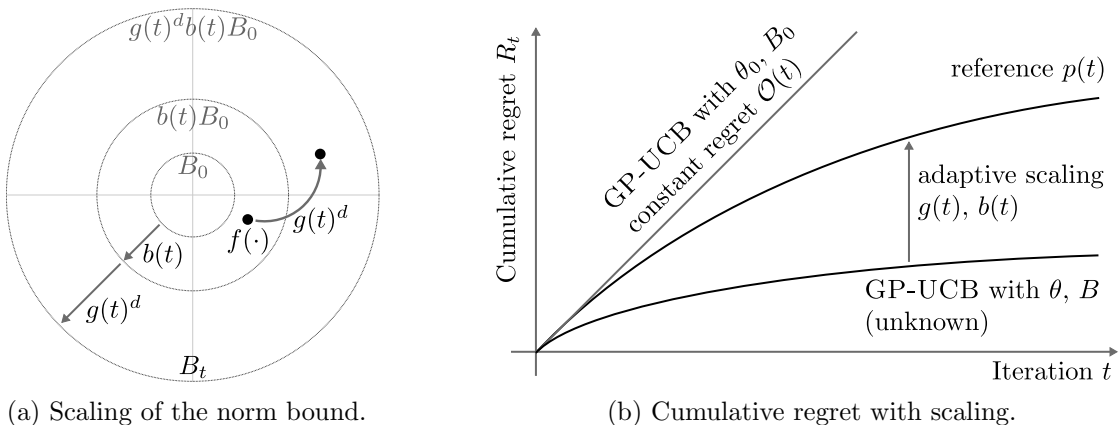
4. The Adaptive GP-UCB Algorithm

In this section, we extend the GP-UCB algorithm to the case where neither the norm bound B nor the lengthscales θ are known. In this case, it is always possible that the local optimum is defined by a local bump based on a kernel with small lengthscales, which has not been encountered by the data points as in Fig. 1b. The only solution to avoid this problem is to keep exploring to eventually cover the input space \mathcal{D} (Bull, 2011). We consider expanding the function space associated with the hyperparameters slowly over time, so that we obtain sublinear regret once the true function class has been identified. Intuitively, this can help BO algorithms avoid premature convergence to local optima caused by misspecified hyperparameters θ and B . For example, in Fig. 2a, the GP-UCB algorithm has converged to a local maximum. By decreasing the lengthscales, we increase the underlying function class, which means that the GP confidence intervals on the function increase. This enables GP-UCB to explore further so that the global optimum is found, as shown in Fig. 2c.

Specifically, we start with an initial guess θ_0 and B_0 for the lengthscales and norm bound on f , respectively. Over the iterations, we scale down the lengthscales and scale up the norm bound,

$$\theta_t = \frac{1}{g(t)} \theta_0, \quad B_t = b(t)g(t)^d B_0, \quad (9)$$

where $g: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ and $b: \mathbb{N} \rightarrow \mathbb{R}_{>0}$ with $b(0) = g(0) = 1$ are functions that can additionally depend on the data collected up to iteration t , \mathcal{A}_t and \mathbf{y}_t . As $g(t)$ increases, the lengthscales θ_t of the kernel become shorter, which enlarges the underlying function space:



(a) Scaling of the norm bound.

(b) Cumulative regret with scaling.

Figure 3: The function f in Fig. 3a has RKHS norm $\|f\|_{\theta_0} > B_0$. To account for this, we expand the norm ball by $b(t)$ over time. When we scale down the lengthscales by $g(t)$, the norm of f in the resulting RKHS is larger, see Lemma 2. We account for this when defining the norm ball B_t in (9). In Fig. 3b, the GP-UCB algorithm based on the misspecified hyperparameters B_0 and θ_0 does not converge (constant regret). Our method scales the lengthscales and norm bound by $g(t)$ and $b(t)$, so that we eventually capture the true model. Scaling the hyperparameters beyond the true ones leads to additional exploration and thus larger cumulative regret than GP-UCB with the true, unknown hyperparameters θ and B . However, as long as the cumulative regret is upper bounded by a sublinear function p , ultimately the A-GP-UCB algorithm converges to the global optimum.

Lemma 2 (Bull, 2011, Lemma 4) *If $f \in \mathcal{H}_\theta$, then $f \in \mathcal{H}_{\theta'}$ for all $0 < \theta' \leq \theta$, and*

$$\|f\|_{\mathcal{H}_{\theta'}}^2 \leq \left(\prod_{i=1}^d \frac{[\theta]_i}{[\theta']_i} \right) \|f\|_{\mathcal{H}_\theta}^2. \quad (10)$$

Lemma 2 states that when decreasing the lengthscales θ , the resulting function space contains the previous one. Thus, as $g(t)$ increases we consider larger RKHS spaces as candidate spaces for the function f . In addition, as we increase $b(t)$, we consider larger norm balls within the function space \mathcal{H}_{θ_t} , which corresponds to more complex functions. However, it follows from (10) that, as we increase $g(t)$, we also increase the norm of any existing function in \mathcal{H}_{θ_0} by at most a factor of $g(t)^d$. This is illustrated in Fig. 3a: as we scale up the norm ball to $b(t)B_0$, we capture f under the initial lengthscales θ_0 . However, by shortening the lengthscales by $g(t)$, the function f has a larger norm in the new function space $\mathcal{H}_{\theta_t} = \mathcal{H}_{\theta_0/g(t)}$. We account for this through the additional scaling factor $g(t)^d$ in the norm bound B_t in (9).

Theoretical analysis Based on the previous derivations together with Lemma 2, it is clear that, if $g(t)$ and $b(t)$ are monotonically increasing functions and $f \in \mathcal{H}_{\theta_{t^*}}$ with $\|f\|_{\theta_{t^*}} \leq B_{t^*}$ for some $t^* > 0$, then $f \in \mathcal{H}_{\theta_t}$ and $\|f\|_{\theta_t} \leq B_t$ for all $t \geq t^*$. That is, once the function f is contained within the norm ball of B_{t^*} for the lengthscales θ_{t^*} , then, for any further increase in $b(t)$ or $g(t)$, the function f is still contained in the candidate space $\{f \in \mathcal{H}_{\theta_t} \mid f \leq B_t\}$.

Based on this insight, we propose A-GP-UCB in Algorithm 1. At iteration t , A-GP-UCB sets the GP lengthscales to θ_t and selects new inputs \mathbf{x}_{t+1} similar to the GP-UCB algorithm, but based on the norm bound B_t . We extend the analysis of GP-UCB and Lemma 1 to obtain our main result.

Theorem 1 *Assume that f has bounded RKHS norm $\|f\|_{k_\theta}^2 \leq B$ in a RKHS that is parametrized by a stationary kernel $k_\theta(\mathbf{x}, \mathbf{x}')$ with unknown lengthscales θ . Based on an initial guess, θ_0 and B_0 , define monotonically increasing functions $g(t) > 0$ and $b(t) > 0$ and run A-GP-UCB with $\beta_t^{1/2} = b(t)g(t)^d B_0 + 4\sigma\sqrt{I_{\theta_t}(\mathbf{y}_t; f) + 1 + \ln(1/\delta)}$ and GP lengthscales $\theta_t = \theta_0/g(t)$. Then, with probability at least $(1 - \delta)$, we obtain a regret bound of*

$$R_t \leq 2B \max \left(g^{-1} \left(\max_i \frac{[\theta_0]_i}{[\theta]_i} \right), b^{-1} \left(\frac{B}{B_0} \right) \right) + \sqrt{C_1 t \beta_t I_{\theta_t}(\mathbf{y}_t; f)}, \quad (11)$$

where I_{θ_t} is the mutual information in (6) based on the GP model with lengthscales θ_t and $C_1 = 8/\log(1 + \sigma^{-2})$.

The proof is given in the appendix. Intuitively, the regret bound in (11) splits the run of the algorithm into two distinct phases. In the first one, either the RKHS space $\mathcal{H}_{\theta_t}(\mathcal{D})$ or the norm bound B_t are too small to contain the true function f . Thus, the GP confidence intervals scaled by $\beta_t^{1/2}$ do not necessarily contain the true function f , as in Fig. 1b. In these iterations, we obtain constant regret that is bounded by $2B$, since $\|f\|_\infty \leq \|f\|_\theta \leq B$. After both g and b have grown sufficiently in order for the considered function space to contain the true function, the confidence bounds are reliable and we can apply the theoretical results of the GP-UCB algorithm. This is illustrated in Fig. 3b: If the initial hyperparameters θ_0 and B_0 are misspecified, the confidence intervals do not contain f and GP-UCB does not converge. We avoid this problem by increasing $b(t)$ and $g(t)$ over time, so that we eventually contain f in our function class. However, increasing the norm ball and decreasing the lengthscales beyond the true ones causes additional exploration and thus additional cumulative regret relative to GP-UCB with the true, unknown hyperparameters. This additional regret represents the cost of not knowing the hyperparameters in advance. As long as the overall regret remains bounded by a sublinear function $p(t)$, our method eventually converges to the global optimum. The regret bound in (11) depends on the true hyperparameters θ and B . However, the algorithm does not depend on them. Theorem 1 provides an instance-specific bound, since the mutual information depends on the inputs in \mathcal{A}_t . One can obtain a worst-case upper bound by bounding $I_{\theta_t}(\mathbf{y}_t; f) \leq \gamma_t(\theta_t)$, which is the worst-case mutual information as in (8), but based on the GP model with lengthscales θ_t . While Theorem 1 assumes that the noise properties are known, the results can be extended to estimate the noise similar to Durand et al. (2018).

For arbitrary functions $g(t)$ and $b(t)$, the candidate function space $\{f \in \mathcal{H}_{\theta_t} \mid f \leq B_t\}$ can grow at a faster rate than it contracts by selecting informative measurements y_t according to (7). In particular, in the regret term $\sqrt{C_1 t \beta_t \gamma_t}$ both β_t and γ_t depend on the scaling factors $g(t)$ and $b(t)$. If these factors grow at a faster rate than \sqrt{t} , the resulting algorithm does not enjoy sublinear regret. We have the following result that explicitly states the dependence of γ_t on the scaling factor $g(t)$.

Proposition 2 *Let k_θ be a stationary kernel parameterized by lengthscales θ as in (2) and define $\gamma_t(\theta)$ for lengthscales θ as in (8). Define the lengthscales as $\theta_t = \theta_0/g(t)$ as in (9).*

Algorithm 1 Adaptive GP-UCB(A-GP-UCB)

- 1: **Input:** Input space \mathcal{D} , $GP(0, k(\mathbf{x}, \mathbf{x}'))$, functions $g(t)$ and $b(t)$
 - 2: Set $B_0 = 1$ and $\theta_0 = \text{diam}(\mathcal{D})$
 - 3: **for all** $t = 0, 1, 2, \dots$ **do**
 - 4: Set the GP kernel lengthscales to $\theta_t = \theta_0/g(t)$
 - 5: $\beta_t^{1/2} \leftarrow B(t) + 4\sigma\sqrt{I_{\theta_t}(\mathbf{y}_t; f) + 1 + \ln(1/\delta)}$ with $B(t) = b(t)g(t)^d B_0$
 - 6: Choose $\mathbf{x}_{t+1} = \text{argmax}_{\mathbf{x} \in \mathcal{D}} \mu_t(\mathbf{x}) + \beta_t^{1/2} \sigma_t(\mathbf{x})$
 - 7: Evaluate $y_{t+1} = f(\mathbf{x}_{t+1}) + \epsilon_{t+1}$
 - 8: Perform Bayesian update to obtain μ_{t+1} and σ_{t+1}
-

- If $k(\mathbf{x}, \mathbf{x}') = \exp(-\frac{1}{2}\|\mathbf{x} - \mathbf{x}'\|_2^2)$ is the squared exponential (Gaussian) kernel, then

$$\gamma_t(\theta_t) = \mathcal{O}\left(g(t)^d (\log t)^{d+1}\right) \quad (12)$$

- If $k(\mathbf{x}, \mathbf{x}') = (2^{1-\nu}/\Gamma(\nu)) r^\nu B_\nu(r)$ is the Matérn kernel, where $r = \sqrt{2\nu}\|\mathbf{x} - \mathbf{x}'\|_2$, B_ν is the modified Bessel function with $\nu > 1$, and Γ is the gamma function. Then

$$\gamma_t(\theta_t) = \mathcal{O}\left(g(t)^{2\nu+d} t^{\frac{d(d+1)}{2\nu+d(d+1)}} \log t\right) \quad (13)$$

Proposition 2 explicitly states the relationship between γ_t and $g(t)$. For the Gaussian kernel, if we scale down the lengthscales by a factor of two, the amount of mutual information that we can gather in the worst case, γ_t , grows by 2^d . Given the dependence of γ_t on $g(t)$, we can refine Theorem 1 to obtain concrete regret bounds for two commonly used kernels.

Corollary 3 *If, under the assumptions of Theorem 1, $g(t)$ and $b(t)$ grow unbounded, then we obtain the following, high-probability regret bounds for Algorithm 1:*

- Squared exponential kernel: $R_t \leq \mathcal{O}\left(b(t)\sqrt{tg(t)^{3d}\gamma_t(\theta_0)} + g(t)^d \gamma_t(\theta_0)\sqrt{t}\right)$;
- Matérn kernel: $R_t \leq \mathcal{O}\left(b(t)\sqrt{tg(t)^{2\nu+3d}\gamma_t(\theta_0)} + g(t)^{\nu+d} \gamma_t(\theta_0)\sqrt{t}\right)$.

If $b(t)$ and $g(t)$ grow unbounded, the first term of the cumulative regret in (11) can be upper bounded by a constant. The remaining result is obtained by plugging in β_t and the bounds from (8). Thus, any functions $g(t)$ and $b(t)$ that render the regret bounds in Corollary 3 sublinear allow the algorithm to converge, even though the true lengthscales and norm bound are unknown.

The specific choices of $b(t)$ and $g(t)$ matter for the regret bound in Theorem 1 in practice. Consider the one-dimensional case $d = 1$ for the Gaussian kernel. Given the true hyperparameters B and θ , if we set $g(t) = \theta_0/\theta$ and $b(t) = B/B_0$ to be constant, we recover the non-adaptive regret bounds of GP-UCB with known hyperparameters. If $g(t)$ depends on t and grows slowly, then the algorithm incurs constant regret during the initial rounds when the model is misspecified, while functions g that grow to values larger than the optimal ones lead to additional exploration and incur an additional $\mathcal{O}(b(t)g(t)^{3d/2})$ factor in the cumulative regret in later rounds, as in Corollary 3. In the following section, we discuss appropriate choices for these functions in practice.

4.1. Choosing the scaling functions $g(t)$ and $b(t)$

It follows from Theorem 1 that A-GP-UCB achieves no-regret for any functions $g(t)$ and $b(t)$ that increase without bound and render (11) sublinear in t . Thus, the corresponding BO routine converges to the optimal value eventually. For example, $b(t) = g(t) = \log(t)$ satisfy this condition. However, the convergence guarantees in Theorem 1 are only meaningful once t has grown sufficiently so that the true function is contained in the confidence intervals. In practice, BO is often used with objective functions f that are expensive to evaluate, which imposes a hard constraint on the number of evaluations. For the regret bounds to be meaningful in this setting, we must choose functions g and b that grow fast enough to ensure that the constant regret period in (11) is small, yet slow enough that the effect of the sublinear regret is visible for small enough t . In the following, we propose two methods to choose $g(t)$ and $b(t)$ *adaptively*, based on the observations seen so far.

For convenience, we fix the relative magnitude of $g(t)$ and $b(t)$. In particular, we define $b(t) = 1 + \epsilon_b(t)$ and $g(t)^d = 1 + \epsilon_g(t)$ together with a weighting factor $\lambda = \epsilon_b(t)/\epsilon_g(t)$ that encodes whether we prefer to scale up the norm bound using $b(t)$ or decrease the lengthscales using $g(t)$. This allows us to reason about the overall magnitude of the scaling $h(t) = (1 + \epsilon_g(t))(1 + \epsilon_b(t)) \geq 1$, which can be uniquely decomposed into $g(t)$ and $b(t)$ given λ . For $\lambda = 0$ we have $g(t) = h(t)$, $b(t) = 1$ and the algorithm prefers to attribute an increase in $h(t)$ to $g(t)$ and shorten the lengthscales, while for $\lambda \rightarrow \infty$ the algorithm prefers to scale up the RKHS norm. The assumptions in Corollary 3 hold for any $\lambda \in (0, \infty)$ if $h(t)$ grows unbounded. Moreover, we have that $g(t)^d \leq h(t)$ and $b(t) \leq h(t)$.

Reference regret While any function $h(t)$ that grows unbounded and renders the cumulative regret in Theorem 1 sublinear makes our method to converge to the optimum eventually, we want to ensure that our method performs well in finite time too. For fixed hyperparameters with $h(t) = 1$, which implies $g(t) = b(t) = 1$, our algorithm reduces to GP-UCB with hyperparameters θ_0 and B_0 and the regret bound term $\sqrt{C_1 \beta_t \gamma_t(\theta_0)}$ is sublinear, which is illustrated by the bottom curve in Fig. 3b. However, this does not imply no-regret if hyperparameters are misspecified as in Fig. 2a, since the first term in Theorem 1 is unbounded in this case. To avoid this, we must increase the scaling factor $h(t)$ to consider larger function classes.

We propose to define a sublinear reference regret $p(t)$, see Fig. 3b, and to scale $h(t)$ to match an estimate of the regret with respect to the current hyperparameters to this reference. As GP-UCB converges, the regret estimate with respect to the current hyperparameters levels off and drops below the reference $p(t)$. In these cases, we increase $h(t)$ to consider larger function classes and explore further. The choice of $p(t)$ thus directly specifies the amount of additional regret one is willing to incur for exploration. Specifically, given a regret estimate $\bar{R}_t(h)$ that depends on the data collected so far and the selected scaling h , we obtain $h(t)$ from matching the reference, $\bar{R}_t(h) = p(t)$, as

$$h^*(t) = \bar{R}_t^{-1}(p(t)), \quad h(t) = \max(h^*(t), h(t-1)). \quad (14)$$

Here we explicitly enforce that $h(t)$ must be an increasing function. In the following, we consider estimators \bar{R}_t that are increasing functions of h , so that (14) can be solved efficiently via a line search.

Whether choosing $h(t)$ according to (14) leads to a sublinear function depends on the regret estimator \bar{R}_t . However, it is always possible to upper bound the $h(t)$ obtained from (14) by a fixed sublinear function. This guarantees sublinear regret eventually. In the following, we consider two estimators that upper bound the cumulative regret experienced so far with respect to the hyperparameters suggested by $h(t)$.

Regret bound As a first estimator for the cumulative regret, we consider the regret bound on R_t in (11). We focus on the Gaussian kernel, but the arguments transfer directly to the case of the Matérn kernel. The term $\sqrt{C_1 t \beta_t I_{\theta_t}(\mathbf{y}_t; f)}$ bounds the regret with respect to the current function class specified by θ_t . In addition to the direct dependence on $b(t)g(t)^d$ in β_t , the regret bound also depends on $g(t)$ implicitly through the mutual information $I_{\theta_t}(\mathbf{y}_t; f)$, where $\theta_t = \theta_0/g(t)$. To make the dependence on $g(t)$ more explicit, we use Theorem 2 and rewrite the mutual information as $(g(t)/g(t-1))^d I_{\theta_{t-1}}(\mathbf{y}_t; f)$ instead. Note that the scaling factor was derived for γ_t , but remains a good indicator of increase in mutual information in practice. With this replacement we use

$$\bar{R}_t(h) = \sqrt{C_1 t \beta_t(b(t), g(t)) g(t)^d I_{\theta_{t-1}}(\mathbf{y}_t; f)} \quad (15)$$

to estimate the regret, where the term $\beta_t(b, g)$ is as in Theorem 1, but with the mutual information similarly replaced with the explicit dependence on $g(t)$. Solving (14) together with (15) is computationally efficient, since computing \bar{R}_t does not require inverting the kernel matrix.

One step predictions While (15) is fast to compute, it requires us to know the dependence of $\gamma_t(\theta_t)$ on $h(t)$. Deriving analytic bounds can be infeasible for many kernels. As an alternative, we estimate the regret one-step ahead directly. In particular, if the considered function class is sufficiently large and our confidence intervals hold at all time steps $t > 0$, then the one-step ahead cumulative regret R_{t+1} for our algorithm at iteration t is bounded from above by

$$\bar{R}_t = 2 \sum_{j=1}^t \beta_j^{1/2} \sigma_j(\mathbf{x}_{j+1}), \quad (16)$$

where each β_t and σ_t is based on the corresponding hyperparameters θ_t . In Theorem 1, R_{t+1} is further upper-bounded by (11). The regret estimate in (16) depends on \mathbf{x}_{t+1} , which is the next input that would be evaluated based on the UCB criterion with GP hyperparameters scaled according to $h(t)$. As the hyperparameters for previous iterations are fixed, the only term that depends on $h(t)$ is the bound on the instantaneous regret, $r_t \leq 2\beta_t \sigma_t(\mathbf{x}_{t+1})$. Unlike (15), (16) is not able to exploit the known dependence of γ_t on $h(t)$, so that it cannot reason about the long-term effects of changing $h(t)$. This means that, empirically, the cumulative regret may overshoot the reference regret, only to settle below it later.

Scaling $h(t)$ according to (16) provides an interesting perspective on the method by Wang and de Freitas (2014). They decrease the kernel lengthscales whenever $\sigma_t(\mathbf{x}_{t+1}) \leq \kappa$. In our framework, this corresponds to $p(t) = \sum_{j=1}^t 2\beta_j \max(\kappa, \sigma_j(\mathbf{x}_{j+1})) \geq \kappa t$, which is not sublinear. As a consequence, while they ultimately bound the cumulative regret using the smallest possible lengthscale, the choice for $p(t)$ forces too much exploration to achieve sublinear regret before the lower bound is reached. In contrast, if we choose $p(t)$ to be sublinear,

then the function class grows slowly enough to ensure more careful exploration. This allows us to achieve sublinear regret in the case when a lower bound on the hyperparameters is not known.

4.2. Practical Considerations and Discussion

In this section, we discuss additional practical considerations and show how to combine the theoretical results with online inference of the hyperparameters.

Online inference and exploration strategies The theoretical results presented in the previous sections extend to the case where the initial guess θ_0 of the GP’s lengthscale is improved online using any estimator, e.g., with MAP estimation to obtain θ_t^{MAP} . Theoretically, as long as the change in θ_0 is bounded, the cumulative regret increases by at most a constant factor. In practice, this bound can always be enforced by truncating the estimated hyperparameters. Moreover, the scaling induced by online inference can be considered to be part of $g(t)$ according to Lemma 2, in which case the norm bound can be adapted accordingly. In practice, online inference improves performance drastically, as it is often difficult to specify an appropriate relative initial scaling of the lengthscales θ_0 .

In more than one dimension, $d > 1$, there are multiple ways that MAP estimation can be combined with the theoretical results of the paper. The simplest one is to enforce an upper bound on the lengthscales based on $g(t)$,

$$\theta_t = \min(\theta_t^{\text{MAP}}, \theta_0 / g(t)), \quad (17)$$

where the min is taken elementwise. This choice is similar to the one by Wang et al. (2016). If all entries of θ_0 have the same magnitude, this scaling can be understood as encouraging additional exploration in the smoothest direction of the input space first. This often makes sense, since MAP estimates tend to assume functions that are too smooth, see Fig. 1. However, it can be undesirable in the case when the true function only depends on a subset of the inputs. In these cases, the MAP estimate would correctly eliminate these inputs from the input space by assigning long lengthscales, but the scaling in (17) would encourage additional exploration in these directions first. Note that eventually exploring the entire input space is unavoidable to avoid getting stuck in local optima (Bull, 2011).

An alternative approach is to instead scale down the MAP estimate directly,

$$\theta_t = \theta_t^{\text{MAP}} / \max(g(t), 1). \quad (18)$$

This scaling can be understood as evenly encouraging additional exploration in all directions. While (18) also explores in directions that have been eliminated by the MAP estimate, unlike (17) it simultaneously explores all directions relative to the MAP estimate. From a theoretical point of view, the choice of exploration strategy does not matter, as in the limit as $t \rightarrow \infty$ all lengthscales approach zero. In the one-dimensional case, the two strategies are equivalent. Both strategies use the MAP lengthscales for BO in the nominal case, but the $g(t)$ factor eventually scales down the lengthscales further. This ensures that our method only improves on the empirical performance of BO with MAP estimation.

In practice, maximum likelihood estimates for the inputs are often good enough when the underlying function resembles a sample from a GP. Thus, the approach presented in

this paper is most relevant when the underlying function has some ‘nonstationarity’. In the literature, other approaches to deal with nonstationarity have been proposed. For example, Snoek et al. (2013) scale the input inputs through a beta function and infer its hyperparameters online. Our approach can easily be combined with any such method, as it works on top of any estimate provided by the underlying inference scheme. Moreover, in high-dimensional spaces one can combine our algorithm with methods to automatically identify a low-dimensional subspace of \mathcal{D} (Djolonga et al., 2013; Wang et al., 2016).

In this paper, we have considered the kernel to be fixed, and only adapted the lengthscales and norm bound. However, often the kernel structure itself is a critical hyperparameter (Duvenaud et al., 2011). The strategy presented in this paper could be used to add rougher kernels over time or, for example, to adapt the ν input of the Matérn kernel, which determines its roughness.

Confidence intervals Empirically, β_t is often set to a constant rather than using the theoretical bounds in Lemma 1, which leads to (point-wise) confidence intervals when f is sampled from a GP model. In particular, typically measurement data is standardized to be zero mean and unit variance and β_t is set to two or three. This often works well in practice, but does not provide any guarantees. However, if one were to believe the resulting confidence bounds, our method can be used to avoid getting stuck in local optima, too. In this case one can set $h(t) = g(t)$ and apply our method as before.

General discussion Knowing how the sample complexity of the underlying BO algorithm depends on the lengthscales also has implications in practice. For example, Wang et al. (2016) and Wabersich and Toussaint (2016) suggest to scale down the lengthscales by a factor of 2 and roughly 1.1, respectively, although not at every iteration. As shown in Sec. 4, this scales the regret bound by a factor of g^d , which quickly grows with the number of dimensions. Exponentiating their factors with $1/d$ is likely to make their approaches more robust when BO is used in high-dimensional input spaces \mathcal{D} .

Lastly, in a comparison of multiple BO algorithms (acquisition functions) on a robotic platform, Calandra et al. (2014) conclude that the GP-UCB algorithm shows the best empirical performance for their problem. They use the theoretical version of the algorithm by Srinivas et al. (2012), in which β_t grows with an additional factor of $\mathcal{O}(\sqrt{\log(t^2)})$ relative to Lemma 1. In our framework with the bounds in Lemma 1, this is equivalent to scaling up the initial guess for the RKHS norm bound for f by the same factor at every iteration, which increases the function class considered by the algorithm over time. We conjecture that this increase of the function class over time is probably responsible for pushing the MAP estimate of the lengthscales out of the local minima, which in turn led to better empirical performance.

5. Experiments

In this section, we evaluate our proposed method on several benchmark problems. As baselines, we consider algorithms based on the UCB acquisition function. We specify a strong gamma prior that encourages short lengthscales, and consider both maximum a posteriori (MAP) point-estimates of the hyperparameters and a Hamiltonian Monte Carlo (HMC) approach that samples from the posterior distribution of the hyperparameters and

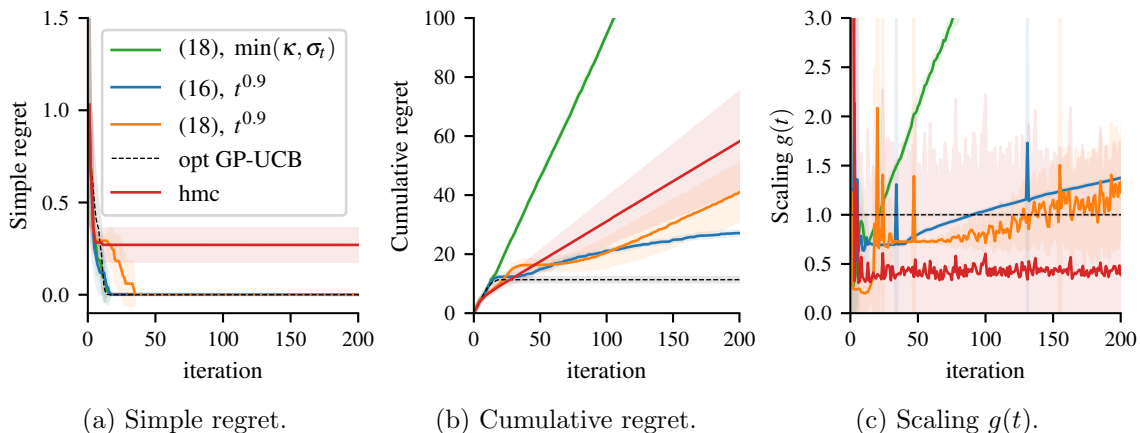


Figure 4: Mean and standard deviation of the empirical simple and cumulative regret over ten different random initializations for the function in Fig. 2. The HMC baseline (red) gets stuck in a local optimum and obtains constant regret in Fig. 4a. GP-UCB with the true hyperparameters (gray dashed) obtains the lowest cumulative regret in Fig. 4b. However, our methods (orange/blue) increase the function class over time, see Fig. 4c, and thus obtain sublinear regret without knowing the true hyperparameters.

marginalizes them out. Unless otherwise specified, the initial lengthscales are set to $\theta_0 = \mathbf{1}$, the initial norm bound is $B_0 = 2$, the confidence bounds hold with probability at least $\delta = 0.9$, and the tradeoff factor between $b(t)$ and $g(t)$ is $\lambda = 0.1$.

We follow several best-practices in BO to ensure a fair comparison with the baselines. We rescale the input space \mathcal{D} to the unit hypercube in order to ensure that both the initial lengthscales and the prior over lengthscales are reasonable for different problems. As is common in practice, the comparison baselines use the empirical confidence intervals suggested in Sec. 4.2, instead of the theoretical bounds in Lemma 1 that are used for our method. Lastly, we initialize all GPs with 2^d measurements that are collected uniformly at random within the domain \mathcal{D} . To measure performance, we use the cumulative regret that has been the main focus of this paper. In addition, we evaluate the different methods in terms of simple regret, which is the regret of the best inputs evaluated so far, $\max_{x \in \mathcal{D}} f(x) - \max_{t' \leq t} f(\mathbf{x}_{t'})$. This metric is relevant when costs during experiments do not matter and BO is only used to determine high-quality inputs by the end of the optimization procedure.

5.1. Synthetic Experiments

Example function We first evaluate all proposed methods on the example function in Fig. 2, which lives inside the RKHS associated with a Gaussian kernel with $\theta = 0.1$ and has norm $\|f\|_{k_\theta} = 2$. We evaluate our proposed method for the sublinear reference function $p(t) = t^{0.9}$ together with maximum a posteriori hyperparameter estimation. We compare against both GP-UCB with the fixed, correct hyperparameters and HMC hyperparameter estimation. Additionally, we consider a modified variant of the method suggested by Wang and de Freitas (2014), see Sec. 4.1. Rather than scaling the lengthscales by a fixed constant, we conduct a line search to find the smallest possible scaling factor that ren-




Figure 5: Simple and cumulative regret over 10 random seeds for samples from a GP with bounded RKHS norm. The GP-UCB algorithm with misspecified hyperparameters (magenta) fails to converge given only a wrong choice of B_0 . In contrast, our methods (blue/orange) converge even though θ_0 is misspecified in addition.

ders $\sigma_t(\mathbf{x}_{t+1}) \geq \kappa = 0.1$. This is the most conservative variant of the algorithm. Note that we do not know a lower bound on the hyperparameters and therefore do not enforce it.

The results of the experiments are shown in Fig. 4. The simple regret plot in Fig. 4a shows that all methods based on hyperparameter adaptation evaluate close-to-optimal inputs eventually, and do so almost as quickly as GP-UCB based on the true hyperparameters (black, dashed). However, the method based on HMC hyperparameter estimation (red) considers functions that are too smooth and gets stuck in local optima, as in Fig. 2. This can also be seen in Fig. 4c, which plots the effective scaling $g(t)$ based on the combination of Bayesian hyperparameter estimation and hyperparameter adaptation through $h(t)$. The HMC hyperparameters consistently over-estimate the lengthscales by a factor of roughly two. In contrast, while the MAP estimation leads to the wrong hyperparameters initially, the adaptation methods in (15) and (16) slowly increase the function class until the true lengthscales are found eventually. It can be seen that the one step estimate (16) (orange) is more noisy than the upper bound in (15) (blue).

While all adaptation methods determine good inputs quickly according to the simple regret, they perform differently in terms of the cumulative regret in Fig. 4b. As expected, the HMC method (red line) converges to a local optimum and experiences constant regret increase equal to the simple regret at every time step. The modified method of Wang and de Freitas (2014) (green line) expands the function class too aggressively and also achieves constant regret. Empirically, their method always explores and never repeatedly evaluates close-to-optimal inputs that would decrease cumulative regret. While the method works well in terms of simple regret, without a lower bound on the hyperparameters it never converges to sublinear regret. As expected from Theorem 1, GP-UCB based on the optimal hyperparameters achieves the lowest cumulative regret. Our two methods expand the function class over time, which allows them to converge to close-to-optimal inputs, even though MAP estimation estimates the hyperparameters wrongly initially. While the regret is sublinear, the additional exploration caused by $g(t)$ means that the cumulative regret is larger. This is the additional cost we incur for not knowing the hyperparameters in advance.

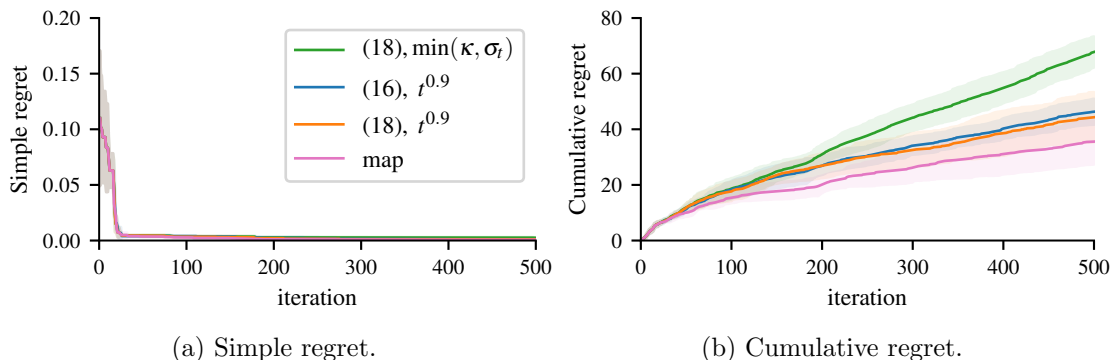


Figure 6: Simple and cumulative regret over 5 random seeds for a logistic regression problem. All methods determine close-to-optimal parameters. However, our methods explore more to counteract misspecified hyperparameters.

Samples from a GP As a second experiment, we compare GP-UCB to A-GP-UCB on samples drawn from a GP when the norm bound B_0 is misspecified. Samples from a GP are not contained in the RKHS. To avoid this technical issue, we sample function values from the posterior GP at only a finite number of discrete gridpoints and interpolate between them using the kernel with the correct lengthscales θ . We rescale these functions to have RKHS norm of $B = 4$, but use $B_0 = 0.25$ as an initial guess for both BO algorithms and do not use any hyperparameter estimation. Even though we use the correct kernel lengthscales for GP-UCB, $\theta_0 = \theta = 0.1$, this discrepancy means that the true function is not contained in the initial confidence intervals. As before, for our method we use the reference regret $p(t) = t^{0.9}$ and additionally misspecify the lengthscales, $\theta_0 = 1$.

The results are shown in Fig. 5. GP-UCB with the correct hyperparameters (black, dashed) obtains the lowest cumulative regret. However, it fails to converge when hyperparameters are misspecified (magenta), since the confidence intervals are too small to encourage any exploration. In contrast, our methods (blue/orange) converge to close-to-optimal inputs as in the previous example.

5.2. Logistic Regression Experiment

Lastly, we use our method to tune a logistic regression problem on the MNIST data set (LeCun, 1998). As in the experiment in Klein et al. (2016), we consider four training inputs: the learning rate, the l_2 regularization constant, the batch size, and the dropout rate. We use the validation loss as the optimization objective.

The results are shown in Fig. 6. Even though the input space is fairly high-dimensional with $d = 4$, all algorithms determine close-to-optimal inputs quickly. In particular, MAP estimation determines that both the dropout rate and the batch size do not influence the validation loss significantly. Since the theoretical results in A-GP-UCB are compatible with MAP estimation, our approach achieves the same empirical performance, but has theoretical worst-case regret bounds. After convergence, the BO baselines repeatedly evaluate the same inputs, without gaining any new information. In contrast, our method continues to explore in order to potentially find better inputs. While it does not occur in this case, this allows us

to be more confident that the global optimum has been identified as t increases. For standard BO methods, there is no guarantee of convergence with misspecified hyperparameters.

6. Conclusion and Future Work

We introduced A-GP-UCB, a BO algorithm that is provably no-regret when hyperparameters are unknown. Our method adapts the hyperparameters online, which causes the underlying BO algorithm to consider larger function spaces over time. Eventually, the function space is large enough to contain the true function, so that our algorithm provably converges. We evaluated our method on several benchmark problems, confirming that, on the one hand, it provably converges even in cases where standard BO algorithms get stuck in local optima, and, on the other hand, enjoys competitive performance as standard BO algorithms that do not have theoretical guarantees in this setting.

The main idea behind our analysis is that adapting the hyperparameters increases the cumulative regret bound, but we do so slowly enough to converge eventually. This idea is fairly general and could also be applied to other no-regret algorithms. Another potential future direction is to investigate alternative strategies to select the scaling factors $b(t)$ and $g(t)$ and consider adapting other parameters such as the kernel structure.

Acknowledgments

This research was supported in part by SNSF grant 200020_159557, ERC grant no. 815943, NSERC grant RGPIN-2014-04634, the Vector Institute, and an Open Philanthropy Project AI fellowship. We would like to thank Johannes Kirschner for valuable discussions.

References

- Yasin Abbasi-Yadkori. *Online learning of linearly parameterized control problems*. PhD thesis, 2012.
- Hany Abdelrahman, Felix Berkenkamp, and Andreas Krause. Bayesian optimization for maximum power point tracking in photovoltaic power plants. In *2016 European Control Conference (ECC)*, pages 2078–2083, 2016.
- James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012.
- Felix Berkenkamp, Andreas Krause, and Angela P. Schoellig. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *arXiv:1602.04450 [cs.RO]*, 2016.
- Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: a unified approach to Bayesian optimization and level-set estimation. In *Advances in Neural Information Processing Systems 29*, pages 1507–1515. Curran Associates, Inc., 2016.

- Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599 [cs]*, 2010.
- Adam D. Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(Oct):2879–2904, 2011.
- Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. An experimental comparison of Bayesian optimization for bipedal locomotion. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1951–1958, 2014.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 844–853. PMLR, 2017.
- Andreas Christmann and Ingo Steinwart. *Support Vector Machines*. Information Science and Statistics. Springer, New York, NY, 2008.
- Josip Djolonga, Andreas Krause, and Volkan Cevher. High-dimensional Gaussian process bandits. In *Advances in Neural Information Processing Systems 26*, pages 1025–1033, 2013.
- Audrey Durand, Odalric-Ambrym Maillard, and Joelle Pineau. Streaming kernel regression with provably adaptive mean, variance, and regularization. *Journal of Machine Learning Research*, 19(17):1–34, 2018.
- David K. Duvenaud, Hannes Nickisch, and Carl Edward Rasmussen. Additive Gaussian processes. In *Advances in Neural Information Processing Systems 24*, pages 226–234, 2011.
- Peter Frazier, Warren Powell, and Savas Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS Journal on Computing*, 21(4):599–613, 2009.
- Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D. Sculley. Google vizier: a service for black-box optimization. In *Proc. of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’17, pages 1487–1495, New York, NY, USA, 2017. ACM.
- I. S. Gradshteyn, I. M. Ryzhik, and Alan Jeffrey. *Table of integrals, series, and products*. Academic Press, Amsterdam, Boston, 7th ed edition, 2007.
- Philipp Hennig and Christian J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(1):1809–1837, 2012.
- José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 918–926. Curran Associates, Inc., 2014.

- Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K. Sriperumbudur. Gaussian processes and kernel methods: a review on connections and equivalences. *arXiv:1807.02582 [stat.ML]*, 2018.
- Johannes Kirschner and Andreas Krause. Information directed sampling and bandits with heteroscedastic noise. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 358–384. PMLR, 2018.
- Aaron Klein, Stefan Falkner, Jost Tobias Springenberg, and Frank Hutter. Bayesian neural network for predicting learning curves. In *NIPS 2016 Bayesian Neural Network Workshop*, 2016.
- Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Daniel J. Lizotte, Tao Wang, Michael H. Bowling, and Dale Schuurmans. Automatic gait optimization with Gaussian process regression. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, volume 7, pages 944–949, 2007.
- Jonas Mockus. *Bayesian approach to global optimization: theory and applications*. Springer Science & Business Media, 2012.
- Jonas Mockus, Vytautas Tiesis, and Antanas Zilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2:117–129, 1978.
- Carl Edward Rasmussen and Christopher K.I Williams. *Gaussian processes for machine learning*. MIT Press, Cambridge MA, 2006.
- Binxin Ru, Michael A. Osborne, Mark Mcleod, and Diego Granziol. Fast information-theoretic Bayesian optimisation. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4384–4392. PMLR, 2018.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 1583–1591. Curran Associates, Inc., 2014.
- M. W. Seeger, S. M. Kakade, and D. P. Foster. Information consistency of nonparametric Gaussian process methods. *IEEE Transactions on Information Theory*, 54(5):2376–2382, 2008.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2951–2959. Curran Associates, Inc., 2012.

- Jasper Snoek, Kevin Swersky, Richard S. Zemel, and Ryan P. Adams. Input warping for Bayesian optimization of non-stationary functions. In *NIPS Workshop on Bayesian Optimization*, 2013.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- V. Vapnik. Principles of risk minimization for learning theory. In J. E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems 4*, pages 831–838. Morgan-Kaufmann, 1992.
- Kim Peter Wabersich and Marc Toussaint. Advancing Bayesian optimization: the mixed-global-local (mgl) kernel and length-scale cool down. *arXiv:1612.03117 [cs, stat]*, 2016.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3627–3635. PMLR, 2017.
- Ziyu Wang and Nando de Freitas. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. *arXiv:1406.7758 [cs, stat]*, 2014.
- Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando de Freitas. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387, 2016.

Appendix A. Proof of Main Theorem

Lemma 3 *Let $f \in \mathcal{H}_{\theta_{t^*}}$ with $\|f\|_{\theta_{t^*}} \leq B_{t^*}$. Then, for any monotonically increasing functions $g(t) \geq 1$ and $b(t) \geq 1$ and for all $t \geq t^*$: $f \in \mathcal{H}_{\theta_t}$ with $\|f\|_{\theta_t} \leq B_t$*

Proof Lemma 2 together with monotonicity of g yields $\mathcal{H}_{\theta_t} \supseteq \mathcal{H}_{\theta_{t^*}}$ so that $f \in \mathcal{H}_{\theta_t}$ and

$$\|f\|_{\theta_t} \leq \prod_{1 \leq i \leq d} \frac{[\theta_{t^*}]_i}{[\theta_t]_i} \|f\|_{\theta_{t^*}} \leq \frac{g(t)^d}{g(t^*)^d} B_{t^*} = \frac{g(t)^d}{g(t^*)^d} g(t^*)^d b(t^*) B_0 = g(t)^d b(t^*) B_0 \leq B_t$$

■

Lemma 4 *Under the assumptions of Lemma 1, let θ_t be a predictable sequence of kernel hyperparameters such that $\|f\|_{k_{\theta_t}} \leq B_t$ and let the GP predictions μ_t and σ_t use the prior covariance k_{θ_t} . If $\beta_t^{1/2} = B_t + 4\sigma \sqrt{I_{\theta_t}(\mathbf{y}_t; f) + 1 + \ln(1/\delta)}$, then $|f(\mathbf{x}) - \mu_t(\mathbf{x})| \leq \beta_t^{1/2} \sigma_t(\mathbf{x})$ holds for all $\mathbf{x} \in D$ and iterations $t \geq 0$ jointly with probability at least $1 - \delta$.*

Proof The proof is the same as the one by Abbasi-Yadkori (2012); Chowdhury and Gopalan (2017), except that the kernel is time-dependent. ■

We are now ready to prove the main result:

Proof [Theorem 1] We split the regret bound into two terms, $R_t = t_0 r_c + r_s(t)$. In the initial rounds, where either $B_t \leq g(t)^d B_0$ or $\max_i [\theta]_i / [\theta]_0 > 1$, the regret is trivially bounded by $r_t \leq 2\|f\|_\infty \leq 2\|f\|_\theta \leq B$. Thus $r_c \leq 2B$. Let $t_0 \in (0, \infty]$ be the first iteration such that $f \in \mathcal{H}_{\theta_{t_0}}$ with $\|f\|_{\theta_{t_0}} \leq B_{t_0}$. From Lemma 3, we have that $f \in \mathcal{H}_{\theta_t}$ with $\|f\|_{\theta_t} \leq B_t$ for all $t \geq t_0$. Thus we can use Lemma 4 to conclude $|f - \mu_t(\mathbf{x})| \leq \beta_t^{1/2} \sigma_t(\mathbf{x})$ for all $\mathbf{x} \in D$ and $t \geq t_0$ jointly with probability at least $(1 - \delta)$. We use Lemmas 5.2-5.4 in Srinivas et al. (2012) to conclude that the second stage has a regret bound of $r_s^2(t) \leq C_1 \beta_t I(\mathbf{y}_t; f)$, which concludes the proof. \blacksquare

Appendix B. Bound on the information capacity γ_t

Theorem 4 (Theorem 8 in Srinivas et al. (2012)) *Suppose that $D \subset \mathbb{R}^d$ is compact, and $k(\mathbf{x}, \mathbf{x}')$ is a covariance function for which the additional assumption of Theorem 2 in Srinivas et al. (2012) hold. Moreover, let $B_k(T_*) = \sum_{s > T_*} \lambda_s$, where $\{\lambda_s\}$ is the operator spectrum of k with respect to the uniform distribution over D . Pick $\tau > 0$, and let $n_T = C_4 T^\tau (\log T)$ with $C_4 = 2\mathcal{V}(D)(2\tau + 1)$. Then, the following bound holds true:*

$$\gamma_T \leq \frac{1/2}{1 - e^{-1}} \max_{r \in \{1, \dots, T\}} T_* \log \left(\frac{r n_T}{\sigma^2} \right) + C_4 \sigma^{-2} \left(1 - \frac{r}{T}\right) (B_k(T_*) T^{\tau+1} + 1) \log T + \mathcal{O}(T^{1-\frac{\tau}{d}}). \quad (19)$$

Theorem 4 allows us to bound γ_t through the operator spectrum of the kernel with respect to the uniform distribution. We now consider this quantity for two specific kernels.

B.1. Bounds for the Squared Exponential Kernel

Lemma 5 *For all $x \in [0, x_{\max}^2]$ it holds that $\log(1 + x^2) \geq \frac{\log(1 + x_{\max}^2)}{x_{\max}^2} x^2$*

In this section, we use Theorem 4 to obtain concrete bounds for the Gaussian kernel. From Seeger et al. (2008), we obtain a bound on the eigenspectrum that is given by

$$\lambda_s \leq c B^{s^{1/d}}, \text{ where } c = \sqrt{\frac{2a}{A}}, \quad b = \frac{1}{2\theta_t^2}, \quad B = \frac{b}{A}, \quad \text{and } A = a + b + \sqrt{a^2 + 2ab}.$$

The constant $a > 0$ parameterizes the distribution $\mu(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}, (4a)^{-1} \mathbf{I}_d)$. As a consequence of $\theta_t > 0$, we have that $b \geq 0$, $0 < B < 1$, $c > 0$, and $A > 0$. In the following, we bound the eigenspectrum. The steps follow the outline of Seeger et al. (2008), but we provide more details and the dependence on the lengthscales θ_t is made explicit:

$$\begin{aligned} B_k(T_*) &= \sum_{s > T_*} \lambda_s \leq c \sum_{s \geq T_*+1} B^{s^{1/d}} = c \sum_{s \geq T_*+1} \exp \log(B^{s^{1/d}}) = c \sum_{s \geq T_*+1} \exp(s^{1/d} \log B), \\ &= c \sum_{s \geq T_*+1} \exp(-s^{1/d} \alpha) \leq c \int_{T_*}^{\infty} \exp(-\alpha s^{1/d}) ds, \end{aligned}$$

where $\alpha = -\log B$. Now substitute $s = \phi(t) = (t/\alpha)^d$. Then $ds = \frac{dt^{d-1}}{\alpha} dt$ and

$$B_k(T_*) \leq c \int_{\alpha T_*^{1/d}}^{\infty} \exp(-t) \frac{dt^{d-1}}{\alpha} dt = cd\alpha^{-d} \Gamma(d, \alpha T_*^{1/d}),$$

where $\Gamma(d, \beta) = \int_{\beta}^{\infty} e^{-t} t^{d-1} dt = (d-1)! e^{-\beta} \sum_{k=0}^{d-1} \beta^k / k!$ for $d \in \mathbb{N}$ as in Gradshteyn et al. (2007, (8.352.4)). Then, with $\beta = \alpha T_*^{1/d}$,

$$B_k(T_*) \leq cd\alpha^{-d} (d-1)! e^{-\beta} \sum_{k=0}^{d-1} \beta^k / k! = c(d!) \alpha^{-d} e^{-\beta} \sum_{k=0}^{d-1} (k!)^{-1} \beta^k.$$

Before we bound the information gain, let us determine how α^{-d} and c depend on the lengthscales. In particular, we want to quantify their upper bounds in terms of $g(t)$.

$$\alpha^{-d} = \log^{-d}(1/B) = \log^{-d}(2\theta_t^2 A) = \log^{-d} \left(1 + 2\theta_t^2 a + 2\theta_t \sqrt{a^2 + \frac{a}{\theta_t^2}} \right) \quad (20)$$

$$\leq \log^{-d} (1 + 2\theta_t^2 a) \leq \left(\frac{\log(1 + 2\theta_0^2 a)}{2\theta_0^2 a} 2\theta_t^2 a \right)^{-d} \quad \text{by Lemma 5} \quad (21)$$

$$= \mathcal{O}(\theta_t^{-2d}) = \mathcal{O}(g^{2d}(t)), \quad (22)$$

where (21) follows from Lemma 5, since $g(t) \geq 1$ for all $t > 0$. Similarly,

$$c = \left(\frac{2a}{a + \frac{1}{2\theta_t^2} + \sqrt{a^2 + \frac{a}{\theta_t^2}}} \right)^{d/2} \leq \left(\frac{2a}{\frac{1}{2\theta_t^2}} \right)^{d/2} = (4a\theta_t^2)^{d/2} = \mathcal{O}(g(t)^{-d}). \quad (23)$$

As in Srinivas et al. (2012), we choose $T_* = (\log(Tn_T)/\alpha)^d$, so that $\beta = \log(Tn_T)$ and therefore does not depend on g_t . Plugging into (19), the first term of (19) dominates and

$$\gamma_T = \mathcal{O} \left(\left[\log(T^{d+1}(\log T)) \right]^{d+1} c \alpha^{-d} \right)^{d/2} = \mathcal{O} \left((\log T)^{d+1} g(t)^d \right). \quad (24)$$

B.2. Matérn kernel

Following the proof for Theorem 2 in the addendum to Seeger et al. (2008), we have that

$$\lambda_s^{(T)} \leq C(1 + \delta) s^{-(2\nu+d)/d} \quad \forall s \geq s_0, \quad (25)$$

For the leading constant we have $C = C_3^{(2\nu+d)/d}$ with $\alpha = \frac{2\pi\theta_t}{\sqrt{2\nu}}$. Hiding terms that do not depend on α and therefore $g(t)$, we have

$$C_t(\alpha, \nu) = \frac{\Gamma(\nu + d/2)}{\pi^{d/2} \Gamma(\nu)} \alpha^d = \mathcal{O}(g(t)^{-d}) \quad c_1 = \frac{1}{(2\pi)^d C_t(\alpha, \nu)} = \mathcal{O}(g(t)^d)$$

$$C_2 = \frac{\alpha^{-d}}{2^d \pi^{d/2} \Gamma(d/2)} = \mathcal{O}(g(t)^d) \quad C_3 = C_2 \frac{2\tilde{C}}{d} c_1^{\frac{-d}{2\nu+d}} = \mathcal{O}(g(t)^d g(t)^{\frac{-d^2}{2\nu+d}}) = \mathcal{O}(g(t)^d),$$

so that $C = \mathcal{O}(g(t)^{2\nu+d})$. The second term in C_3 must be over-approximated as a consequence of the proof strategy. It follows that $B_k(T_*) = \mathcal{O}(g(t)^{2\nu d} T_*^{1-(2\nu+d)/d})$ and, as in Srinivas et al. (2012), that $\gamma_T = \mathcal{O}(T^{\frac{d(d+1)}{2\nu+d(d+1)}} (\log T) g(t)^{2\nu d})$.