On Lower Bounds for Standard and Robust Gaussian Process Bandit Optimization

Xu Cai [1] Jonathan Scarlett [1 2]

Abstract

In this paper, we consider algorithm-independent lower bounds for the problem of black-box optimization of functions having a bounded norm is some Reproducing Kernel Hilbert Space (RKHS), which can be viewed as a non-Bayesian Gaussian process bandit problem. In the standard noisy setting, we provide a novel proof technique for deriving lower bounds on the regret, with benefits including simplicity, versatility, and an improved dependence on the error probability. In a robust setting in which every sampled point may be perturbed by a suitably-constrained adversary, we provide a novel lower bound for deterministic strategies, demonstrating an inevitable joint dependence of the cumulative regret on the corruption level and the time horizon, in contrast with existing lower bounds that only characterize the individual dependencies. Furthermore, in a distinct robust setting in which the final point is perturbed by an adversary, we strengthen an existing lower bound that only holds for target success probabilities very close to one, by allowing for arbitrary success probabilities above $\frac{2}{3}$.

1. Introduction

The use of Gaussian process (GP) methods for black-box function optimization has seen significant advances in recent years, with applications including hyperparameter tuning, robotics, molecular design, and many more. On the theoretical side, a variety of algorithms have been developed with provable regret bounds (Srinivas et al., 2010; Bull, 2011; Contal et al., 2013; Wang et al., 2016; Bogunovic et al., 2016; Wang & Jegelka, 2017; Janz et al., 2020), and algorithm-independent lower bounds have been given in several settings of interest (Bull, 2011; Scarlett et al., 2017; Scarlett, 2018; Chowdhury & Gopalan, 2019; Wang et al., 2020).

These theoretical works can be broadly categorized into one of two types: In the *Bayesian setting*, one adopts a Gaussian process prior according to some kernel function, whereas in the *non-Bayesian setting*, the function is assumed to lie in some Reproducing Kernel Hilbert Space (RKHS) and be upper bounded in terms of the corresponding RKHS norm.

In this paper, we focus on the non-Bayesian setting, and seek to broaden the existing understanding of algorithm-independent lower bounds on the regret, which have received significantly less attention than upper bounds. Our main contributions are briefly summarized as follows:

- In the standard noisy GP optimization setting, we provide an alternative proof strategy for the existing lower bounds of (Scarlett et al., 2017), which we believe to be of significant importance in itself due to the lack of techniques in the literature. We additionally show that our approach strengthens the dependence on the error probability, and give scenarios in which our approach is simpler and/or more versatile.
- We provide a novel lower bound for a robust setting in which the sampled points are adversarially corrupted (Bogunovic et al., 2020). Our bound demonstrates that the cumulative regret of any deterministic algorithm must incur a certain *joint* dependence on the corruption level and time horizon, strengthening results from (Bogunovic et al., 2020) stating that certain *separate* dependencies are unavoidable.
- We provide an improvement on an existing lower bound for a distinct robust setting (Bogunovic et al., 2018a), in which the *final point returned* is perturbed by an adversary. While the lower bound of (Bogunovic et al., 2018a) shows that a certain number of samples is needed to attain a certain level of regret with probability very close to one, we show that the same number of samples (up to constant factors) is

[1]Department of Computer Science, National University of Singapore [2]Department of Mathematics & Institute of Data Science, National University of Singapore. Correspondence to: Xu Cai <caix@u.nus.edu>, Jonathan Scarlett <scarlett@comp.nus.edu.sg>.

Proceedings of the 38th International Conference on Machine Learning, PMLR 139, 2021.

required just to succeed with probability at least $\frac{2}{3}$.

The relevant existing results are highlighted throughout the paper, with further details in Appendix A.

2. Problem Setup

The three problem settings considered throughout the paper are formally described as follows. We additionally informally summarize the existing lower bounds in each of these settings, with formal statements given in Appendix A along with existing upper bounds. The existing and new bounds are summarized in Table 1 below.

2.1. Standard Setting

Let f be a function on the compact domain $D = [0,1]^d$; by simple re-scaling, the results that we state readily extend to other rectangular domains. The smoothness of f is modeled by assuming that $\|f\|_k \le B$, where $\|\cdot\|_k$ is the RKHS norm associated with some kernel function $k(\mathbf{x},\mathbf{x}')$ (Rasmussen, 2006). The set of all functions satisfying $\|f\|_k \le B$ is denoted by $\mathcal{F}_k(B)$, and $\mathbf{x}^*$ denotes an arbitrary maximizer of f.

At each round indexed by t, the algorithm selects some $\mathbf{x}_t \in D$, and observes a noisy sample $y_t = f(\mathbf{x}_t) + z_t$. Here the noise term is distributed as $N(0,\sigma^2)$, with $\sigma^2 > 0$ and independence between times.

We measure the performance using the following two widespread notions of regret:

- **Simple regret:** After T rounds, an additional point $\mathbf{x}^{(T)}$ is returned, and the simple regret is given by $r(\mathbf{x}^{(T)}) = f(\mathbf{x}^*) - f(\mathbf{x}^{(T)})$.
- **Cumulative regret:** After T rounds, the cumulative regret incurred is $R_T = \sum_{t=1}^T r_t$, where $r_t = f(\mathbf{x}^*) - f(\mathbf{x}_t)$.

As with the previous work on noisy lower bounds (Scarlett et al., 2017), we focus on the squared exponential (SE) and Matérn kernels, defined as follows with length-scale $l > 0$ and smoothness $\nu > 0$ (Rasmussen, 2006):

$$k_{\text{SE}}(\mathbf{x},\mathbf{x}') = \exp\left(-\frac{r^2_{\mathbf{x},\mathbf{x}'}}{2l^2}\right) \tag{1}$$

$$k_{\text{Matérn}}(\mathbf{x},\mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)}\left(\frac{\sqrt{2\nu}\,r_{\mathbf{x},\mathbf{x}'}}{l}\right)^{\nu} J_\nu\left(\frac{\sqrt{2\nu}\,r_{\mathbf{x},\mathbf{x}'}}{l}\right), \tag{2}$$

where $r_{\mathbf{x},\mathbf{x}'} = \|\mathbf{x}-\mathbf{x}'\|$, and J_ν denotes the modified Bessel function.

Existing lower bounds. The results of (Scarlett et al., 2017) are informally summarized as follows:

- Attaining (average or constant-probability) simple regret ϵ requires the time horizon to satisfy $T = \Omega\big(\frac{1}{\epsilon^2}\big(\log\frac{1}{\epsilon}\big)^{d/2}\big)$ for the SE kernel, and $T = \Omega\big(\big(\frac{1}{\epsilon}\big)^{2+d/\nu}\big)$ for the Matérn kernel.
- The (average or constant-probability) cumulative regret is lower bounded according to $R_T = \Omega\big(\sqrt{T(\log T)^{d/2}}\big)$ for the SE kernel, and $R_T = \Omega\big(T^{\frac{\nu+d}{2\nu+d}}\big)$ for the Matérn kernel.

The SE kernel bounds have near-matching upper bounds (Srinivas et al., 2010), and while standard results yield wider gaps for the Matérn kernel, these have been tightened in recent works; see Appendix A for details.

In Sections 4.2–4.5, we will present novel analysis techniques that can both simplify the proofs and strengthen the dependence on the error probability compared to the lower bounds in (Scarlett et al., 2017).[3]

2.2. Robust Setting – Corrupted Samples

In the robust setting studied in (Bogunovic et al., 2020), the optimization goal is similar, but each sampled point is further subject to adversarial noise; for $t = 1,\dots,T$:

- Based on the previous samples $\{(\mathbf{x}_i,\tilde{y}_i)\}_{i=1}^{t-1}$, the player selects a distribution $\Phi_t(\cdot)$ over D.
- Given knowledge of the true function f, the previous samples $\{(\mathbf{x}_i,y_i)\}_{i=1}^{t-1}$, and the player's distribution $\Phi_t(\cdot)$, an adversary selects a function $c_t(\cdot) : D \to [-B_0, B_0]$, where $B_0 > 0$ is constant.
- The player draws $\mathbf{x}_t \in D$ from the distribution Φ_t, and observes the corrupted sample

$$\tilde{y}_t = y_t + c_t(\mathbf{x}_t), \tag{3}$$

where y_t is the noisy non-corrupted observation $y_t = f(\mathbf{x}_t) + z_t$ as in Section 2.1.

Note that in the special case that $\Phi_t(\cdot)$ is deterministic, the adversary knowing Φ_t also implies knowledge of $\mathbf{x}_t$.

For this problem to be meaningful, the adversary must be constrained. Following (Bogunovic et al., 2020), we assume the following constraint for some corruption level C:

$$\sum_{t=1}^T \max_{\mathbf{x}\in D} |c_t(\mathbf{x})| \le C. \tag{4}$$

When $C = 0$, we reduce to the setup of Section 2.1.

While both the simple regret and cumulative regret could be considered here, we focus entirely on the latter, as it has been the focus of the related existing works (Bogunovic

[3] We are not aware of any way to adapt the analysis of (Scarlett et al., 2017) to obtain a high-probability lower bound that grows unbounded as the target error probability approaches zero.

SE kernel

	Upper Bound	Existing Lower Bound	Our Lower Bound
Standard Cumulative Regret[1,2]	$O^*\left(\sqrt{T(\log T)^{2d}\log\frac{1}{\delta}}\right)$	$\Omega\left(\sqrt{T(\log T)^{d/2}}\right)$	$\Omega^*\left(\sqrt{T(\log T)^{d/2}\log\frac{1}{\delta}}\right)$
Corrupted Samples Cumul. Regret, $\delta=\Theta(1)$	$O^*\left(\overline{R}_T^{\text{std}}+C\sqrt{T(\log T)^d}\right)$	$\Omega\left(\underline{R}_T^{\text{std}}+C\right)$	$\Omega\left(\underline{R}_T^{\text{std}}+C(\log T)^{d/2}\right)$
Corrupted Final Point Time to ϵ-optimality[2]	$O^*\left(\frac{1}{\epsilon^2}\left(\log\frac{1}{\epsilon}\right)^{2d}\log\frac{1}{\delta}\right)$	$\Omega\left(\frac{1}{\epsilon^2}\left(\log\frac{1}{\epsilon}\right)^{\frac{d}{2}}\right)$ **(only for $\delta\le O(\xi^d)$)**	$\Omega\left(\frac{1}{\epsilon^2}\left(\log\frac{1}{\epsilon}\right)^{\frac{d}{2}}\log\frac{1}{\delta}\right)$

Matérn-ν kernel

	Upper Bound	Existing Lower Bound	Our Lower Bound
Standard Cumulative Regret[1]	$O^*\left(T^{\frac{\nu+d}{2\nu+d}}\sqrt{\log\frac{1}{\delta}}\right)$	$\Omega\left(T^{\frac{\nu+d}{2\nu+d}}\right)$	$\Omega\left(T^{\frac{\nu+d}{2\nu+d}}\left(\log\frac{1}{\delta}\right)^{\frac{\nu}{2\nu+d}}\right)$
Corrupted Samples Cumul. Regret, $\delta=\Theta(1)$	$O^*\left(\overline{R}_T^{\text{std}}+CT^{\frac{\nu+d}{2\nu+d}}\right)$	$\Omega\left(\underline{R}_T^{\text{std}}+C\right)$	$\Omega\left(\underline{R}_T^{\text{std}}+C^{\frac{\nu}{d+\nu}}T^{\frac{d}{d+\nu}}\right)$
Corrupted Final Point Time to ϵ-optimality	$O^*\left(\left(\frac{1}{\epsilon}\right)^{\frac{2(2\nu+d)}{2\nu-d}}+\left(\frac{\log\frac{1}{\delta}}{\epsilon^2}\right)^{1+\frac{d}{2\nu}}\right)$ **(only for $d<2\nu$)**	$\Omega\left(\frac{1}{\epsilon^2}\left(\frac{1}{\epsilon}\right)^{d/\nu}\right)$ **(only for $\delta\le O(\xi^d)$)**	$\Omega\left(\frac{1}{\epsilon^2}\left(\frac{1}{\epsilon}\right)^{d/\nu}\log\frac{1}{\delta}\right)$

Table 1. Summary of new and existing regret bounds. T denotes the time horizon, d denotes the dimension, ξ denotes the corruption radius, and δ denotes the allowed error probability. In the middle row, $\overline{R}_T^{\text{std}}$ and $\underline{R}_T^{\text{std}}$ denote upper and lower bounds on the standard cumulative regret. The existing upper and lower bounds are from (Srinivas et al., 2010; Chowdhury & Gopalan, 2017; Scarlett et al., 2017; Bogunovic et al., 2018a; 2020), with the partial exception of the Matérn kernel upper bounds, which are detailed at the end of Appendix A.4. The notation $O^*(\cdot)$ and $\Omega^*(\cdot)$ hides dimension-independent $\log T$ factors, as well as $\log\log\frac{1}{\delta}$ factors.

et al., 2020; 2021; Lykouris et al., 2018; Gupta et al., 2019; Li et al., 2019). See also (Bogunovic et al., 2020, App. C) for discussion on the use of simple regret in this setting.

Existing lower bound. The only lower bound stated in (Bogunovic et al., 2020) states that $R_T=\Omega(C)$ for any algorithm, whereas the upper bound therein essentially amounts to *multiplying* (rather than adding) the uncorrupted regret bound by C. Thus, significant gaps remain in terms of the *joint* dependence on C and T, which our lower bound in Section 4.6 will partially address.

2.3. Robust Setting – Corrupted Final Point

Here we detail a different robust setting, previously considered in (Bogunovic et al., 2018a), in which the samples themselves are only subject to random (non-adversarial) noise, but the *final point* returned may be adversarially perturbed. For a real-valued function $\text{dist}(\mathbf{x},\mathbf{x}')$ and constant ξ, we define the set-valued function

$$\Delta_\xi(\mathbf{x})=\left\{\mathbf{x}'-\mathbf{x}\,:\,\mathbf{x}'\in D\ \text{ and }\ \text{dist}(\mathbf{x},\mathbf{x}')\le\xi\right\}\quad(5)$$

representing the set of perturbations of $\mathbf{x}$ such that the newly obtained point $\mathbf{x}'$ is within a "distance" ξ of $\mathbf{x}$.

We seek to attain a function value as high as possible following the worst-case perturbation within $\Delta_\xi(\cdot)$; in particular, the global robust optimizer is given by

$$\mathbf{x}^*_\xi\in\arg\max_{\mathbf{x}\in D}\min_{\boldsymbol{\delta}\in\Delta_\xi(\mathbf{x})}f(\mathbf{x}+\boldsymbol{\delta}).\quad(6)$$

Then, if the algorithm returns $\mathbf{x}^{(T)}$, the performance is measured by the ξ-*regret*:

$$r_\xi(\mathbf{x})=\min_{\boldsymbol{\delta}\in\Delta_\xi(\mathbf{x}^*_\xi)}f(\mathbf{x}^*_\xi+\boldsymbol{\delta})-\min_{\boldsymbol{\delta}\in\Delta_\xi(\mathbf{x})}f(\mathbf{x}+\boldsymbol{\delta}).\quad(7)$$

We focus our attention on the primary case of interest in which $\text{dist}(\mathbf{x},\mathbf{x}')=\|\mathbf{x}-\mathbf{x}'\|_2$, meaning that achieving low ξ-regret amounts to favoring *broad peaks* instead of narrow ones, particularly for higher ξ.

While robust cumulative regret notions are possible (Kirschner et al., 2020), we focus on the (simple) ξ-regret, as it was the focus of (Bogunovic et al., 2018a) and extensive related works (Sessa et al., 2020; Nguyen et al., 2020; Bertsimas et al., 2010).

Existing lower bound. A lower bound is proved in (Bogunovic et al., 2018a) for the case of constant $\xi>0$, with the same scaling as the standard setting. However, (Bogunovic et al., 2018a) only proves this hardness result for succeeding with probability very close to one; our lower bound in Section A.3 overcomes this limitation.

[1]Analogous results are also given for the standard simple regret (time to ϵ-optimality).

[2]Here we have presented simplified and slightly loosened forms; the refined variants are stated at the end of Appendix A.4.

3. Main Results

In this section, we formally state the new lower bounds that are summarized in Table 1.

3.1. Standard Setting

Our first contribution is to provide a new approach to establishing lower bounds in the standard setting (Section 2.1), with several advantages compared to (Scarlett et al., 2017) discussed in Section 4.4.

In the standard multi-armed bandit problem with a finite number of independent arms, (Kaufmann et al., 2016, Lemma 1) gives a versatile tool for deriving regret bounds based on the data processing inequality for KL divergence (e.g., see (Polyanskiy & Wu, 2014, Sec. 6.2)). The idea is that if two bandit instances must produce different outcomes (e.g., a different final point $\mathbf{x}^{(T)}$ must be returned) in order to succeed, but their sample distributions are close in KL divergence, then the time horizon must be large.

While (Kaufmann et al., 2016, Lemma 1) is only stated for a finite number of arms, the proof technique therein readily yields the variant in Lemma 1 below for a continuous input space, with the KL divergence quantities defined by maximizing within each of a finite number of regions partitioning the space. See also (Aziz et al., 2018) for an extension of (Kaufmann et al., 2016, Lemma 1) to a different infinite-arm problem.

In the following, we let $\mathbb{P}_f[\cdot]$ denote probabilities (with respect to the random noise) when the underlying function is f, and we let $P_f(y|\mathbf{x})$ be the conditional distribution $N(f(\mathbf{x}), \sigma^2)$ according to the Gaussian noise model.

Lemma 1. (Relating Two Instances – Adapted from (Kaufmann et al., 2016, Lemma 1)) *Fix $f, f' \in \mathcal{F}_k(B)$, let $\{\mathcal{R}_j\}_{j=1}^M$ be a partition of the input space into M disjoint regions, and let $\mathcal{A}$ be any event depending on the history up to some almost-surely finite stopping time τ.*[3] *Then, for $\delta \in \big(0, \frac{1}{3}\big)$, if $\mathbb{P}_f[\mathcal{A}] \ge 1-\delta$ and $\mathbb{P}_{f'}[\mathcal{A}] \le \delta$, we have*

$$\sum_{j=1}^{M} \mathbb{E}_f[N_j(\tau)]\overline{D}^j_{f,f'} \ge \log\frac{1}{2.4\delta}, \tag{8}$$

where $N_j(\tau)$ is the number of selected points in the j-th region up to time τ, and

$$\overline{D}^j_{f,f'} = \max_{\mathbf{x}\in\mathcal{R}_j} D\big(P_f(\cdot|\mathbf{x}) \,\|\, P_{f'}(\cdot|\mathbf{x})\big) \tag{9}$$

is the maximum KL divergence between samples (i.e., noisy function values) from f and f' in the j-th region.

[3]Following (Kaufmann et al., 2016), we state this result for general algorithms that are allowed to choose when to stop. Our focus in this paper is on the fixed-length setting in which the time horizon is pre-specified, and this setting is recovered by simply setting $\tau = T$ deterministically.

In Section 4, we will use Lemma 1 to prove the following lower bounds on the simple regret and cumulative regret, which are similar to those of (Scarlett et al., 2017) but enjoy an improved $\log\frac{1}{\delta}$ dependence on the target error probability δ. Despite this improvement, we highlight that the key contribution in this part of the paper is the novel lower bounding techniques for GP bandits via Lemma 1, rather than the results themselves. See Section 4.4 for a comparison to the approach of (Scarlett et al., 2017).

Theorem 1. (Simple Regret Lower Bound – Standard Setting) *Fix $\delta \in \big(0, \frac{1}{3}\big)$, $\epsilon \in \big(0, \frac{1}{2}\big)$, $B > 0$, and $T \in \mathbb{Z}$. Suppose there exists an algorithm that, for any $f \in \mathcal{F}_k(B)$, achieves average simple regret $r(\mathbf{x}^{(T)}) \le \epsilon$ with probability at least $1-\delta$. Then, if $\frac{\epsilon}{B}$ is sufficiently small, we have the following:*

1. *For $k = k_{\text{SE}}$, it is necessary that*

$$T = \Omega\bigg(\frac{\sigma^2}{\epsilon^2}\Big(\log\frac{B}{\epsilon}\Big)^{d/2}\log\frac{1}{\delta}\bigg). \tag{10}$$

2. *For $k = k_{\text{Matérn}}$, it is necessary that*

$$T = \Omega\bigg(\frac{\sigma^2}{\epsilon^2}\Big(\frac{B}{\epsilon}\Big)^{d/\nu}\log\frac{1}{\delta}\bigg). \tag{11}$$

Here, the implied constants may depend on (d, l, ν).

Theorem 2. (Cumulative Regret Lower Bound – Standard Setting) *Given $T \in \mathbb{Z}$, $\delta \in \big(0, \frac{1}{3}\big)$, and $B > 0$, for any algorithm, we must have the following:*

1. *For $k = k_{SE}$, there exists $f \in \mathcal{F}_k(B)$ such that the following holds with probability at least δ:*[4]

$$R_T = \Omega\left(\sqrt{T\sigma^2\Big(\log\frac{B^2T}{\sigma^2\log\frac{1}{\delta}}\Big)^{d/2}\log\frac{1}{\delta}}\right) \tag{12}$$

provided that[5] *$\frac{\sigma^2\log\frac{1}{\delta}}{B^2} = O(T)$ with a sufficiently small implied constant.*

2. *For $k = k_{Matérn}$, there exists $f \in \mathcal{F}_k(B)$ such that the following holds with probability at least δ:*

$$R_T = \Omega\bigg(B^{\frac{d}{2\nu+d}}T^{\frac{\nu+d}{2\nu+d}}\Big(\sigma^2\log\frac{1}{\delta}\Big)^{\frac{\nu}{2\nu+d}}\bigg) \tag{13}$$

provided that $\frac{\sigma^2\log\frac{1}{\delta}}{B^2} = O\big(T^{\frac{1}{2+d/\nu}}\big)$ with a sufficiently small implied constant.

[4]This "failure" event occurring with probability δ implies that the algorithm is unable to attain a $(1-\delta)$-probability of "success".

[5]As discussed in (Scarlett et al., 2017), scaling assumptions of this kind are very mild, and are needed to avoid the right-hand side of (12) contradicting a trivial $O(BT)$ upper bound.

Here, the implied constants may depend on (d, l, ν).

In Section 4.5, we show that for the Matérn kernel, the analysis can be simplified even further by using a function class proposed in (Bull, 2011) (which studied the noiseless setting) with bounded support.

3.2. Robust Setting – Corrupted Samples

In the general setup studied in (Bogunovic et al., 2020) and presented in Section 2.2, the player may randomize the choice of action, and the adversary can know the distribution but not the specific action. However, if the player's actions are deterministic (given the history), then knowing the distribution is equivalent to knowing the specific action. In this section, we provide a lower bound for such scenarios. While a lower bound that only holds for deterministic algorithms may seem limited, it is worth noting that the smallest regret upper bound in (Bogunovic et al., 2020) (see Theorem 9 in Appendix A) is established using such an algorithm. More generally, it is important to know to what extent randomization is needed for robustness, so bounds for both deterministic and randomized algorithms are of significant interest.

Theorem 3. (Lower Bound – Corrupted Samples) *In the setting of corrupted samples with a corruption level satisfying* $\Theta(1) \le C \le T^{1-\Omega(1)}$, *even in the noiseless setting* ($\sigma^2 = 0$)*, any deterministic algorithm (including those having knowledge of* C*) yields the following with probability one for some* $f \in \mathcal{F}_k(B)$:

- *Under the SE kernel,* $R_T = \Omega\big(C(\log T)^{d/2}\big)$;
- *Under the Matérn-*ν *kernel,* $R_T = \Omega\big(C^{\frac{\nu}{d+\nu}} T^{\frac{d}{d+\nu}}\big)$.

We provide a proof outline in Section 4.6, and the full details in Appendix B. We note that the assumption $\Theta(1) \le C \le T^{1-\Omega(1)}$ primarily rules out the case $C = \Theta(T)$ in which the adversary can corrupt every point by a constant amount. This assumption also ensures that the bound $R_T = \Omega\big(C^{\frac{\nu}{d+\nu}} T^{\frac{d}{d+\nu}}\big)$ is stronger than the bound $R_T = \Omega(C)$ from (Bogunovic et al., 2020). Note also that any lower bound for the standard setting applies here, since the adversary can choose not to corrupt.

Theorem 3 addresses a question posed in (Bogunovic et al., 2020) on the joint dependence of the cumulative regret on C and T. The upper bounds established therein (one of which we replicate in Theorem 9 in Appendix A) are of the form $O(C\overline{R}^{(0)})$, where $\overline{R}^{(0)}$ is a standard (non-corrupted) regret bound, whereas analogous results from the multi-armed bandit literature (Gupta et al., 2019) suggest that $\tilde{O}(\overline{R}^{(0)} + C)$ may be possible, where the $\tilde{O}(\cdot)$ notation hides dimension-independent logarithmic factors.

Theorem 3 shows that, at least for deterministic algorithms, such a level of improvement is impossible in the RKHS setting. On the other hand, further gaps remain between the lower bounds in Theorem 3 and the $O(C\overline{R}^{(0)})$ upper bounds of (Bogunovic et al., 2020) (e.g., for the SE kernel, the latter introduces an $\tilde{O}(C\sqrt{T(\log T)^{2d}})$ term, whereas Theorem 3 gives an $\Omega\big(C(\log T)^{d/2}\big)$ lower bound). Recent results for the linear bandit setting (Bogunovic et al., 2021) suggest that the looseness here may be in the upper bound; this is left for future work.

3.3. Robust Setting – Corrupted Final Point

Here we provide improved variant of the lower bound in (Bogunovic et al., 2018a) (replicated in Theorem 12 in Appendix A) for the adversarially robust setting with a corrupted final point, described in Section 2.3.

Theorem 4. (Improved Lower Bound – Corrupted Final Point) *Fix* $\xi \in \big(0, \frac{1}{2}\big)$, $\epsilon \in \big(0, \frac{1}{2}\big)$, $B > 0$, *and* $T \in \mathbb{Z}$, *and set* $\mathrm{dist}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2$. *Suppose that there exists an algorithm that, for any* $f \in \mathcal{F}_k(B)$, *reports* $\mathbf{x}^{(T)}$ *achieving* ξ*-regret* $r_\xi(\mathbf{x}^{(T)}) \le \epsilon$ *with probability at least* $1 - \delta$. *Then, provided that* $\frac{\epsilon}{B}$ *is sufficiently small, we have the following:*

1. *For* $k = k_{\mathrm{SE}}$, *it is necessary that* $T = \Omega\big(\frac{\sigma^2}{\epsilon^2}\big(\log\frac{B}{\epsilon}\big)^{d/2}\log\frac{1}{\delta}\big)$.
2. *For* $k = k_{\text{Matérn}}$, *it is necessary that* $T = \Omega\big(\frac{\sigma^2}{\epsilon^2}\big(\frac{B}{\epsilon}\big)^{d/\nu}\log\frac{1}{\delta}\big)$.

Here, the implied constants may depend on (ξ, d, l, ν).

Compared to the existing lower bound in (Bogunovic et al., 2018a) (Theorem 12 in Appendix A), we have removed the restrictive requirement that δ is sufficiently small (see Appendix A.3 for further discussion), giving the same scaling laws even when the algorithm is only required to succeed with a small probability such as 0.01 (i.e, $\delta = 0.99$). In addition, for small δ, we attain a $\log\frac{1}{\delta}$ factor improvement similar to Theorem 1.

4. Mathematical Analysis and Proofs

4.1. Preliminaries

In this section, we introduce some preliminary auxiliary results from (Scarlett et al., 2017) that will be used throughout our analysis. While we utilize the function class and auxiliary results from this existing work, we apply them in a significantly different manner in order to broaden the limited techniques known for GP bandit lower bounds, and to reap the advantages outlined above and in Section 4.4.

We proceed as follows (Scarlett et al., 2017):

- We lower bound the worst-case regret within $\mathcal{F}_k(B)$ by the regret averaged over a finite collection $\{f_1, \dots, f_M\} \subset \mathcal{F}_k(B)$ of size M.

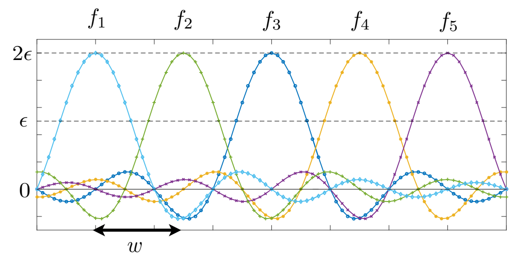

Figure 1. Illustration of functions $f_1, \dots, f_5$ such that any given point is ϵ-optimal for at most one function.

- Except where stated otherwise, we choose each $f_m(\mathbf{x})$ to be a shifted version of a common function $g(\mathbf{x})$ on $\mathbb{R}^d$. Specifically, each $f_m(\mathbf{x})$ is obtained by shifting $g(\mathbf{x})$ by a different amount, and then cropping to $D = [0,1]^d$. For our purposes, we require $g(\mathbf{x})$ to satisfy the following properties:

  1. The RKHS norm in $\mathbb{R}^d$ satisfies $\|g\|_k \le B$;
  2. We have (i) $g(\mathbf{x}) \in [-2\epsilon, 2\epsilon]$ with maximum value $g(0) = 2\epsilon$, and (ii) there is a "width" w such that $g(\mathbf{x}) < \epsilon$ for all $\|\mathbf{x}\|_\infty \ge \frac{w}{2}$;
  3. There are absolute constants $h_0 > 0$ and $\zeta > 0$ such that $g(\mathbf{x}) = \frac{2\epsilon}{h_0} h\big(\frac{\mathbf{x}\zeta}{w}\big)$ for some function $h(\mathbf{z})$ that decays faster than any finite power of $\|\mathbf{z}\|_2^{-1}$ as $\|\mathbf{z}\|_2 \to \infty$.

  Letting $g(\mathbf{x})$ be such a function, we construct the M functions by shifting $g(\mathbf{x})$ so that each $f_m(\mathbf{x})$ is centered on a unique point in a uniform grid, with points separated by w in each dimension. Since $D = [0,1]^d$, one can construct

$$M = \left\lfloor \left(\frac{1}{w}\right)^d \right\rfloor \tag{14}$$

  such functions; we will always consider $w \ll 1$, so that the case $M = 0$ is avoided. See Figure 1 for an illustration of the function class.

- It is shown in (Scarlett et al., 2017) that the above properties can be achieved with

$$M = \left\lfloor \left(\frac{\sqrt{\log \frac{B(2\pi l^2)^{d/4} h(0)}{2\epsilon}}}{\zeta \pi l} \right)^d \right\rfloor \tag{15}$$

  in the case of the SE kernel, and with

$$M = \left\lfloor \left(\frac{Bc_3}{\epsilon}\right)^{d/\nu} \right\rfloor \tag{16}$$

  in the case of the Matérn kernel, where $c_3 := \big(\frac{1}{\zeta}\big)^\nu \cdot \big(\frac{c_2^{-1/2}}{2(8\pi^2)^{(\nu+d/2)/2}}\big)$, and where $c_2 > 0$ is an absolute constant. Note that these values of M amount to choosing w in (14), and we will always consider $\frac{\epsilon}{B}$ to be sufficiently small, thus ensuring that $M \gg 1$ and $w \ll 1$ as stated above.

In addition, we introduce the following notation:

- The probability density function of the output sequence $\mathbf{y} = (y_1, \dots, y_T)$ when $f = f_m$ is denoted by $P_m(\mathbf{y})$ (and implicitly depends on the arbitrary underlying bandit algorithm). We also define $f_0(\mathbf{x}) = 0$ to be the zero function, and define $P_0(\mathbf{y})$ analogously for the case that the optimization algorithm is run on f_0. Expectations and probabilities (with respect to the noisy observations) are similarly written as $\mathbb{E}_m$, $\mathbb{P}_m$, $\mathbb{E}_0$, and $\mathbb{P}_0$ when the underlying function is f_m or f_0. On the other hand, in the absence of a subscript, $\mathbb{E}$ and $\mathbb{P}$ are taken with respect to the noisy observations *and* the random function f drawn uniformly from $\{f_1, \dots, f_M\}$. In addition, $\mathbb{P}_f$ and $\mathbb{E}_f$ will sometimes be used for generic f.

- Let $\{\mathcal{R}_m\}_{m=1}^M$ be a partition of the domain into M regions according to the above-mentioned uniform grid, with f_m taking its minimum value of -2ϵ in the center of $\mathcal{R}_m$. Moreover, let j_t be the index at time t such that $\mathbf{x}_t$ falls into $\mathcal{R}_{j_t}$; this can be thought of as a quantization of $\mathbf{x}_t$.

- Define the maximum absolute function value within a given region $\mathcal{R}_j$ as

$$\overline{v}_m^j := \max_{\mathbf{x} \in \mathcal{R}_j} |f_m(\mathbf{x})|, \tag{17}$$

  and the maximum KL divergence to P_0 within $\mathcal{R}_j$ as

$$\overline{D}_m^j := \max_{\mathbf{x} \in \mathcal{R}_j} D(P_0(\cdot|\mathbf{x}) \| P_m(\cdot|\mathbf{x})), \tag{18}$$

  where $P_m(y|\mathbf{x})$ is the distribution of an observation y for a given selected point $\mathbf{x}$ under the function f_m, and similarly for $P_0(y|\mathbf{x})$.

- Let $N_j \in \{0, \dots, T\}$ be a random variable representing the number of points from $\mathcal{R}_j$ that are selected throughout the T rounds.

Finally, the following auxiliary lemmas will be useful.

Lemma 2. (Scarlett et al., 2017, Eq. (36)) *For P_1 and P_2 being Gaussian with means (μ_1, μ_2) and a common variance σ^2, we have $D(P_1\|P_2) = \frac{(\mu_1-\mu_2)^2}{2\sigma^2}$.*

Lemma 3. (Scarlett et al., 2017, Lemma 7) *The functions $\{f_m\}$ corresponding to (15)–(16) are such that the quantities $\overline{v}_m^j$ in (17) satisfy (i) $\sum_{j=1}^M \overline{v}_m^j = O(\epsilon)$ for all m; (ii) $\sum_{m=1}^M \overline{v}_m^j = O(\epsilon)$ for all j; and (iii) $\sum_{m=1}^M (\overline{v}_m^j)^2 = O(\epsilon^2)$ for all j.*

4.2. Standard Setting – Simple Regret

For fixed $\epsilon > 0$, we consider the function class $\{f_1, \dots, f_M\}$ described in Section 4.1, with B replaced by $\frac{B}{3}$. This change should be understood as applying to all previous equations and auxiliary results that we use, e.g., replacing B by $\frac{B}{3}$ in (16), but this only affects constant factors, which we do not attempt to optimize anyway. Before continuing, we recall the important property that any $\mathbf{x} \in [0,1]^d$ can be ϵ-optimal for at most one function.

For fixed m and m', we will apply Lemma 1 with $f(\mathbf{x}) = f_m(\mathbf{x})$ and $f'(\mathbf{x}) = f_m(\mathbf{x}) + 2f_{m'}(\mathbf{x})$. Intuitively, this choice is made so that f and f' have different maximizers, but remain near-identical except in a small region around the peak of f'. It will be useful to characterize the quantity $\overline{D}^j_{f,f'}$ in (9); by Lemma 2,

$$\overline{D}^j_{f,f'} = \max_{\mathbf{x} \in \mathcal{R}_j} \frac{|2f_{m'}(\mathbf{x})|^2}{2\sigma^2} = \frac{2(\overline{v}^j_{m'})^2}{\sigma^2}, \tag{19}$$

using the definition of $\overline{v}^j_m$ in (17).

In the following, let $\mathcal{A}$ be the event that the returned point $\mathbf{x}^{(T)}$ lies in the region $\mathcal{R}_m$ (defined just above (17)). Suppose that an algorithm attains simple regret at most ϵ for both f and f' (note that $\|f'\|_k \le B$ by the triangle inequality and $\max_j \|f_j\|_k \le \frac{B}{3}$), each with probability at least $1-\delta$. We claim that this implies $\mathbb{P}_f[\mathcal{A}] \ge 1-\delta$ and $\mathbb{P}_{f'}[\mathcal{A}] \le \delta$. Indeed, by construction in Section 4.1, only points in $\mathcal{R}_m$ can be ϵ-optimal under $f = f_m$, and only points in $\mathcal{R}_{m'}$ can be ϵ-optimal under $f' = f_m + 2f_{m'}$. Hence, Lemma 1 and (19) give

$$\frac{2}{\sigma^2} \sum_{j=1}^M \mathbb{E}_m[N_j] \cdot (\overline{v}^j_{m'})^2 \ge \log \frac{1}{2.4\delta}, \tag{20}$$

and summing over all $m' \ne m$ gives

$$\frac{2}{\sigma^2} \sum_{m' \ne m} \sum_{j=1}^M \mathbb{E}_m[N_j] \cdot (\overline{v}^j_{m'})^2 \ge (M-1) \log \frac{1}{2.4\delta}. \tag{21}$$

Swapping the summations, using Lemma 3 to upper bound $\sum_{m' \ne m} (\overline{v}^j_{m'})^2 \le O(\epsilon^2)$, and applying $\sum_{j=1}^M \mathbb{E}_m[N_j(\tau)] = T$, we obtain

$$\frac{2c_0\epsilon^2 T}{\sigma^2} \ge (M-1) \log \frac{1}{2.4\delta} \tag{22}$$

for some constant c_0, or equivalently,

$$T \ge \frac{(M-1)\sigma^2}{2c_0\epsilon^2} \log \frac{1}{2.4\delta}. \tag{23}$$

Theorem 1 now follows using (15) (with $\frac{B}{3}$ in place of B) for the SE kernel, or (16) for the Matérn-ν kernel.

Remark 1. The preceding analysis can easily be adapted to show that when T is allowed to have variable length (i.e., the algorithm is allowed to choose when to stop), $\mathbb{E}[T]$ is lower bounded by the right-hand side of (23). See (Gabillon et al., 2012) for a discussion on analogous variations in the context of multi-armed bandits.

4.3. Standard Setting – Cumulative Regret

We fix some $\epsilon > 0$ to be specified later, consider the function class $\{f_1, \dots, f_M\}$ from Section 4.1, and show that it is not possible to attain $R_T \le \frac{T\epsilon}{2}$ with probability at least $1-\delta$ for all functions with $\|f\|_k \le B$.

Assuming by contradiction that the preceding goal is possible, this class of functions includes the choices of f and f' at the start of Section 4.2. However, if we let $\mathcal{A}$ be the event that at least $\frac{T}{2}$ of the sampled points lie in $\mathcal{R}_m$, it follows that $\mathbb{P}_f[\mathcal{A}] \ge 1-\delta$ and $\mathbb{P}_{f'}[\mathcal{A}] \le \delta$, since (i) each sample outside $\mathcal{R}_m$ incurs regret at least ϵ under f; and (ii) each sample within $\mathcal{R}_m$ incurs regret at least ϵ under f'.

Hence, despite being derived with a different choice of $\mathcal{A}$, (20) still holds in this case, and (23) follows. This was derived under the assumption that $R_T \le \frac{T\epsilon}{2}$ with probability at least $1-\delta$ for all functions with $\|f\|_k \le B$; the contrapositive statement is that when

$$T < \frac{(M-1)\sigma^2}{2c_0\epsilon^2} \log \frac{1}{2.4\delta}, \tag{24}$$

it must be the case that some function yields $R_T > \frac{T\epsilon}{2}$ with probability at least δ.

The remainder of the proof of Theorem 2 follows that of (Scarlett et al., 2017, Sec. 5.3), but with $\sigma^2 \log \frac{1}{\delta}$ in place of σ^2, and the final regret expressions adjusted accordingly (e.g., compare Theorem 2 with Theorem 8 in Appendix A). Due to this similarity, we only outline the details:

(i) Consider (24) nearly holding with equality (e.g., $T = \frac{M\sigma^2}{4c_0\epsilon^2} \log \frac{1}{2.4\delta}$ suffices).

(ii) Substitute M from (15) or (16) (with $\frac{B}{3}$ in place of B) into this choice from T, and solve to get an asymptotic expression for ϵ in terms of T.

(iii) Substitute this expression for ϵ into $R_T > \frac{T\epsilon}{2}$ to obtain the final regret bound.

See also Appendix B for similar steps given in more detail, albeit in the robust setting.

4.4. Comparison of Proof Techniques

The above analysis borrows ingredients from (Scarlett et al., 2017) and establishes similar final results; the key difference is in the use of Lemma 1 in place of an additive change-of-measure result (see Lemma 8 in Appendix D). We highlight the following advantages of our approach:

- While both approaches can be used to lower bound the average or constant-probability regret,[6] the above analysis gives the more precise $\log \frac{1}{\delta}$ dependence when the algorithm is required to succeed with probability at least $1-\delta$.
- As highlighted in Remark 1, the above simple regret analysis extends immediately to provide a lower bound on $\mathbb{E}[T]$ in the varying-T setting, whereas attaining this via the approach of (Scarlett et al., 2017) appears to be less straightforward.
- Although we do not explore it in this paper, we expect our approach to be more amenable to deriving *instance-dependent* regret bounds, rather than worst-case regret bounds over the function class. The idea, as in the multi-armed bandit setting (Kaufmann et al., 2016), is that if we can "perturb" one function/instance to another so that the set of near-optimal points changes significantly, we can use Lemma 1 to infer bounds on the required number of time steps in the original instance.
- As evidence of the versatility of our approach, we will use it in Section 4.7 to derive an improved result over that of (Bogunovic et al., 2018a) in the robust setting with a corrupted final point.

4.5. Simplified Analysis – Matérn Kernel

In the function class from (Scarlett et al., 2017) used above, each function is a bump function (with bounded support) in the frequency domain, meaning that it is non-zero almost everywhere in the spatial domain. In contrast, the earlier work of Bull (Bull, 2011) for the noiseless setting directly adopts a bump function in the spatial domain, permitting a simple analysis for the Matérn kernel.

It was noted in (Scarlett et al., 2017) that such a choice is infeasible for the SE kernel, since its RKHS norm is infinite. Nevertheless, in this section, we show that the (spatial) bump function is indeed much simpler to work with under the Matérn kernel, not only in the noiseless setting of (Bull, 2011), but also in the presence of noise.

The following result is stated in (Bull, 2011, Sec. A.2), and follows using Lemma 5 therein.

Lemma 4. (Bounded-Support Function Construction (Bull, 2011)) *Let* $h(\mathbf{x}) = \exp\big(\frac{-1}{1-\|\mathbf{x}\|^2}\big)\mathbb{1}\{\|\mathbf{x}\|_2 < 1\}$ *be the* d*-dimensional bump function, and define* $g(\mathbf{x}) = \frac{2\epsilon}{h(\mathbf{0})} h\big(\frac{\mathbf{x}}{w}\big)$ *for some* $w > 0$ *and* $\epsilon > 0$*. Then,* g *satisfies the following properties:*

- $g(\mathbf{x}) = 0$ *for all* $\mathbf{x}$ *outside the* ℓ_2*-ball of radius* w *centered at the origin;*
- $g(\mathbf{x}) \in [0, 2\epsilon]$ *for all* $\mathbf{x}$*, and* $g(\mathbf{0}) = 2\epsilon$.
- $\|g\|_k \le c_1 \frac{2\epsilon}{h(0)} \big(\frac{1}{w}\big)^\nu \|h\|_k$ *when* k *is the Matérn-*ν *kernel on* $\mathbb{R}^d$*, where* c_1 *is constant. In particular, we have* $\|g\|_k \le B$ *when* $w = \big(\frac{2\epsilon c_1 \|h\|_k}{h(\mathbf{0})B}\big)^{1/\nu}$.

This function can be used to simplify both the original analysis in (Scarlett et al., 2017), and the alternative proof in Sections 4.2–4.3. We focus on the latter, and on the simple regret; the cumulative regret can be handled similarly.

We consider functions $\{f_1, \dots, f_M\}$ constructed similarly to Section 4.1, but with each f_m being a shifted and scaled version of $g(\mathbf{x})$ in Lemma 4, using the choice of w in the third statement of the lemma. By the first part of the lemma, the functions have disjoint support as long as their center points are separated by at least w. We also use $\frac{B}{3}$ in place of B in the same way as Section 4.2. By forming a regularly spaced grid in each dimension, it follows that we can form

$$M = \Big\lfloor \frac{1}{w} \Big\rfloor^d = \Big\lfloor \frac{h(\mathbf{0})B}{6\epsilon c_1 \|h\|_k} \Big\rfloor^{d/\nu} \tag{25}$$

such functions.[7] Observe that this matches the $O\big(\big(\frac{B}{\epsilon}\big)^{d/\nu}\big)$ scaling in (16).

We clearly still have the property that any point $\mathbf{x}$ is ϵ-optimal for at most one function. The additional useful property here is that any point $\mathbf{x}$ yields *any non-zero value* for at most one function. Letting $\{\mathcal{R}_j\}_{j=1}^M$ be the partition of the domain induced by the above-mentioned grid (so that each f_m's support is a subset of $\mathcal{R}_m$), we notice that (20) still holds (with the choice of function in the definition of $v^j_{m'}$ suitably modified), but now simplifies to

$$\mathbb{E}_m[N_{m'}(\tau)] \cdot (v^{m'}_{m'})^2 \ge \frac{\sigma^2}{2} \log \frac{1}{2.4\delta}, \tag{26}$$

since there is no difference between $f(\mathbf{x}) = f_m(\mathbf{x})$ and $f'(\mathbf{x}) = f_m(\mathbf{x}) + 2f_{m'}(\mathbf{x})$ outside region $\mathcal{R}_{m'}$.

Since the maximum function value is 2ϵ, we have $(v^{m'}_{m'})^2 \le 4\epsilon^2$, and substituting into (26) and summing over $m' \neq m$ gives $T \ge \frac{\sigma^2(M-1)}{8\epsilon^2} \log \frac{1}{2.4\delta}$. This matches (23) up to modified constant factors, but is proved via a simpler analysis.

4.6. Robust Setting – Corrupted Samples

The high-level ideas behind proving Theorem 3 are outlined as follows, with the details in Appendix B:

[6] Under the new proof given here, this is achieved by setting $\delta = \frac{1}{2}$, or any other fixed constant in $(0, 1)$.

[7] We can seemingly fit significantly more points using a sphere packing argument (e.g., (Duchi, Sec. 13.2.3)), but this would only increase M by a constant factor depending on d, and we do not attempt to optimize constants in this paper.

- We consider an adversary that pushes all function values down to zero until its budget is exhausted.

- The function class chosen is similar to that in Figure 1, and we observe that (i) the adversary does not utilize much of its budget unless the sampled point is near the function's peak, and (ii) as long as the adversary is still active, the regret incurred at each time instant is typically $O(\epsilon)$ (since the algorithm has only observed $y_1 = \ldots = y_t = 0$, and thus has not learned where the peak is).

- We choose ϵ in a manner such that the adversary is still active at time T for at least half of the functions in the class, yielding $R_T = \Omega(T\epsilon)$. With M functions in the class, we show that occurs when $T\epsilon = \Theta(CM)$.

- Combining $T\epsilon = \Theta(CM)$ with the choices of M in (15) and (16) yields the desired result.

4.7. Robust Setting – Corrupted Final Point

To prove Theorem 4, we introduce a new function class that overcomes the limitation of that of (Bogunovic et al., 2018a) (illustrated in Figure 3 in Appendix A) in only handling success probabilities very close to one. Here we only present an idealized version of the function class that cannot be used directly due to yielding infinite RKHS norm. In Appendix C, we provide the proof details and the precise function class used. The idealized function class is depicted in Figure 2 for both $d = 1$ and $d = 2$.

We consider a class of functions of size $M + 1$, denoted by $\{f_0, f_1, \ldots, f_M\}$. For every function in the class, most points are within distance ξ of a point with value -2ϵ. However, there is a narrow region (depicted in plain color in Figure 2) where this may not be the case. The functions $f_1, \ldots, f_M$ are distinguished only by the existence of one additional narrow spike going down to -4ϵ in this region (see the 1D case in Figure 2), whereas for the function f_0, the spike is absent. For instance, in the 1D case, if the narrow spike has width w', then the number of functions is $M + 1 = \frac{\xi}{w'} + 1$.

With this class of functions, we have the following crucial observations on when the algorithm returns a point with ξ-stable regret at most ϵ:

- Under f_0, the returned point $\mathbf{x}^{(T)}$ must lie within the plain region of diameter ξ;

- Under any of $f_1, \ldots, f_M$, the returned point $\mathbf{x}^{(T)}$ must lie outside that plain region;

- The only way to distinguish between f_0 and a given f_i is to sample within the associated narrow spike in which the function value is -4ϵ.

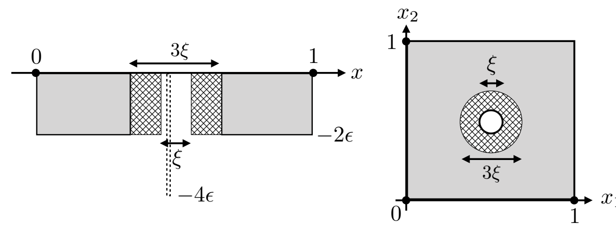

Figure 2. Idealized version of the function class under corrupted final points, in 1D (left) and 2D (right). The shaded regions have value -2ϵ; the checkered regions have value 0 but their points can be perturbed into the shaded region; and the plain regions have value 0 and cannot be. If another spike is present, as per the dashed curve in the 1D case, then this creates a region (covering the entire plain region) that can be perturbed down to -4ϵ.

Due to the $N(0, \sigma^2)$ noise, this roughly amounts to needing to take $\Omega\big(\frac{\sigma^2}{\epsilon^2}\big)$ samples within the narrow spike, and since there are M possible spike locations, this means that $\Omega\big(\frac{M\sigma^2}{\epsilon^2}\big)$ samples are needed. We therefore have a similar lower bound to (23), and a similar regret bound to the standard setting follows (with modified constants additionally depending on ξ).

5. Conclusion

We have provided novel techniques and results for algorithm-independent lower bounds in non-Bayesian GP bandit optimization. In the standard setting, we have provided a new proof technique whose benefits include simplicity, versatility, and improved dependence on the error probability.

In the robust setting with corrupted samples, we have provided the first lower bound characterizing *joint* dependence on the corruption level and time horizon. In the robust setting with a corrupted final point, we have overcome a limitation of the existing lower bound, demonstrating the impossibility of attaining any non-trivial constant error probability rather than only values close to one.

An immediate direction for future work is to further close the gaps in the upper and lower bounds, particularly in the robust setting with corrupted samples.

Acknowledgement

This work was supported by the Singapore National Research Foundation (NRF) under grant number R-252-000-A74-281.

References

Aronszajn, N. Theory of reproducing kernels. *Trans. Amer. Math. Soc*, 68(3):337–404, 1950.

Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *IEEE Conf. Found. Comp. Sci. (FOCS)*, 1995.

Aziz, M., Anderton, J., .Kaufmann, E., and .Aslam, J. Pure exploration in infinitely-armed bandit models with fixed-confidence. In *Alg. Learn. Theory (ALT)*, 2018.

Beland, J. J. and Nair, P. B. Bayesian optimization under uncertainty. NIPS BayesOpt 2017 workshop, 2017.

Bertsimas, D., Nohadani, O., and Teo, K. M. Nonconvex robust optimization for problems with constraints. *INFORMS journal on Computing*, 22(1):44–58, 2010.

Bogunovic, I., Scarlett, J., Krause, A., and Cevher, V. Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*, 2016.

Bogunovic, I., Scarlett, J., Jegelka, S., and Cevher, V. Adversarially robust optimization with Gaussian processes. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*, 2018a.

Bogunovic, I., Zhao, J., and Cevher, V. Robust maximization of non-submodular objectives. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2018b.

Bogunovic, I., Krause, A., and Scarlett, J. Corruption-tolerant Gaussian process bandit optimization. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2020.

Bogunovic, I., Losalka, A., Krause, A., and Scarlett, J. Stochastic linear bandits robust to adversarial attacks. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2021.

Bull, A. D. Convergence rates of efficient global optimization algorithms. *J. Mach. Learn. Res.*, 12(Oct.):2879–2904, 2011.

Calandriello, D., Carratino, L., Lazaric, A., Valko, M., and Rosasco, L. Gaussian process optimization with adaptive sketching: Scalable and no regret. In *Conf. Learn. Theory (COLT)*, 2019.

Chowdhury, S. R. and Gopalan, A. On kernelized multi-armed bandits. In *Int. Conf. Mach. Learn. (ICML)*, 2017.

Chowdhury, S. R. and Gopalan, A. Bayesian optimization under heavy-tailed payoffs. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*, 2019.

Contal, E., Buffoni, D., Robicquet, A., and Vayatis, N. *Machine Learning and Knowledge Discovery in Databases*, chapter Parallel Gaussian Process Optimization with Upper Confidence Bound and Pure Exploration, pp. 225–240. Springer Berlin Heidelberg, 2013.

Dai Nguyen, T., Gupta, S., Rana, S., and Venkatesh, S. Stable Bayesian optimization. In *Pacific-Asia Conf. Knowl. Disc. and Data Mining*, 2017.

De Freitas, N., Smola, A. J., and Zoghi, M. Exponential regret bounds for Gaussian process bandits with deterministic observations. In *Int. Conf. Mach. Learn. (ICML)*, 2012.

Duchi, J. Lecture notes for statistics 311/electrical engineering 377. http://stanford.edu/class/stats311/.

Gabillon, V., Ghavamzadeh, M., and Lazaric, A. Best arm identification: A unified approach to fixed budget and fixed confidence. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*, 2012.

Grill, J.-B., Valko, M., and Munos, R. Optimistic optimization of a Brownian. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*. 2018.

Grünewälder, S., Audibert, J.-Y., Opper, M., and Shawe-Taylor, J. Regret bounds for Gaussian process bandit problems. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, pp. 273–280, 2010.

Gupta, A., Koren, T., and Talwar, K. Better algorithms for stochastic bandits with adversarial corruptions. In *Conf. Learn. Theory (COLT)*, 2019.

Janz, D., Burt, D. R., and González, J. Bandit optimisation of functions in the Matérn kernel RKHS. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2020.

Kaufmann, E., Cappé, O., and Garivier, A. On the complexity of best-arm identification in multi-armed bandit models. *J. Mach. Learn. Res. (JMLR)*, 17(1):1–42, 2016.

Kawaguchi, K., Kaelbling, L. P., and Lozano-Pérez, T. Bayesian optimization with exponential convergence. In *Conf. Neur. Inf. Proc. Sys. (NeurIPS)*, 2015.

Kirschner, J., Bogunovic, I., Jegelka, S., and Krause, A. Distributionally robust Bayesian optimization. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2020.

Li, Y., Lou, E. Y., and Shan, L. Stochastic linear optimization with adversarial corruption. https://arxiv.org/abs/1909.02109, 2019.

Lykouris, T., Mirrokni, V., and Paes Leme, R. Stochastic bandits robust to adversarial corruptions. In *ACM Symp. Theory Comp. (STOC)*, 2018.

Lyu, Y., Yuan, Y., and Tsang, I. W. Efficient batch black-box optimization with deterministic regret bounds. https://arxiv.org/abs/1905.10041, 2019.

Martinez-Cantin, R., Tee, K., and McCourt, M. Practical Bayesian optimization in the presence of outliers. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2018.

Nguyen, T. T., Gupta, S., Ha, H., Rana, S., and Venkatesh, S. Distributionally robust bayesian quadrature optimization. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2020.

Nogueira, J., Martinez-Cantin, R., Bernardino, A., and Jamone, L. Unscented Bayesian optimization for safe robot grasping. In *IEEE/RSJ Int. Conf. Intel. Robots and Systems (IROS)*, 2016.

Polyanskiy, Y. and Wu, Y. Lecture notes on information theory. http://www.stat.yale.edu/~yw562/ln.html, 2014.

Rasmussen, C. E. Gaussian processes for machine learning. MIT Press, 2006.

Scarlett, J. Tight regret bounds for Bayesian optimization in one dimension. In *Int. Conf. Mach. Learn. (ICML)*, 2018.

Scarlett, J., Bogunovic, I., and Cevher, V. Lower bounds on regret for noisy Gaussian process bandit optimization. In *Conf. Learn. Theory (COLT)*. 2017.

Sessa, P. G., Bogunovic, I., Kamgarpour, M., and Krause, A. Mixed strategies for robust optimization of unknown objectives. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2020.

Shekhar, S. and Javidi, T. Gaussian process bandits with adaptive discretization. *Elec. J. Stats.*, 12(2):3829–3874, 2018.

Shekhar, S. and Javidi, T. Multi-scale zero-order optimization of smooth functions in an RKHS. https://arxiv.org/abs/2005.04832, 2020.

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Int. Conf. Mach. Learn. (ICML)*, 2010.

Vakili, S., Picheny, V., and Durrande, N. Regret bounds for noise-free Bayesian optimization. *https://arxiv.org/abs/2002.05096*, 2020.

Vakili, S., Khezeli, K., and Picheny, V. On information gain and regret bounds in Gaussian process bandits. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2021.

Valko, M., Korda, N., Munos, R., Flaounas, I., and Cristianini, N. Finite-time analysis of kernelised contextual bandits. In *Conf. Uncertainty in AI (UAI)*, 2013.

Wang, Z. and Jegelka, S. Max-value entropy search for efficient Bayesian optimization. In *Int. Conf. Mach. Learn. (ICML)*, pp. 3627–3635, 2017.

Wang, Z., Zhou, B., and Jegelka, S. Optimization as estimation with Gaussian processes in bandit settings. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2016.

Wang, Z., Tan, V. Y. F., and Scarlett, J. Tight regret bounds for noisy optimization of a Brownian motion. https://arxiv.org/abs/2001.09327, 2020.