A Robust Phased Elimination Algorithm for Corruption-Tolerant Gaussian Process Bandits

Ilija Bogunovic
University College London
i.bogunovic@ucl.ac.uk

Zihan Li
National University of Singapore
lizihan@u.nus.edu

Andreas Krause
ETH Zürich
krausea@ethz.ch

Jonathan Scarlett
National University of Singapore
scarlett@comp.nus.edu.sg

Abstract

We consider the sequential optimization of an unknown, continuous, and expensive to evaluate reward function, from noisy and adversarially corrupted observed rewards. When the corruption attacks are subject to a suitable budget C and the function lives in a Reproducing Kernel Hilbert Space (RKHS), the problem can be posed as *corrupted Gaussian process (GP) bandit optimization*. We propose a novel robust elimination-type algorithm that runs in epochs, combines exploration with infrequent switching to select a small subset of actions, and plays each action for multiple time instants. Our algorithm, *Robust GP Phased Elimination (RGP-PE)*, successfully balances robustness to corruptions with exploration and exploitation such that its performance degrades minimally in the presence (or absence) of adversarial corruptions. When T is the number of samples and γ_T is the maximal information gain, the corruption-dependent term in our regret bound is $O(C\gamma_T^{3/2})$, which is significantly tighter than the existing $O(C\sqrt{T\gamma_T})$ for several commonly-considered kernels. We perform the first empirical study of robustness in the corrupted GP bandit setting, and show that our algorithm is robust against a variety of adversarial attacks.

1 Introduction

Black-box optimization is a fundamental problem with broad applications including hyperparameter tuning [42], robotics [34], and chemical design [20], among others. To make the problem tractable, a variety of smoothness properties have been adopted, and Reproducing Kernel Hilbert Space (RKHS) functions have proved to provide a versatile framework that can be tackled via Gaussian process (GP) methods [43, 15]. This problem is referred to as *GP bandits* or *kernelized bandits*.

While an extensive line of works have established GP bandit algorithms and regret bounds, settings with adversarial corruptions have only arisen relatively recently. Such corruptions may come in the form of outliers [38, 41], perturbations of sampled inputs [5, 40, 16], adversarial noise in the rewards [8], or perturbations of the final recommendation [7]. In this work, we are interested in the setting of adversarial noise in the rewards, in which the performance of standard non-robust GP bandit algorithms can deteriorate significantly (see Fig. 1).

The first work considering this setting [8] established regret bounds for various algorithms depending on the degree of knowledge on the corruption level C (defined formally in Section 2). A key limitation in their regret bound is that the main corruption-dependent term, C, and the usual uncorrupted regret term, which is $\sqrt{T}$ or higher (with time horizon T), are *multiplied together*. That is, the dependence

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

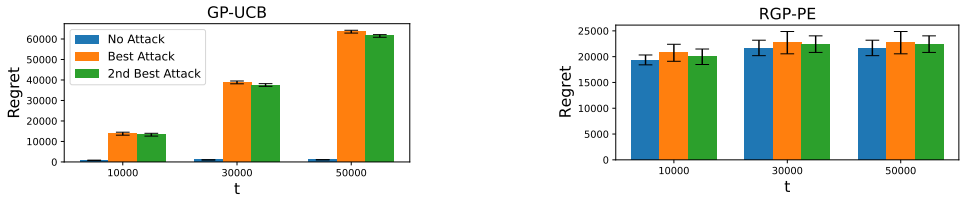

Figure 1: Performance of GP-UCB [43] and Robust GP Phased Elimination (RGP-PE, this work) with no attacks and the two most effective corruption attacks on the Robot3D pushing task. As the number of samples t increases, the performance of non-robust GP-UCB deteriorates significantly under both attacking strategies, while the performance of the proposed algorithm remains robust.

on C is multiplicative with respect to the uncorrupted bound. Analogous studies of bandits with independent arms [35, 21] or linear rewards [9] suggest that *additive* dependence may be possible, but this has remained very much open in the GP bandit setting.

In this paper, we address this fundamental gap in the literature by introducing a novel algorithm in which the uncorrupted term and the C-dependent term are clearly decoupled, and the latter is only multiplied by a kernel-dependent function of T that can be much smaller than $\sqrt{T}$.

Related work. The closest work to ours is [8], which also considers the corruption-tolerant GP bandit setting. In that work, the authors propose a confidence-bound-based algorithm with enlarged confidence. As outlined above, the regret bound therein scales as $O(C\sqrt{T\gamma_T})$, and the possibility of additive C dependence was left as an open problem.

The question of additive vs. multiplicative dependence first arose in multi-armed bandits with independent arms, with an initial work [35] being multiplicative, and a subsequent work [21] improving to additive. Closer to our setup (and in fact a special case of it via the linear kernel) is the case of corrupted stochastic linear bandits, in which additive dependence was obtained in [9], with the corruption term more precisely being $O(Cd^{3/2}\log T)$ under mild assumptions.[1] Our main result will achieve a similar bound as a special case, while being much more general due to handling general kernels, and adopting GP-based algorithmic and mathematical techniques that have minimal overlap with the linear setting. Other less related results for corrupted linear bandits (e.g., contextual or instance-dependent) are given in [31] and [50].

Adversarial corruptions of the rewards were also considered for GP bandits in [28], but with the key difference of considering a weaker adversary that does not know the chosen action when choosing the corruption term. This distinction has a considerable effect on the problem, leading to significantly different algorithms, and with the setting of [28] leading to a GP-UCB-style regret bound in which the corruptions only impact the constant factors. In our setting, the effect of corruptions is much more significant, and we know from [9] that this is unavoidable in general.

Other less related notions of robustness in GP bandits have included outliers [38], misspecification [13, 6], input noise [5, 40, 16], risk-aversion [39, 11, 37], and corruptions in the final recommendation [7, 29]. Moreover, other settings with adversarial corruptions have included multi-armed bandit and online [21, 25, 24, 3], active [14], reinforcement learning [36, 49, 4], and multi-agent RL [33].

Corruption-robustness have been considered in other sequential decision making problems including multi-armed bandits and predictioin with expert advice / online learning.

Contributions. We provide a novel algorithm for GP bandit optimization with adversarial corruptions, that attains the first regret bound to avoid multiplying the uncorrupted part by the corruption level C. Our algorithm crucially incorporates a *rare switching* idea, along with a non-standard robust estimator, enlarged confidence bounds, and a minimal number of plays of each selected action; see Sections 2.1 and 3 for details. To our knowledge, we are the first to use rare switching to achieve adversarial robustness; previous works instead used it for reducing computational complexity.

We show that our regret bound is *provably near-optimal* for the SE kernel, and recovers recently-established bounds for stochastic linear bandits [9] that are also known to be near-optimal. For the Matérn kernel, the degree of tightness depends on the dimension and smoothness parameter, but

[1] In a paper concurrent with ours, the $Cd^{3/2}$ dependence has been improved to Cd for linear bandits [23]. We leave it for future work to determine whether a similar improvement is possible for GP bandits.

our bound strictly improves on that of [8] in all scaling regimes where the latter is non-trivial (i.e., sub-linear in T); see Table 1 on Page 7 for a summary. We demonstrate that our algorithm is able to successfully defend against various attacks, including those proposed in [22].

On the technical side, we note that the GP setting dictates the use of a significantly different algorithm compared to linear bandits, and a technical analysis with only minor overlap. To highlight this, in Appendix E, we explore an approach based on a direct reduction to linear bandits (followed by using the algorithm in [9]), and show that it yields strictly worse regret scaling than our main result.

2 Problem Setting and Preliminaries

We consider the Gaussian process bandit (i.e., kernelized bandit) problem, in which the goal of the learner is to maximize the collected rewards by sequentially querying the unknown reward function $f : \mathcal{X} \to \mathbb{R}$ over T rounds. In particular, at every time t, the learner selects $x_t \in \mathcal{X}$ and receives

$$y_t = f(x_t) + \epsilon_t, \tag{1}$$

where ϵ_t is assumed to be σ-sub-Gaussian with independence over time steps, and σ is also known.

We consider the corrupted setting in which, besides the stochastic noise, the observations at every time step are adversarially corrupted, so that the learner observes

$$\widetilde{y}_t = y_t + c_t. \tag{2}$$

Following [8], we make the following assumptions on the adversary:

- The adversary knows the true reward function $f(\cdot)$, and, at every round t, it observes x_t before deciding upon the corruption c_t.
- The total adversarial corruption budget over T rounds is bounded as follows:

$$\sum_{t=1}^{T} |c_t| \leq C. \tag{3}$$

  In this paper, we focus primarily on the case where C is known to the learner, but we also discuss in Section 3.4 how our results have implications for the case of unknown C.

The domain $\mathcal{X}$ is assumed to either be finite, or a compact subset of $\mathbb{R}^d$ for some dimension d (e.g., $\mathcal{X} = [0, 1]^d$). In either case, $\mathcal{X}$ is endowed with a continuous, positive semidefinite kernel function $k(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ that is normalized to satisfy $k(x, x') \leq 1$ for all $x, x' \in \mathcal{X}$. We further assume that f has a bounded norm in the corresponding Reproducing Kernel Hilbert Space (RKHS) $\mathcal{H}_k$, i.e., $\|f\|_k \leq B$ (see Appendix A for more details). This assumption permits the construction of confidence bounds via Gaussian process (GP) models (Section 3.2).

The learner's performance is measured using the widely-considered notion of cumulative regret:

$$R_T = \sum_{t=1}^{T} \Big(\max_{x \in \mathcal{X}} f(x) - f(x_t) \Big), \tag{4}$$

and we are interested in the *joint* dependence of R_T on C and T. As noted in [35] and [8], one could alternatively define the cumulative regret with respect to the corrupted values (i.e., $f(x) + c_t$), and these notions coincide to within an additive term of $2C$.

2.1 Gaussian Process Model under Corruptions

In the standard (non-corrupted) setting, previous algorithms use (i) zero-mean GP priors for modeling the uncertainty in f (i.e., they assume $f \sim GP(0, k)$), and (ii) Gaussian likelihood models for the observations. As more data points become available, Bayesian posterior updates are then performed according to a misspecified model in which the noise variables $\epsilon_t = y_t - f(x_t)$ are assumed to be drawn independently across t from $\mathcal{N}(0, \lambda)$, where λ is a hyperparameter that may differ from the true noise variance σ^2. In particular, in the absence of corruptions, given a sequence of points $\{x_1, \ldots, x_t\}$ and their noisy observations $\{y_1, \ldots, y_t\}$, the posterior mean and variance are given by

$$\mu_t(x) = k_t(x)^T \big(K_t + \lambda I_t\big)^{-1} Y_t, \tag{5}$$

$$\sigma_t^2(x) = k(x, x) - k_t(x)^T \big(K_t + \lambda I_t\big)^{-1} k_t(x), \tag{6}$$

where $k_t(x) = \big[k(x_i, x)\big]_{i=1}^t$, $K_t = \big[k(x_t, x_{t'})\big]_{t,t'}$ is the kernel matrix, and $Y_t \in \mathbb{R}^t$ contains the non-corrupted observations up to time t, i.e., $Y_t[i] = y_i$ for $i \in [t]$.

In the corrupted setting, given the inputs $\{x_1, \dots, x_t\}$ and their corrupted observations $\{\widetilde{y}_1, \dots, \widetilde{y}_t\}$ (with $\widetilde{y}_i = y_i + c_i$), we propose the following non-standard robust posterior mean estimator:

$$\widetilde{\mu}_t(x) = k_t(x)^T (K_t + \lambda I_t)^{-1} \widetilde{Y}_t, \tag{7}$$

where $\widetilde{Y}_t \in \mathbb{R}^t$ and $\widetilde{Y}_t[i] = \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}\widetilde{y}_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}}$ for $i \in [t]$. Intuitively, the averaging of terms corresponding to identical actions is done in order to diminish the impact of corruption, and this will be a crucial component of our analysis. In our algorithm, besides $\widetilde{\mu}_t(\cdot)$, we will also make use of the standard posterior variance $\sigma_t^2(\cdot)$ as given in Eq. (6); the use of this quantity is intuitively reasonable because GP posterior variances do not depend on the observations.

The main quantity that characterizes the regret bounds in the non-corrupted setting (and is also useful in our setting) is the *maximum information gain* [43], defined at time t as

$$\gamma_t = \max_{x_1, \dots, x_t} \frac{1}{2} \ln \det(I_t + \lambda^{-1} K_t). \tag{8}$$

3 Robust GP Phased Elimination

3.1 Algorithm and Confidence Bounds

Our algorithm works in epochs indexed by $h = 0, 1, \dots, H-1$, each of which consists of sampling a batch of points. The epoch lengths may be chosen adaptively, and hence H may not be deterministic, but we will ensure with probability one that $H \le \bar{H}$ with $\bar{H} = \log_2 T$. The length of epoch h is denoted by u_h, so that $\sum_{h=0}^{H-1} u_h = T$.

The algorithm and analysis are based on the widespread notion of confidence bounds. While our confidence bounds will be expanded to account for corruptions, it is useful to consider the following generic assumption regarding non-corrupted observations (although the algorithm cannot access these, they will appear in our mathematical analysis).

Assumption 1 (Regular confidence bounds)**.** *Let $\mu^{(h)}(x)$ and $\sigma^{(h)}(x)$ denote the posterior mean and standard deviation computed (hypothetically) using only the non-corrupted observations $\{(x_i, y_i)\}_{i=1}^{u_h}$ in epoch h using Eqs.* (5) *and* (6)*. We assume that given $\delta \in (0, 1)$, there exists a sequence of parameters $\beta_h = \beta_h(\delta)$ which is non-decreasing in h and yields with probability at least $1 - \delta$ that*

$$|\mu^{(h)}(x) - f(x)| \le \beta_h \sigma^{(h)}(x) \tag{9}$$

simultaneously for all $h \ge 0$ and $x \in \mathcal{X}$.

Specific choices of β_h satisfying this assumption will be considered in Section 3.2.

Similarly to previous kernelized algorithms (e.g., [8, 6]), our proposed algorithm makes use of enlarged confidence bounds. Hence, our first result concerns concentration of an RKHS member under corrupted observations, where we make use of the proposed estimator from Eq. (7).

Lemma 2 (Corrupted confidence bounds)**.** *Under Assumption 1, let $\widetilde{\mu}^{(h)}(x)$ denote the posterior mean based on only the corrupted observations $\{(x_i, \widetilde{y}_i)\}_{i=1}^{u_h}$ in epoch h using Eq.* (7)*, and let $u_{\min} \ge 1$ denote the minimum number of times any single action from $\{x_i\}_{i=1}^{u_h}$ is played, i.e., $u_{\min} = \min_{x \in \{x_1, \dots, x_{u_h}\}} \sum_{i=1}^{u_h} \mathbb{1}\{x_i = x\}$. Then, with probability at least $1 - \delta$, it holds for all $x \in \mathcal{X}$ and $h \ge 0$ that*

$$|\widetilde{\mu}^{(h)}(x) - f(x)| \le \Big(\beta_h + \tfrac{C\sqrt{u_h}}{u_{\min}\lambda}\Big)\sigma^{(h)}(x). \tag{10}$$

The confidence-bound enlargement is proportional to the total amount of corruption C. This bears some similarity to the confidence intervals used in [8, Lemma 2], but we note the following important differences:

- We make use of a novel kernelized mean estimator (Eq. (7)) that takes average over rewards corresponding to the same played action;

Algorithm 1 Robust GP Phased Elimination (RGP-PE)

Input: Domain $\mathcal{X} \subset \mathbb{R}^d$, truncation parameter $\psi > 0$, corruption budget C, switching parameter $\eta > 1$, regularization parameter $\lambda > 0$

1: Initialize $l_0 = 2$, and $h = 0$ and $\mathcal{X}_h = \mathcal{X}$
2: Set $\mathcal{S}_h = \emptyset$, $t' = 0$, $\sigma_0(x) = 1$ for all $x \in \mathcal{X}_h$
3: **for** $t = 1, 2, \ldots, l_h$ **do**
4: Select $x_t = \arg\max_{x \in \mathcal{X}_h} \sigma_{t'}(x)$
5: Update $\mathcal{S}_h \leftarrow \mathcal{S}_h \cup \{x_t\}$
6: **if** $\det(I_t + \lambda^{-1} K_t) > \eta \det(I_{t'} + \lambda^{-1} K_{t'})$ **then**
7: Set $t' \leftarrow t$
8: Compute $\sigma_{t'}(\cdot)$ via Eq. (6) by using $\{x_i\}_{i=1}^{t'}$
9: **end if**
10: **end for**
11: Set $\xi_h(x) = \frac{\sum_{i=1}^{l_h} \mathbb{1}\{x = x_i\}}{l_h}$ for every $x \in \mathcal{S}_h$
12: Set $u_h(x) = \lceil l_h \max\{\xi_h(x), \psi\} \rceil$ for every $x \in \mathcal{S}_h$
13: Take each action $x \in \mathcal{S}_h$ exactly $u_h(x)$ times with corresponding rewards $(\widetilde{y}_j)_{j=1}^{u_h}$ where $u_h = \sum_{x \in \mathcal{S}_h} u_h(x)$
14: Estimate $\widetilde{\mu}^{(h)}(\cdot)$ and $\sigma^{(h)}(\cdot)$ according to Eq. (7) and Eq. (6) using only the u_h points from the current epoch.
15: Update the active set of actions to:

$$\mathcal{X}_{h+1} \leftarrow \Big\{ x \in \mathcal{X}_h : \widetilde{\mu}^{(h)}(x) + \big(\beta_h + \tfrac{C\sqrt{u_h}}{l_h \psi \lambda}\big)\sigma^{(h)}(x) \geq \max_{x \in \mathcal{X}_h} \widetilde{\mu}^{(h)}(x) - \big(\beta_h + \tfrac{C\sqrt{u_h}}{l_h \psi \lambda}\big)\sigma^{(h)}(x) \Big\}$$

16: Set $l_{h+1} \leftarrow 2l_h$, $h \leftarrow h + 1$ and return to Step 2 (terminating after T total actions are played).

- Our enlargement term is $O(C\frac{\sqrt{u_h}}{u_{\min}})$, as opposed to $O(C)$ used in [8, Lemma 2]. We will typically apply this lemma with $\frac{\sqrt{u_h}}{u_{\min}} \ll 1$, so that our confidence width is much smaller.

For the second of these, the intuition is that if the same action is played multiple times, it becomes harder for the adversary to hide the true value (i.e., since the rewards of the same played actions are averaged, the adversary needs to spend more of its budget corrupting the reward).

The Robust GP-Phased Elimination algorithm (Algorithm 1) proceeds in epochs (indexed by h) of exponentially increasing length u_h. At every round t (where $t \in \{1, \ldots, l_h\}$ and $l_h = 2^{h+1}$) within an epoch h, the algorithm selects an action maximizing a posterior uncertainty computed at some (possibly strictly earlier) time t':

$$x_t = \arg\max_{x \in \mathcal{X}_h} \sigma_{t'}(x), \tag{11}$$

where $\mathcal{X}_h$ denotes the set of active actions in epoch h. The selected action is then added to $\mathcal{S}_h$ which is a set that contains distinct actions selected in epoch h.

The key idea behind using t' instead of t in Eq. (11) is to ensure that our algorithm *rarely switches*, based on a condition relating to the information gain (Line 6), meaning that the same action x_t is typically selected multiple times. Whenever there are ties, they are resolved arbitrarily but consistently over rounds (i.e., if $\sigma_{t'}(\cdot)$ does not change, the same points are selected). Based on Lines 6 to 9, we update t' and recompute $\sigma_{t'}(x)$ only when $\det(I_t + \lambda^{-1} K_t)$ increases by a constant factor η.

Related ideas of rare switching have appeared in the literature (e.g., [1, 47, 19]), but to our knowledge we are the first to use this idea in the kernelized bandit problem to provide an algorithm that includes an explicit switching condition for improving robustness. Intuitively, by rarely switching, we obtain more samples of the same point, allowing us to average more of them together and making the "averaged" observation harder to corrupt. Concurrent work also used rare switching to reduce GP posterior computation, noting that the computation time can be made to scale (cubically) with the number of *unique* points [12]. This benefit also applies directly to our algorithm, and we exploit it to run large-T experiments in Section 4.

After the set $\mathcal{S}_h$ is constructed, we define $\xi_h(x) = \frac{\sum_{i=1}^{l_h} \mathbb{1}\{x=x_i\}}{l_h}$ for every $x \in \mathcal{S}_h$, representing the empirical frequency of selecting $x_t \in \mathcal{X}_h$ in l_h rounds. The algorithm then plays actions from $\mathcal{S}_h$ only, where the number of times each action x from $\mathcal{S}_h$ is played is denoted by $u_h(x) = \lceil l_h \max\{\xi_h(x), \psi\} \rceil$. Here, the *truncation parameter* ψ ensures that each action from $\mathcal{S}_h$ is played sufficiently many times; this idea was used for corrupted linear bandits in [9]. Our theory suggests a particular choice of ψ; see Theorem 3. Each action $x \in \mathcal{S}_h$ is played for $u_h(x)$ times in an arbitrary order, leading to the total epoch length $u_h = \sum_{x \in \mathcal{S}_h} u_h(x)$.

Based on the received noisy and potentially corrupted rewards $\{x_j, \widetilde{y}_j\}_{j=1}^{u_h}$, the algorithm updates its estimates $\widetilde{\mu}^{(h)}(\cdot)$ and $\sigma^{(h)}(\cdot)$ according to Eq. (7) and Eq. (6). Finally, each epoch h ends by updating the set of active actions $\mathcal{X}_{h+1}$. To do so, we use the confidence bounds from Lemma 2 with $u_{\min} = l_h\psi$, where $l_h\psi$ is a lower bound on the number of times each distinct action from $\mathcal{S}_h$ is played. These confidence bounds are valid in the sense that the true function is contained within the confidence bounds with high probability. The definition of $\mathcal{X}_{h+1}$ (Line 15) ensures that with high probability, the optimal action is never eliminated.

Besides the standard exploration/exploitation trade-off (controlled via β_h), our algorithm additionally balances robustness to corruptions. This is done via two parameters: the switching parameter η and truncation parameter ψ. We set these parameter to ensure that the number of distinct actions played per epoch is sufficiently small, while the number of plays per each such action is sufficiently large. This trade-off is non-trivial; for example, in the case that $C = 0$ (i.e., the non-corrupted setting), resampling the same actions (controlled via ψ) increases the regret.

Main result. We now present our main theoretical result, where we use $O^*(\cdot)$ notation to hide constants and dimension-independent log factors. We treat the RKHS norm bound B as being fixed, so its dependence is also hidden in $O(\cdot)$ or $O^*(\cdot)$ notation.

Theorem 3 (Main result)**.** *Under the preceding setup and Assumption 1, for any corruption budget $C \geq 0$, Algorithm 1 with a constant switching parameter $\eta > 1$ and truncation parameter $\psi = \frac{\ln \eta}{2\gamma_T}$ satisfies the following with probability at least $1 - \delta$:*

$$R_T = O^*\big(\beta_{\bar{H}} \sqrt{T\gamma_T} + C\gamma_T^{3/2}\big). \tag{12}$$

3.2 Applications to Specific Confidence Bounds

Now we discuss specific choices of β_h satisfying Assumption 1, and the resulting final regret bounds.

We observe that the actions in each fixed epoch are sampled non-adaptively, and the resulting GP posterior formed only depends on the points in that epoch. As noted in [32], these conditions are sufficient to make use of the following confidence bounds for non-adaptive sampling.

Lemma 4. [45, Theorem 1] *When $\{x_i\}_{i=1}^t$ are selected independently of all the observations $\{y_i\}_{i=1}^t$, it holds for any fixed $x \in \mathcal{X}$ and any $t \geq 1$ with probability at least $1 - \delta$ that $|\mu_t(x) - f(x)| \leq \big(B + \frac{\sigma}{\sqrt{\lambda}}\sqrt{2\log\frac{1}{\delta}}\big)\sigma_t(x)$.*

For finite domains, applying the union bound leads to a choice of β_h for the proposed algorithm such that $\beta_{\bar{H}}$ only contributes to logarithmic terms in the cumulative regret.

Corollary 5. *Defining $\bar{\beta}_h(\delta) = B + \frac{\sigma}{\sqrt{\lambda}}\sqrt{2\log\frac{|\mathcal{X}|}{\delta}}$, we have that Assumption 1 holds with $\beta_h = \bar{\beta}_h(\delta_h)$ and $\delta_h = \frac{6\delta}{(h+1)^2\pi^2}$. Hence, with probability at least $1 - \delta$, Algorithm 1 with switching parameter $\eta > 1$, truncation parameter $\psi = \frac{\ln \eta}{2\gamma_T}$, and β_h as above achieves*

$$R_T = O^*\big(\sqrt{T\gamma_T} + C\gamma_T^{3/2}\big). \tag{13}$$

This corollary is obtained by noting that the error probability is at most δ as desired, since a union bound over $\mathcal{X}$ gives a per-epoch term of at most δ_h, and $\sum_{h=0}^{H-1} \delta_h \leq \sum_{h=0}^{\infty} \frac{6\delta}{(h+1)^2\pi^2} = (\sum_{h=0}^{\infty} \frac{1}{(h+1)^2})\frac{6\delta}{\pi^2} \leq \frac{\pi^2}{6} \cdot \frac{6\delta}{\pi^2} = \delta$.

For general (possibly continuous) domains, one option is to set β_h according to a widely-used confidence bound as follows, though we will shortly discuss improved choices.

Kernel	Lower Bound	Existing	Ours
Linear	$\sqrt{Td} + Cd$	$\sqrt{Td} + Cd^{3/2}$	$\sqrt{Td} + Cd^{3/2}$
SE	$\sqrt{T(\log T)^{d/2}} + C(\log T)^{d/2}$	$\sqrt{T}(\log T)^d + C\sqrt{T}(\log T)^{d/2}$	$\sqrt{T(\log T)^d} + C(\log T)^{3d/2}$
Matérn	$T^{\frac{\nu+d}{2\nu+d}} + C^{\frac{\nu}{d+\nu}} T^{\frac{d}{d+\nu}}$	$T^{\frac{2\nu+3d}{4\nu+2d}} + CT^{\frac{\nu+d}{2\nu+d}}$	$T^{\frac{\nu+d}{2\nu+d}} + CT^{\frac{3d}{4\nu+2d}}$

Table 1: Summary of regret bounds with constants and dimension-independent log factors omitted. For the SE and Matérn kernels, the upper bounds are from [8] and the lower bounds are from [10]. For the linear kernel, the existing bounds are from [9], except the $\sqrt{Td}$ lower bound which is from [17].

Lemma 6. [15, Theorem 2] *For any (possibly adaptive) sampling strategy, it holds with probability at least* $1-\delta$ *that* $|\mu_t(x) - f(x)| \leq \big(B + \sigma\sqrt{2(\gamma_t + 1 + \ln(1/\delta))}\big)\sigma_t(x)$ *for all* $x \in \mathcal{X}$ *and* $t \geq 1$.

By a similar argument to Corollary 5 and the fact that γ_t is increasing in t, we obtain the following.

Corollary 7. *If* $u_h \leq \bar{u}_h$ *almost surely, then defining* $\check{\beta}_h(\delta) = B + \sigma\sqrt{2(\gamma_{\bar{u}_h} + 1 + \ln(1/\delta))}$*, we have that Assumption 1 holds with* $\beta_h = \check{\beta}_h(\delta_h)$ *and* $\delta_h = \frac{6\delta}{(h+1)^2\pi^2}$*. Hence, with probability at least* $1-\delta$*, Algorithm 1 with a constant switching parameter* $\eta > 1$*, truncation parameter* $\psi = \frac{\ln \eta}{2\gamma_T}$*, and* β_h *as above achieves*

$$R_T = O^*\big(\sqrt{T}\gamma_T + C\gamma_T^{3/2}\big), \tag{14}$$

where we crudely selected $\bar{u}_h = T$.

While this regret bound can be significantly weaker than Corollary 5 due to the $O^*(\sqrt{T}\gamma_T)$ term, we can also obtain an analog of Corollary 5 (i.e., attaining the improved dependence in Eq. (13)) for continuous domains, under the mild assumption that functions in the RKHS are Lipschitz continuous (which is true for the kernels we consider below). A crude approach is to have the algorithm use a very fine discretization [26, 32], and a more sophisticated approach is to only discretize as part of the analysis [45]. The details can be found in the preceding references, and we avoid repeating them.

3.3 Comparisons to Existing Bounds

We specialize our regret bound in Eq. (13) to specific kernels by substituting $\gamma_T = O^*(d)$ for the linear kernel, $\gamma_T = O^*((\log T)^d)$ for the SE kernel, and $\gamma_T = O^*(T^{\frac{d}{2\nu+d}})$ for the Matérn kernel [43]. The resulting regret bounds are shown in Table 1 (omitting constants and dimension-independent log factors), along with the best known existing upper and lower bounds. We observe the following:

- For the linear kernel, we recover the recent upper bound of [9], and this is tight up to the presence of d vs. $d^{3/2}$ in the corrupted part.
- For the SE kernel, we match the lower bound of [10] up to small changes in the implied constant in each $(\log T)^{\Theta(d)}$ term. In contrast, the existing upper bound of [8] incurs a much larger $\sqrt{T}$ term in the corrupted part.
- For the Matérn kernel, compared to the existing result in [8], we obtain an improvement in the non-corrupted part recently established in [32], matching the non-corrupted lower bound. In the corrupted part, the existing result has a better exponent to T when $\nu < \frac{d}{2}$, whereas ours is better when $\nu > \frac{d}{2}$, in particular approaching zero (instead of $\frac{1}{2}$) as $\nu \to \infty$ and nearly matching the lower bound in this limit. However, when $\nu < \frac{d}{2}$ we find that the non-corrupted part in [8] is super-linear in T, making the bound trivial. Hence, our bound is better whenever non-trivial scaling is attained.

The bounds based on a reduction to linear bandits, which we derive in Appendix E, are omitted in Table 1. We briefly note that they are able to provide a similar upper bound to our main one under the SE kernel, but are always strictly worse under the Matérn kernel.

3.4 Implications for the Unknown C Setting

While we have focused on the case of known C, an idea from a concurrent work [23] (on linear bandits) can be used to transfer our main result to a setting with unknown C.

The idea is that if the parameter C is used by the algorithm but C_{true} is the amount of corruption actually used by the adversary, then the analysis goes through unchanged as long as $C \geq C_{\text{true}}$.

Hence, we may cautiously choose a large value of C to cover more values of C_{true}. As an important special case, we may choose C such that the corrupted and uncorrupted regret terms are of the same order; for instance, in (13), setting $C = O\big(\frac{\sqrt{T}}{\gamma_T}\big)$ gives $R_T = O^*(\sqrt{T\gamma_T})$. Hence, we find that any corruption level C_{true} up to $O\big(\frac{\sqrt{T}}{\gamma_T}\big)$ only affects the constant (or possibly logarithmic) factors, and the precise corruption level does not need to be known.

For particularly smooth kernels such as linear and SE (with constant dimension), the scaling $O\big(\frac{\sqrt{T}}{\gamma_T}\big)$ reduces to $O^*(\sqrt{T})$. This may not seem as high as ideal, but at least in the case of linear bandits, it is known to be the best we can hope for unless the algorithm attains significantly higher uncorrupted regret [9, 23]. Specifically, if optimal $O^*(\sqrt{T})$ uncorrupted regret is attained, then linear regret is unavoidable when $C = \omega(\sqrt{T})$. See [23] for similar statements with the dependence on d included.

The overall picture remains less complete for general kernels, but the preceding discussion reveals that our results for known C do have important implications for the unknown C setting.

4 Experiments

We experimentally evaluate the performance of our proposed algorithm, along with two baselines, one robust and one non-robust. Our experiments serve as a proof of concept for our proposed approach, but also highlight possible remaining gaps between theory and practice, e.g., arising from large constant factors in the regret bounds. We emphasize that our contributions are primarily theoretical.

Algorithms. We consider the following three algorithms:

1. RGP-PE: Robust GP-Phased Elimination with constant β_h; this is a slight variation of Corollary 5 in which the number of epochs H turns out to be a small constant in our experiments.
2. GP-UCB: a representative non-robust fully sequential algorithm with slowly growing β_t, where $t \in [T]$ [43, Algorithm 1].
3. RGP-UCB: the robust version of GP-UCB with slowly growing β_t [8, Algorithm 1], where the only difference from GP-UCB is that the theoretical coefficient of σ_{t-1} in the UCB is $\beta_t + \frac{C}{\sqrt{\lambda}}$.

We found the term $\beta_h + \frac{C\sqrt{u_h}}{l_h \psi \lambda}$ multiplying $\sigma^{(h)}$ in Algorithm 1 to be overly conservative, so we instead replace it by $\beta_h + b \cdot \frac{C}{\sqrt{u_h}}$ (since l_h and u_h are similar, we replace $\frac{\sqrt{u_h}}{l_h}$ by $\frac{1}{\sqrt{u_h}}$), where $b \in (0, 1]$ is an additional parameter controlling the degree of exploration and robustness. Similarly, in RGP-UCB we use the coefficient $\beta_t + b \cdot \frac{C}{\sqrt{\lambda}}$. The remaining parameters β_h and β_t are specified below.

Synthetic Function. We produce a synthetic 2D function f_1, shown in Figure 4 of the supplementary material, which is randomly sampled from a Gaussian Process with zero mean and the SE kernel with lengthscale $l = 0.5$. The domain $\mathcal{X}$ of f_1 contains 100 points obtained by evenly splitting $[-5, 5]^2$ into a 10×10 grid. We use the true kernel as the prior for all three algorithms, and use $\beta_h = 4$ for RGP-PE, and $\beta_t = \sqrt{\log t}/2$ for GP-UCB and RGP-UCB.

Robot Pushing Objective Function. We consider the deterministic robot pushing objective function on a 2D plane introduced in [48], which aims to find suitable parameters to push an object to the target location r_g. We use the Robot3d function, which takes the robot location (r_x, r_y) and pushing duration t_r as a 3D input, and outputs the reversed distance between the pushed robot location and the target location r_g, i.e.,

$$\text{Robot3D}(r_x, r_y, t_r) = 5 - \|\text{push}(r_x, r_y, t_r) - r_g\|,$$

where $\text{push}(\cdot)$ outputs the pushed robot location.

We let the domain $\mathcal{X}$ contain 100 points (r_x, r_y, t_r) randomly sampled from $[-5, 5]^2 \times [1, 30]$, and the target location r_g is set to be $(3, 2)$. Since the lengthscale of the SE kernel with maximum likelihood given the noiseless data is $1.94 \approx 2$, we use the SE kernel with $l = 2$ as prior for all three algorithms. We found it beneficial for all algorithms to be slightly more explorative for this function, and accordingly use $\beta_h = 6$ for KE and $\beta_t = 2\sqrt{\log t}$ for GP-UCB and RGP-UCB.

Attack Methods. We consider the following five attack methods, which continue until the corruption budget is exhausted:

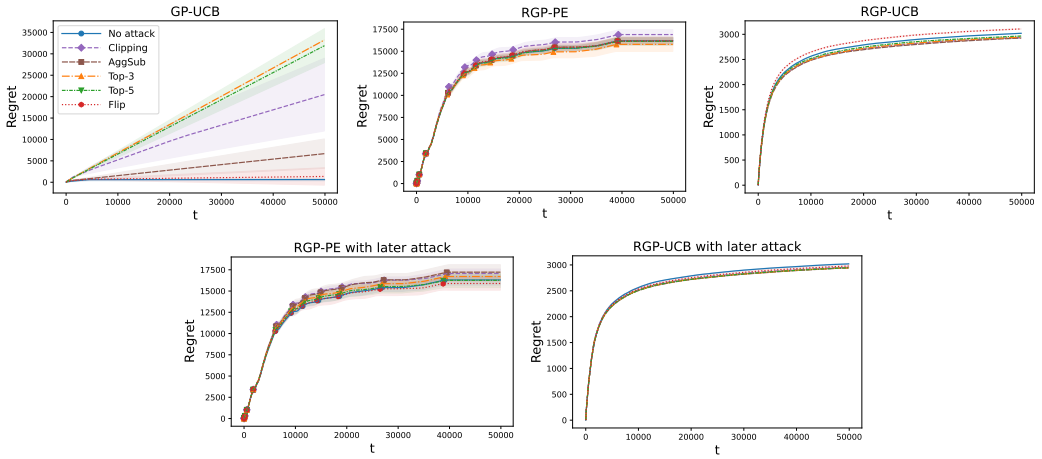

Figure 2: Performance on f_1 with $C = 50$. We observe that GP-UCB incurs linear regret for several attacks, whereas the other algorithms exhibit robustness to all of the attacks.

- Clipping: This attack proposed in [22] perturbs f and produces another reward function $\widetilde{f}$ whose optima are in some region $\mathcal{R}_{\text{target}}$ that does not contain x^* by setting

$$\widetilde{f}(x) = \begin{cases} f(x) & x \in \mathcal{R}_{\text{target}}, \\ \min\{f(x), f(\widetilde{x}^*) - \Delta\} & x \notin \mathcal{R}_{\text{target}}, \end{cases}$$

  where $\widetilde{x}^* = \arg\max_{x \in \mathcal{R}_{\text{target}}} f(x)$. We let $\Delta = 0.5$ and choose $\mathcal{R}_{\text{target}} = \{(x_1, x_2) \in \mathcal{X} : x_1 \leq x_2\}$ for f_1, and $\mathcal{R}_{\text{target}} = \{(r_x, r_y, t_r) \in \mathcal{X} : r_x \geq 0\}$ for the function Robot3D.
- Aggressive Subtraction (AggSub): This attack proposed in [22] sets

$$\widetilde{f}(x) = \begin{cases} f(x) & x \in \mathcal{R}_{\text{target}}, \\ f(x) - h_{\max} & x \notin \mathcal{R}_{\text{target}}, \end{cases}$$

  for some $h_{\max} > f(x^*) - f(\widetilde{x}^*)$. We use the same $\mathcal{R}_{\text{target}}$ as the Clipping attack, and let $h_{\max} = 1$ for f_1 and $h_{\max} = 3$ for Robot3D.
- Top-K: When x is one of the top K remaining actions, this attack perturbs the reward down to -1. We consider both $K = 3$ and $K = 5$.
- Flip: This attack simply flips the reward from $f(x)$ to $-f(x)$. Both this attack and the previous one are variations of attacks considered for linear bandits in [9].

For the algorithms, we consider $C = 50$ and $C = 100$. By default, the attack starts at $t = 1$, but for the robust algorithms RGP-PE and RGP-UCB, we also conduct experiments with a *later* attack, where (i) the attack in RGP-PE starts when at least one action is eliminated from the domain; and (ii) the attack in RGP-UCB starts when at least one action has UCB strictly lower than $\max_{x \in \mathcal{X}} \text{LCB}(x)$.

We let $T = 50000$,[2] $\sigma = 0.02$, and $\lambda = 1$ for all three algorithms, $b = 0.1$ for RGP-PE and RGP-UCB, and $\psi = 0.5, \eta = 2$ for RGP-PE.The results are produced by performing 10 trials and plotting the average cumulative regret, with error bars indicating one standard deviation.

Comparison of Algorithms. As shown in Figures 2 and 3, the non-robust algorithm GP-UCB succeeds when no attack is applied. However, the cumulative regret for f_1 associated with the Clipping, AggSub, Top-3, and Top-5 attacks grow linearly, indicating that these four attacks succeed in driving GP-UCB towards a suboptimal action. Similarly, the Top-3 and Top-5 attacks incur linear regret for Robot3D. In contrast, we find that RGP-PE has only one action remaining at the end of the 13th epoch, and manages to defend against all five attack methods for both functions.

The baseline robust algorithm RGP-UCB also successfully defends against all the attacks, and generally has lower cumulative regret than RGP-PE, despite RGP-PE having a stronger regret

[2] As we mentioned previously, this large value of T is feasible due to the computation time only scaling with respect to the number of *unique* points [12], which is much smaller than T.

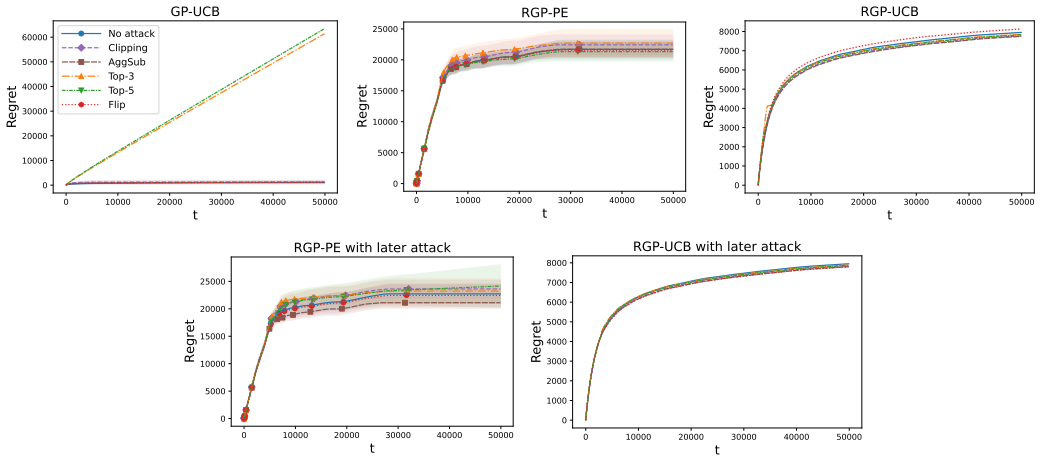

Figure 3: Performance on Robot3D with $C = 100$. We observe that GP-UCB incurs linear regret for two attacks, whereas the other algorithms exhibit robustness to all of the attacks.

guarantee. There are at least two possibly reasons for this: (i) The analysis of RGP-UCB in [8] could be loose, with a tighter analysis potentially giving an additive dependence similar to Theorem 3, and (ii) the strong scaling laws in our theory may still leave room for improvements in the constant factors (or logarithmic). Further addressing these findings remains an interesting direction for future work. We note that even in the more specialized problem of corrupted stochastic linear bandits, analogous practical limitations of a phased elimination algorithm were observed in [9].

Later Attack. We observe that RGP-PE and RGP-UCB are also able to defend against the later attack, and their performance is similar to when the attack starts from the beginning. There are only two trials of RGP-PE (budget $C = 100$ and Top-5 attack on Robot3D in Figure 3), in which the only action remaining at the end of the 13th epoch is slightly suboptimal. In Appendix F, we additionally show the experiment results for f_1 with $C = 100$, and Robot3D with $C = 50$.

5 Conclusion

We have provided a new algorithm for corruption-tolerant GP bandits based on phased elimination, incorporating a key idea of *rare switching* based on a certain condition relating to the information gain, along with a robust estimator, enlarged confidence bounds, and truncation to ensure a minimal number of plays of each selected action. Our regret bound recovers the best known existing bound under the linear kernel, is provably near-optimal under the SE kernel, and improves on the best existing bound in all cases where the latter is non-trivial.

Acknowledgment

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme grant agreement No 815943. J. Scarlett was supported by the Singapore National Research Foundation (NRF) under grant number R-252-000-A74-281.

References

[1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Conference on Neural Information Processing Systems*, 2011.

[2] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. pages 2312–2320, 2011.

[3] Idan Amir, Idan Attias, Tomer Koren, Yishay Mansour, and Roi Livni. Prediction with corrupted expert advice. *Conference on Neural Information Processing Systems*, 2020.

[4] Kiarash Banihashem, Adish Singla, and Goran Radanovic. Defense against reward poisoning attacks in reinforcement learning. *arXiv preprint arXiv:2102.05776*, 2021.

[5] Justin J. Beland and Prasanth B. Nair. Bayesian optimization under uncertainty. NIPS BayesOpt 2017 workshop, 2017.

[6] Ilija Bogunovic and Andreas Krause. Misspecified Gaussian process bandit optimization. *Conference on Neural Information Processing Systems*, 34, 2021.

[7] Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. Adversarially robust optimization with Gaussian processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5760–5770, 2018.

[8] Ilija Bogunovic, Andreas Krause, and Jonathan Scarlett. Corruption-tolerant Gaussian process bandit optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

[9] Ilija Bogunovic, Arpan Losalka, Andreas Krause, and Jonathan Scarlett. Stochastic linear bandits robust to adversarial attacks. In *International Conference on Artificial Intelligence and Statistics*, pages 991–999, 2021.

[10] Xu Cai and Jonathan Scarlett. On lower bounds for standard and robust Gaussian process bandit optimization. In *International Conference on Machine Learning*, 2021.

[11] Sait Cakmak, Raul Astudillo Marban, Peter Frazier, and Enlu Zhou. Bayesian optimization of risk measures. In *Conference on Neural Information Processing Systems*, 2020.

[12] Daniele Calandriello, Luigi Carratino, Alessandro Lazaric, Michal Valko, and Lorenzo Rosasco. Scaling Gaussian process optimization by evaluating a few unique candidates multiple times. 2022.

[13] Romain Camilleri, Kevin Jamieson, and Julian Katz-Samuels. High-dimensional experimental design and kernel bandits. In *International Conference on Machine Learning*, 2021.

[14] Yifang Chen, Simon Shaolei Du, and Kevin Jamieson. Corruption robust active learning. In *Conference on Neural Information Processing Systems*, 2021.

[15] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning (ICML)*, pages 844–853, 2017.

[16] Thanh Dai Nguyen, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Stable Bayesian optimization. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 578–591. Springer, 2017.

[17] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, 2008.

[18] Audrey Durand, Odalric-Ambrym Maillard, and Joelle Pineau. Streaming kernel regression with provably adaptive mean, variance, and regularization. *The Journal of Machine Learning Research*, 19(1):650–683, 2018.

[19] Minbo Gao, Tianle Xie, Simon S Du, and Lin F Yang. A provably efficient algorithm for linear markov decision process with low switching cost. *arXiv preprint arXiv:2101.00494*, 2021.

[20] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chem. Sci."*, 11:577–586, 2020.

[21] Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Conference on Learning Theory (COLT)*, 2019.

[22] Eric Han and Jonathan Scarlett. Adversarial attacks on Gaussian process bandits. In *International Conference on Machine Learning*, 2021.

[23] Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. Nearly optimal algorithms for linear contextual bandits with adversarial corruptions. *Conference on Neural Information Processing Systems*, 2022.

[24] Shinji Ito. On optimal robustness to adversarial corruption in online decision problems. *Conference on Neural Information Processing Systems*, 2021.

[25] Shinji Ito, Taira Tsuchiya, and Junya Honda. Adversarially robust multi-armed bandit algorithm with variance-dependent regret bounds. In *Conference on Learning Theory (COLT)*, 2022.

[26] David Janz, David R. Burt, and Javier González. Bandit optimisation of functions in the Matérn kernel RKHS. In *International Conference on Artificial Intelligence and Statistics*, 2020.

[27] Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. https://arxiv.org/abs/1807.02582, 2018.

[28] Johannes Kirschner and Andreas Krause. Bias-robust bayesian optimization via dueling bandits. In *International Conference on Machine Learning*, 2021.

[29] Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2174–2184. PMLR, 2020.

[30] Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in RL with a generative model. In *International Conference on Machine Learning*, 2020.

[31] Yingkai Li, Edmund Y Lou, and Liren Shan. Stochastic linear optimization with adversarial corruption. *arXiv preprint arXiv:1909.02109*, 2019.

[32] Zihan Li and Jonathan Scarlett. Gaussian process bandit optimization with few batches. In *International Conference on Artificial Intelligence and Statistics*, 2022.

[33] Junyan Liu, Shuai Li, and Dapeng Li. Cooperative stochastic multi-agent multi-armed bandits robust to adversarial corruptions. *arXiv preprint arXiv:2106.04207*, 2021.

[34] Daniel J Lizotte, Tao Wang, Michael H Bowling, and Dale Schuurmans. Automatic gait optimization with Gaussian process regression. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 944–949, 2007.

[35] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *ACM Symposium on Theory of Computing (STOC)*, pages 114–122. ACM, 2018.

[36] Thodoris Lykouris, Max Simchowitz, Alex Slivkins, and Wen Sun. Corruption-robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, pages 3242–3245. PMLR, 2021.

[37] Anastasia Makarova, Ilnura Usmanova, Ilija Bogunovic, and Andreas Krause. Risk-averse heteroscedastic Bayesian optimization. *Conference on Neural Information Processing Systems*, 2021.

[38] Ruben Martinez-Cantin, Kevin Tee, and Michael McCourt. Practical Bayesian optimization in the presence of outliers. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.

[39] Quoc Phong Nguyen, Zhongxiang Dai, Bryan Kian Hsiang Low, and Patrick Jaillet. Value-at-risk optimization with Gaussian processes. In *International Conference on Machine Learning*.

[40] J. Nogueira, R. Martinez-Cantin, A. Bernardino, and L. Jamone. Unscented Bayesian optimization for safe robot grasping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.

[41] Sayak Ray Chowdhury and Aditya Gopalan. Bayesian optimization under heavy-tailed payoffs. *Conference on Neural Information Processing Systems*, 2019.

[42] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In *Conference on Neural information Processing Systems*, pages 2951–2959, 2012.

[43] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning (ICML)*, 2010.

[44] Sho Takemori and Masahiro Sato. Approximation theory based methods for rkhs bandits. In *International Conference on Machine Learning*, 2021.

[45] Sattar Vakili, Nacime Bouziani, Sepehr Jalali, Alberto Bernacchia, and Da shan Shiu. Optimal order simple regret for Gaussian process bandits. In *Conference on Neural information Processing Systems*, 2021.

[46] Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in Gaussian process bandits. In *Conference on Neural information Processing Systems*, 2021.

[47] Tianhao Wang, Dongruo Zhou, and Quanquan Gu. Provably efficient reinforcement learning with linear function approximation under adaptivity constraints. *Conference on Neural Information Processing Systems*, 2021.

[48] Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *International Conference on Machine Learning (ICML)*, pages 3627–3635, 2017.

[49] Chen-Yu Wei, Christoph Dann, and Julian Zimmert. A model selection approach for corruption robust reinforcement learning. In *Conference on Algorithmic Learning Theory*, 2022.

[50] Heyang Zhao, Dongruo Zhou, and Quanquan Gu. Linear contextual bandits with adversarial corruptions. *arXiv preprint arXiv:2110.12615*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]

(b) Did you mention the license of the assets? [N/A]

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary Material

A Robust Phased Elimination Algorithm for Corruption-Tolerant Gaussian Process Bandits

(Bogunovic/Li/Krause/Scarlett, NeurIPS 2022)

A Preliminaries

Here, we outline some useful and well-known results and definitions typically used in kernelized/GP bandit (Bayesian optimization) algorithms.

RKHS and kernel functions. We denote by $\mathcal{H}_k$ the reproducing kernel Hilbert space (RKHS) corresponding to the kernel k, defined as a Hilbert space of functions equipped with an inner product $\langle \cdot, \cdot \rangle_k$, satisfying the reproducing property, i.e., $\langle f(\cdot), k(\cdot, x) \rangle_k = f(x), \forall x \in \mathcal{X}, \forall f \in \mathcal{H}_k$.

Since we assume that the kernel is bounded (i.e., $k(x, x') \leq 1$), continuous, and has a compact domain (namely, $D = [0, 1]^d$), the conditions of Mercer's theorem are satisfied [27], and the kernel admits a countably infinite (or finite) dimensional feature space, i.e., there exists $\{(\lambda_m, \phi_m)\}_{m=1}^{\infty}$ such that $k(x, x') = \sum_{m=1}^{\infty} \lambda_m \phi_m(x) \phi_m(x')$ where the $\phi_m(\cdot)$ are eigenfunctions, and the $\lambda_m \geq 0$ are eigenvalues. We form an infinite-dimensional feature vector as follows:

$$\phi(x) = (\sqrt{\lambda_1}\phi_1(x), \sqrt{\lambda_2}\phi_2(x), \dots), \tag{15}$$

which yields $k(x, x') = \phi(x)^T \phi(x')$. As stated in the main text, we assume that the RKHS norm is upper bounded by some constant $B > 0$.

The following lemma provides a useful expression for $\sigma_t^2(x)$. This result is fairly standard, but for completeness, we provide a short proof. Here and subsequently, we use I to denote the infinite-dimensional identity matrix in feature space.

Lemma 8. *Defining $\Phi_t = [\phi(x_1), \dots, \phi(x_t)]^T$, we have*

$$\sigma_t^2(x) = \lambda \phi(x)^T (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x). \tag{16}$$

Proof. We can rewrite $\sigma_t^2(x)$ as follows,

$$\begin{aligned}
\sigma_t^2(x) &= k(x, x) - k_t(x)^T (K_t + \lambda I_t)^{-1} k_t(x) && (17)\\
&= \phi(x)^T \phi(x) - \phi(x)^T \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} \Phi_t \phi(x) && (18)\\
&= \phi(x)^T \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} \Phi_t \phi(x) + \lambda \phi(x)^T (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x) \\
&\qquad - \phi(x)^T \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} \Phi_t \phi(x) && (19)\\
&= \lambda \phi(x)^T (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x), && (20)
\end{aligned}$$

where Eq. (19) uses $\phi(x) = \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} \Phi_t \phi(x) + \lambda (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x)$, which can be obtained as follows.

$$\begin{aligned}
(\Phi_t^T \Phi_t + \lambda I)\phi(x) &= \Phi_t^T \Phi_t \phi(x) + \lambda \phi(x) && (21)\\
\phi(x) &= (\Phi_t^T \Phi_t + \lambda I)^{-1} \Phi_t^T \Phi_t \phi(x) + \lambda (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x) && (22)\\
&= \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} \Phi_t \phi(x) + \lambda (\Phi_t^T \Phi_t + \lambda I)^{-1} \phi(x), && (23)
\end{aligned}$$

where the last step follows from the standard push-through identity $(\Phi_t^T \Phi_t + \lambda I)\Phi_t^T = \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)$ (e.g., [15, Eq. (12)]), which implies $\Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1} = (\Phi_t^T \Phi_t + \lambda I)^{-1} \Phi_t^T$.

□

Some of the most commonly used kernels are:

- Linear kernel: $k_{\text{lin}}(x, x') = x^T x'$,
- Squared exponential kernel: $k_{\text{SE}}(x, x') = \exp\big(-\frac{\|x - x'\|^2}{2l^2}\big)$,

- Matérn kernel: $k_{\text{Mat}}(x, x') = \frac{2^{1-\nu}}{\Gamma(\nu)}\Big(\frac{\sqrt{2\nu}\|x-x'\|}{l}\Big)J_\nu\Big(\frac{\sqrt{2\nu}\|x-x'\|}{l}\Big)$,

where l denotes the length-scale hyperparameter, $\nu > 0$ is an additional hyperparameter that dictates the smoothness, and $J(\cdot)$ and $\Gamma(\cdot)$ denote the modified Bessel function and the Gamma function, respectively.

Maximum information gain. The maximum information gain is defined as [43]

$$\begin{aligned}\gamma_t &:= \max_{A\subseteq\mathcal{X}:|A|=t} I(f_A; y_A)\\ &= \max_{x_1,\ldots,x_t} \frac{1}{2}\log\det(I_t + \lambda^{-1}K_t),\end{aligned}$$

where $f_A = [f(x_t)]_{x_t\in A}$, $y_A = [y_t]_{x_t\in A}$, and $I(\cdot;\cdot)$ denotes mutual information. The maximum information gain quantifies the maximum reduction in uncertainty about f after t observations. The following upper bounds for specific kernels have been shown previously [43, 46]:

- Linear kernel: $\gamma_t^{\text{lin}} = O^*(d\log t)$,
- Squared exponential kernel: $\gamma_t^{\text{SE}} = O^*((\log t)^d)$,
- Matérn kernel: $\gamma_t^{\text{Mat}} = O^*\big(t^{\frac{d}{2\nu+d}}\big)$.

The following lemma shows that $\sum_{t=1}^T \sigma_{t-1}(x_t)$ can be upper bounded in terms of γ_T.

Lemma 9. *With $\sigma_{t-1}(x_t)$ denoting the posterior standard deviation at x_t based on $(x_1, \ldots, x_{t-1})$, we have*

$$\sum_{t=1}^T \sigma_{t-1}(x_t) \le \sqrt{T\sum_{t=1}^T \sigma_{t-1}^2(x_t)} \le \sqrt{\frac{2}{\log(1+\lambda^{-1})}T\gamma_T} \le \sqrt{(2\lambda+1)T\gamma_T}.$$

Proof. The first inequality follows by Cauchy-Schwartz inequality; the second inequality follows from Lemma 5.4 of [43]; the last inequality follows since $(2\lambda+1)\log(1+\lambda^{-1}) > 2$ for $\lambda > 0$. □

B Corrupted Confidence Bounds

For convenience, we first restate our main assumption regarding non-corrupted confidence bounds.

Assumption 1 (Regular confidence bounds)**.** *Let $\mu^{(h)}(x)$ and $\sigma^{(h)}(x)$ denote the posterior mean and standard deviation computed (hypothetically) using only the non-corrupted observations $\{(x_i, y_i)\}_{i=1}^{u_h}$ in epoch h using Eqs.* (5) *and* (6)*. We assume that given $\delta \in (0,1)$, there exists a sequence of parameters $\beta_h = \beta_h(\delta)$ which is non-decreasing in h and yields with probability at least $1-\delta$ that*

$$|\mu^{(h)}(x) - f(x)| \le \beta_h\sigma^{(h)}(x) \tag{9}$$

simultaneously for all $h \ge 0$ and $x \in \mathcal{X}$.

In this appendix, we prove Lemma 2, which is restated as follows.

Lemma 2 (Corrupted confidence bounds)**.** *Under Assumption 1, let $\widetilde{\mu}^{(h)}(x)$ denote the posterior mean based on only the corrupted observations $\{(x_i, \widetilde{y}_i)\}_{i=1}^{u_h}$ in epoch h using Eq.* (7)*, and let $u_{\min} \ge 1$ denote the minimum number of times any single action from $\{x_i\}_{i=1}^{u_h}$ is played, i.e., $u_{\min} = \min_{x\in\{x_1,\ldots,x_{u_h}\}}\sum_{i=1}^{u_h}\mathbb{1}\{x_i = x\}$. Then, with probability at least $1-\delta$, it holds for all $x \in \mathcal{X}$ and $h \ge 0$ that*

$$|\widetilde{\mu}^{(h)}(x) - f(x)| \le \Big(\beta_h + \frac{C\sqrt{u_h}}{u_{\min}\lambda}\Big)\sigma^{(h)}(x). \tag{10}$$

Proof. For simplicity, we denote the epoch length u_h by t in this proof, and use $\mu_t(\cdot)$, $\widetilde{\mu}_t(\cdot)$, and $\sigma_t(\cdot)$ to denote $\mu^{(h)}(\cdot)$, $\widetilde{\mu}^{(h)}(\cdot)$, and $\sigma^{(h)}(\cdot)$, respectively. Thus, here $\sigma_t(\cdot)$ is defined with respect to the $t = u_h$ sampled points, whereas Algorithm 1 only computes the posterior variance with respect to the points selected in the for loop, of which there are l_h (possibly strictly fewer than u_h). This part of the analysis only requires the former notion, so there should be no confusion between the two.

We first recall the definition of the robust-corrupted mean estimator from Eq. (7), i.e.,

$$\widetilde{\mu}_t(x) = k_t(x)^T (K_t + \lambda I_t)^{-1} \widetilde{Y}_t, \tag{24}$$

where $\widetilde{Y}_t \in \mathbb{R}^t$ and $\widetilde{Y}_t[i] = \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}\widetilde{y}_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}}$ for $i \in [t]$. We use $z_t(x;\lambda) \in \mathbb{R}^t$ to denote $k_t(x)^T(K_t + \lambda I_t)^{-1}$ which implies $\widetilde{\mu}_t(x) = \sum_{i=1}^t z_t(x;\lambda)[i] \cdot \widetilde{Y}_t[i]$.

We will also use the following equivalent feature-based expression: $z_t(x;\lambda) = k_t(x)^T(K_t + \lambda I)^{-1} = \phi(x)^T \Phi_t^T (\Phi_t \Phi_t^T + \lambda I_t)^{-1}$, where $k(x,x') = \phi(x)^T\phi(x')$, $\phi(x) \in \mathcal{H}_k(\mathcal{X})$ for every $x \in \mathcal{X}$, and $\Phi_t = (\phi(x_{t'}))_{t' \le t}$ denotes the matrix of (potentially infinite-dimensional) features placed in t rows. Finally, recalling that I denotes the infinite-dimensional identity matrix in feature space, we also have

$$z_t(x;\lambda) = \phi(x)^T (\Phi_t^T \Phi_t + \lambda I)^{-1} \Phi_t^T, \tag{25}$$

which follows from the standard push-through identity $\Phi_t^T \left(\Phi_t \Phi_t^T + \lambda I_t\right)^{-1} = \left(\Phi_t^T \Phi_t + \lambda I\right)^{-1} \Phi_t^T$ (e.g., see Eq. (12) of [15]).

We proceed to analyze the corrupted estimator $\widetilde{\mu}_t(x)$:

$$\begin{aligned}
\widetilde{\mu}_t(x) &= \sum_{i=1}^t z_t(x;\lambda)[i]\ \widetilde{Y}_t[i] && (26)\\
&= \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}\widetilde{y}_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] && (27)\\
&= \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}(f(x_i) + \epsilon_j + c_j)}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] && (28)\\
&= \sum_{i=1}^t f(x_i) z_t(x;\lambda)[i] + \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}\epsilon_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] + \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] && (29)\\
&= \sum_{i=1}^t f(x_i) z_t(x;\lambda)[i] + \sum_{i=1}^t \epsilon_i z_t(x;\lambda)[i] + \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] && (30)\\
&= \sum_{i=1}^t (f(x_i) + \epsilon_i) z_t(x;\lambda)[i] + \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i] && (31)\\
&= \mu_t(x) + \sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i = x_j\}} z_t(x;\lambda)[i]. && (32)
\end{aligned}$$

Here, we used the definition of $\widetilde{Y}_t[i]$ in Eq. (27) and the corrupted observation $\widetilde{y}_j$ corresponding to $x_j = x_i$ at time j in Eq. (28), while Eq. (29) follows from rearranging. The proof of Eq. (30) is deferred to the next paragraph. Finally, Eq. (32) follows from the definition of the noisy stochastic observation $y_i = f(x_i) + \epsilon_i$ and the definition of the standard (non-corrupted) mean estimator from Eq. (5).

To prove Eq. (30), we define $\widetilde{\epsilon}_t \in \mathbb{R}^t$ such that $\widetilde{\epsilon}_t[i] = \frac{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}\epsilon_j}{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}}$ for $i \in [t]$, and use $u_t(x)$ to denote $\sum_{j=1}^t \mathbb{1}\{x = x_j\}$, i.e., the number of times action x was played during the t rounds. Then,

$$\sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}\epsilon_j}{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}} z_t(x;\lambda)[i] = z_t(x;\lambda)\widetilde{\epsilon}_t \tag{33}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\Phi_t^T\widetilde{\epsilon}_t \tag{34}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\sum_{i=1}^t \phi(x_i)\widetilde{\epsilon}_t[i] \tag{35}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\sum_{x\in\mathcal{X}, u_t(x)\neq 0} u_t(x)\phi(x)\frac{\sum_{j=1}^t \mathbb{1}\{x=x_j\}\epsilon_j}{u_t(x)} \tag{36}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\sum_{x\in\mathcal{X}, u_t(x)\neq 0} \phi(x)\sum_{j=1}^t \mathbb{1}\{x = x_j\}\epsilon_j \tag{37}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\sum_{j=1}^t \phi(x_j)\epsilon_j \tag{38}$$

$$= \phi(x)^T(\Phi_t^T\Phi_t + \lambda I)^{-1}\Phi_t^T\epsilon_t \tag{39}$$

$$= z_t(x;\lambda)\epsilon_t = \sum_{i=1}^t \epsilon_i z_t(x;\lambda)[i], \tag{40}$$

where Eq. (34) holds due to Eq. (25), and Eq. (36) uses the definitions of $\widetilde{\epsilon}_t$ and $u_t(x)$, and (38)–(40) are analogous to (33)–(35) in the opposite order.

By rearranging Eq. (32), it follows that we can bound the absolute difference between the corrupted mean estimator and the standard one as follows:

$$|\widetilde{\mu}_t(x) - \mu_t(x)| \leq \Big|\sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}} z_t(x;\lambda)[i]\Big|. \tag{41}$$

Next, we proceed to analyze the right hand side term. We use C_t to denote a vector in $\mathbb{R}^t$ such that $C_t[i] = \frac{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}}$ for every $i \in [t]$. Then, continuing from Eq. (41), we have

$$\Big|\sum_{i=1}^t \frac{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}c_j}{\sum_{j=1}^t \mathbb{1}\{x_i=x_j\}} z_t(x;\lambda)[i]\Big| = \Big|\phi(x)^T \underbrace{(\Phi_t^T\Phi_t + \lambda I)^{-1}}_{:=\Gamma_t^{-1}}\Phi_t^T C_t\Big| \tag{42}$$

$$= \Big|\sum_{i=1}^t C_t[i]\phi(x)^T\Gamma_t^{-1}\phi(x_i)\Big|, \tag{43}$$

where we again used the form of z_t given in Eq. (25).

Let $C_t(x) = \sum_{j=1}^{t} \mathbb{1}\{x = x_j\}c_j$ for $x \in \mathcal{X}$. Then, we can rewrite (43) as

$$\Big|\sum_{i=1}^{t} \frac{\sum_{j=1}^{t} \mathbb{1}\{x_i=x_j\}c_j}{\sum_{j=1}^{t} \mathbb{1}\{x_i=x_j\}} z_t(x;\lambda)[i]\Big| = \Big|\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} \frac{C_t(x')}{u_t(x')} u_t(x')\phi(x)^T\Gamma_t^{-1}\phi(x')\Big| \tag{44}$$

$$\leq \sum_{x'\in\mathcal{X}, u_t(x')\neq 0} \frac{C}{u_t(x')} u_t(x')\Big|\phi(x)^T\Gamma_t^{-1}\phi(x')\Big| \tag{45}$$

$$\leq \frac{C}{u_{\min}} \sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\Big|\phi(x)^T\Gamma_t^{-1}\phi(x')\Big| \tag{46}$$

$$\leq \frac{C}{u_{\min}} \sqrt{\Big(\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\Big)\phi(x)^T \sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\Gamma_t^{-1}\phi(x')\phi(x')^T\Gamma_t^{-1}\phi(x)} \tag{47}$$

$$\leq \frac{C}{u_{\min}} \sqrt{\Big(\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\Big)\phi(x)^T \sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\Gamma_t^{-1}\big(\phi(x')\phi(x')^T + \tfrac{\lambda}{t}I\big)\Gamma_t^{-1}\phi(x)} \tag{48}$$

$$= \frac{C}{u_{\min}} \sqrt{\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\|\phi(x)\|^2_{\Gamma_t^{-1}}} \tag{49}$$

$$= \frac{C}{u_{\min}} \sqrt{t}\|\phi(x)\|_{\Gamma_t^{-1}} = \frac{C\sqrt{t}}{\lambda u_{\min}}\sigma_t(x), \tag{50}$$

where:

- Eq. (45) holds since $C \geq |C_t(x)|$ for every $x \in \mathcal{X}$.
- Eq. (46) follows from the definition of $u_{\min}$ in the lemma statement.
- To obtain Eq. (47), we multiply and divide by $\sum_{x\in\mathcal{X}, u_t(x)\neq 0} u_t(x)$ and apply $\mathbb{E}[|X|] \leq \sqrt{\mathbb{E}[X^2]}$ considering the distribution $\frac{u_t(x')}{\sum_{x\in\mathcal{X}, u_t(x)\neq 0} u_t(x)}$. (Note also that, in generic vector-matrix notation, $(a^TMb)^2 = a^TMbb^TMa$ when M is a symmetric matrix.)
- To obtain Eq. (49), we use $\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\frac{\lambda}{t}I = \lambda I$ (i.e., $\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x') = t$), and note that $\Gamma_t = \big(\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\phi(x')\phi(x')^T\big) + \lambda I$. Combining these facts gives $\sum_{x'\in\mathcal{X}, u_t(x')\neq 0} u_t(x')\big(\phi(x')\phi(x')^T + \frac{\lambda}{t}I\big) = \Gamma_t$, which cancels with one of the Γ_t^{-1} terms. The remaining quantity $\phi(x)^T\Gamma_t^{-1}\phi(x)$ is precisely the definition of $\|\phi(x)\|^2_{\Gamma_t^{-1}}$.
- Finally, Eq. (50) holds since
$$\|\phi(x)\|^2_{\Gamma_t^{-1}} = \phi(x)^T\Gamma_t^{-1}\phi(x) = \lambda^{-1}\sigma_t^2(x), \tag{51}$$
which holds due to Eq. (16).

Conditioned on the event in Assumption 1, the final result then follows since

$$|\widetilde{\mu}_t(x) - f(x)| \leq |\mu_t(x) - f(x)| + |\widetilde{\mu}_t(x) - \mu_t(x)| \leq \Big(\beta_h + \tfrac{C\sqrt{t}}{\lambda u_{\min}}\Big)\sigma_t(x), \tag{52}$$

where we apply Assumption 1 and Eq. (50) to upper bound $|\mu_t(x) - f(x)|$ and $|\widetilde{\mu}_t(x) - \mu_t(x)|$, respectively.

□

C Auxiliary Results

In the following, we recall the notation in Algorithm 1, particularly the truncation parameter $\psi > 0$. In addition, in accordance with the algorithm statement, quantities such as $\sigma_t(\cdot)$ and K_t implicitly

depend on h, and are defined with respect to the $t \leq l_h$ points chosen up to time t in the for loop (as opposed to the $u_h \geq l_h$ points sampled *after* the for loop).

We first formalize the claim that the number of epochs is at most $\bar{H} = \log_2 T$.

Lemma 10. *For any time horizon T, Algorithm 1 terminates after at most $\log_2 T$ epochs.*

Proof. This follows immediately from the fact that we initialize $l_0 = 2$, double l_h after each epoch, and take at least l_h actions in epoch h (see Line 12 with $\sum_x \xi_h(x) = 1$) until T actions have been played. □

Next, we state a simple result regarding the epoch lengths.

Lemma 11. *The length u_h of epoch h in Algorithm 1 satisfies $u_h \leq l_h(2 + |\mathcal{S}_h|\psi)$.*

Proof. The number of times each action from $\mathcal{S}_h$ is played is $u_h(x)$, and is given in Algorithm 1 (Line 12). Hence, we have

$$u_h = \sum_{x \in \mathcal{S}_h} \lceil l_h \max\{\xi_h(x), \psi\} \rceil \tag{53}$$

$$\leq \sum_{x \in \mathcal{S}_h} (l_h \max\{\xi_h(x), \psi\} + 1) \tag{54}$$

$$\leq |\mathcal{S}_h| + \sum_{x \in \mathcal{S}_h} (l_h \xi_h(x) + l_h \psi) \tag{55}$$

$$\leq 2l_h + l_h \psi |\mathcal{S}_h| = l_h(2 + \psi|\mathcal{S}_h|), \tag{56}$$

where in the last inequality, we use $|\mathcal{S}_h| \leq l_h$ and $\sum_{x \in \mathcal{S}_h} \xi_h(x) = 1$. □

The following result characterizes the posterior uncertainty of points sampled in between the switching events in Algorithm 1, and may be of independent interest for problems in RKHS function spaces, particularly in settings where infrequent action switching is desirable.

Lemma 12. *Consider any epoch h, the corresponding set of actions $\mathcal{X}_h$, and the regularization parameter $\lambda > 0$. Let $t, t' \in [l_h]$ denote two rounds in epoch h such that $t \geq t'$, and for which*

$$\det(I_t + \lambda^{-1} K_t) \leq \eta \det(I_{t'} + \lambda^{-1} K_{t'}) \tag{57}$$

(i.e., the condition in Line 6 in Algorithm 1 does not hold), where $\eta > 1$. Then, for every $x \in \mathcal{X}_h$, it holds that

$$\sigma_{t'}(x) \leq \sqrt{\eta} \sigma_t(x). \tag{58}$$

Proof. We first consider the case that $k(x, x') = \phi(x)^T \phi(x')$ for every $x, x' \in \mathcal{X}$ with finite-dimensional features: $\phi(x) \in \mathbb{R}^{d_\phi}$ for some $d_\phi < \infty$. We let $\Phi_t = (\phi(x_{t'}))_{t' \leq t} \in \mathbb{R}^{t \times d_\phi}$ denote the matrix of features placed in t rows. We will later drop the assumption of finite dimensionality to obtain the result in our original setup.

We also note that if $\phi(x)$ contains all zeros for some input $x \in \mathcal{X}$, the statement in Equation (58) trivially holds (i.e., both sides are zero), so in the rest of the analysis, we assume that this is not the case.

In the following, let x be any fixed point in the domain. From Eq. (57), we have:

$$\eta \geq \frac{\det(\lambda^{-1}K_t + I_t)}{\det(\lambda^{-1}K_{t'} + I_{t'})} \tag{59}$$

$$= \frac{\det(K_t + \lambda I_t)}{\det(K_{t'} + \lambda I_{t'})} \tag{60}$$

$$= \frac{\det\left(\Phi_t^T\Phi_t + \lambda I_d\right)}{\det\left(\Phi_{t'}^T\Phi_{t'} + \lambda I_d\right)} \tag{61}$$

$$= \frac{\det\left(\left(\Phi_{t'}^T\Phi_{t'} + \lambda I_d\right)^{-1}\right)}{\det\left(\left(\Phi_t^T\Phi_t + \lambda I_d\right)^{-1}\right)} \tag{62}$$

$$\geq \frac{\phi(x)^T\left(\Phi_{t'}^T\Phi_{t'} + \lambda I_d\right)^{-1}\phi(x)}{\phi(x)^T\left(\Phi_t^T\Phi_t + \lambda I_d\right)^{-1}\phi(x)} \tag{63}$$

$$= \frac{\sigma_{t'}^2(x)}{\sigma_t^2(x)}. \tag{64}$$

Here, Eq. (61) holds due to the Weinstein–Aronszajn identity (i.e., $\det(I + AB) = \det(I + BA)$), and in Eq. (62) we use the fact that $\det(A) = (\det(A^{-1}))^{-1}$ for any invertible matrix A. Eq. (63) is proved in the following paragraph, and Eq. (64) follows from the alternative definition of $\sigma_t(\cdot)$ in Eq. (16).

It remains to prove the inequality in Eq. (63), which closely follows the proof of [2, Lemma 12]. For any $i \in [t]$, let $V_i := \lambda^{-1}\Phi_i^T\Phi_i + I$. We first show that

$$\frac{\phi(x)^T V_t \phi(x)}{\phi(x)^T V_{t-1}\phi(x)} \leq 1 + \|\lambda^{-1/2}\phi(x_t)\|^2_{V_{t-1}^{-1}}. \tag{65}$$

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

$$\phi(x)^T V_t \phi(x) = \phi(x)^T V_{t-1}\phi(x) + \phi(x)^T\left(\lambda^{-1}\phi(x_t)\phi(x_t)^T\right)\phi(x) \tag{66}$$

$$= \phi(x)^T V_{t-1}\phi(x) + \lambda^{-1}\left(\phi(x)^T\phi(x_t)\right)^2 \tag{67}$$

$$= \phi(x)^T V_{t-1}\phi(x) + \lambda^{-1}\left(\phi(x)^T V_{t-1}^{1/2} V_{t-1}^{-1/2}\phi(x_t)\right)^2 \tag{68}$$

$$\leq \phi(x)^T V_{t-1}\phi(x) + \lambda^{-1}\|\phi(x)^T V_{t-1}^{1/2}\|_2^2 \|V_{t-1}^{-1/2}\phi(x_t)\|_2^2 \tag{69}$$

$$= \phi(x)^T V_{t-1}\phi(x) + \lambda^{-1}(\phi(x)^T V_{t-1}\phi(x))(\phi(x_t)V_{t-1}^{-1}\phi(x_t)) \tag{70}$$

$$= \left(1 + \|\lambda^{-1/2}\phi(x_t)\|^2_{V_{t-1}^{-1}}\right)\phi(x)^T V_{t-1}\phi(x), \tag{71}$$

where Eq. (69) follows from Cauchy-Schwarz inequality. Hence, Eq. (65) follows by rearranging.

Since $t > t'$, we have:

$$\frac{\phi(x)^T V_t \phi(x)}{\phi(x)^T V_{t'}\phi(x)}$$

$$= \frac{\phi(x)^T V_t \phi(x)}{\phi(x)^T V_{t-1}\phi(x)} \cdot \frac{\phi(x)^T V_{t-1} \phi(x)}{\phi(x)^T V_{t-2}\phi(x)} \cdot \ldots \frac{\phi(x)^T V_{t'+1}\phi(x)}{\phi(x)^T V_{t'}\phi(x)} \tag{72}$$

$$\leq \left(1 + \|\lambda^{-1/2}\phi(x_t)\|^2_{V_{t-1}^{-1}}\right) \cdot \left(1 + \|\lambda^{-1/2}\phi(x_{t-1})\|^2_{V_{t-2}^{-1}}\right) \cdot \ldots \left(1 + \|\lambda^{-1/2}\phi(x_{t'+1})\|^2_{V_{t'}^{-1}}\right) \tag{73}$$

$$= \frac{\det(V_t)}{\det(V_{t-1})} \cdot \frac{\det(V_{t-1})}{\det(V_{t-2})} \cdot \ldots \frac{\det(V_{t'+1})}{\det(V_{t'})} \tag{74}$$

$$= \frac{\det(V_t)}{\det(V_{t'})}, \tag{75}$$

where Eq. (73) follows from Eq. (65), and Eq. (74) uses the fact that

$$\frac{\det(V_t)}{\det(V_{t-1})} = 1 + \|\lambda^{-1/2}\phi(x_t)\|^2_{V_{t-1}^{-1}}, \tag{76}$$

which is shown in the proof of [18, Theorem 2.2].

It remains to handle the possibly infinite feature dimension. Consider $k(x,x') = \sum_{i=1}^{\infty} \lambda_i \phi_i(x)\phi_i(x')$ and let $k_{d_\phi}(x,x') = \sum_{i=1}^{d_\phi} \lambda_i \phi_i(x)\phi_i(x')$ denote the finite dimensional kernel that corresponds to the d_ϕ-dimensional feature space such that $\lim_{d_\phi \to \infty} k_{d_\phi}(x,x') = k(x,x')$ for every $x, x' \in \mathcal{X}$. We use K_{t,d_ϕ} and $\sigma^2_{t,d_\phi}(\cdot)$ to denote the restriction of the corresponding quantities when the kernel $k_{d_\phi}(\cdot,\cdot)$ is used. First, we note that Eq. (60) still holds. Moreover, we have $\frac{\det(K_t+\lambda I_t)}{\det(K_{t'}+\lambda I_{t'})} = \lim_{d_\phi\to\infty} \frac{\det(K_{t,d_\phi}+\lambda I_t)}{\det(K_{t',d_\phi}+\lambda I_{t'})}$ and $\frac{\sigma^2_{t'}(x)}{\sigma^2_t(x)} = \lim_{d_\phi\to\infty} \frac{\sigma^2_{t',d_\phi}(x)}{\sigma^2_{t,d_\phi}(x)}$, and the former limit is lower bounded by the latter due to the fact that Eqs. (61) to (63) are all valid for the finite d_ϕ-feature approximation. Thus, the final result still holds for infinite dimensional kernels.

□

Next, we uniformly bound the posterior variance for the points remaining after a given epoch.

Lemma 13. *For any epoch h and the corresponding set of actions $\mathcal{X}_h$, it holds that*

$$\max_{x\in\mathcal{X}_h} \sigma^{(h)}(x) \le \sqrt{\frac{\eta(2\lambda+1)\gamma_{l_h}}{l_h}}. \tag{77}$$

Proof. Recall that u_h corresponds to the length of epoch h and that we can $\sigma^{(h)}(x)$ represents a posterior variance $\sigma_{u_h}(x)$ taken with respect to the u_h sampled points after the epoch. We first relate this to the posterior variance $\sigma_{l_h}(x)$ (abusing notation slightly) taken only with respect to the l_h points in the for loop in Algorithm 1. In particular, we claim that the former is upper bounded by the latter, and so it suffices to work with $\sigma_{l_h}(x)$. To see this, we recall that each x is sampled $u_h(x) = \lceil l_h \max\{\xi_h(x), \psi\}\rceil$ times, and the definition $\xi_h(x) = \frac{\sum_{i=1}^{l_h} \mathbb{1}\{x=x_i\}}{l_h}$ gives $l_h\xi_h(x) = \sum_{i=1}^{l_h} \mathbb{1}\{x = x_i\}$. Thus, the number of times each point is sampled is at least as high as the number of times it is selected in the for loop. Since conditioning on a higher number of points always decreases (or at least does not increase) the posterior variance in a Gaussian process, the desired claim follows.

We proceed to upper bound $\max_{x\in\mathcal{X}_h} \sigma_{l_h}(x)$. Let $\mathcal{T}_h = \{t \in [l_h] : \det(I_t + \lambda^{-1}K_t) > \eta \det(I_{t'} + \lambda^{-1}K_{t'})\}$ be the rounds in which the condition in Line 6 (Algorithm 1) is satisfied. Moreover, let $\bar{\mathcal{T}}_h = \mathcal{T}_h \cup \{0\}$ and let its elements $\bar{\mathcal{T}}_h = \{t'_0, \dots, t'_i, \dots, t'_{|\mathcal{T}_h|}\}$ be increasingly ordered. We note that $\max_{x\in\mathcal{X}_h} \sigma_{l_h}(x) \le \sigma_{t'_i}(x_{t'_i+1})$ for every $t'_i \in \bar{\mathcal{T}}_h$ according to the selection rule in Algorithm 1 (Line 4) and the fact that $\sigma_t(\cdot)$ is decreasing with respect to t. It follows that

$$l_h\big(\max_{x\in\mathcal{X}_h} \sigma_{l_h}(x)\big) \le \Big(\sum_{i=0}^{|\mathcal{T}_h|-1} (t'_{i+1} - t'_i)\sigma_{t'_i}(x_{t'_i+1})\Big) + (l_h - t'_{|\mathcal{T}_h|})\sigma_{t'_{|\mathcal{T}_h|}}(x_{t'_{|\mathcal{T}_h|}+1}). \tag{78}$$

Observe that by definition, we have $x_{t'_i+1} = x_{t'_i+2} = \cdots = x_{t'_{i+1}}$, i.e., these form a chain of identical points up to when the switching condition in Line 6 holds. Accordingly, by Lemma 12, it holds that $\sigma_{t'_i}(x_{t'_i+1}) \le \sqrt{\eta}\sigma_t(x_{t+1})$ for every $t \in \{t'_i, \dots, t'_{i+1} - 1\}$. By combining this with Eq. (78), we obtain

$$l_h\big(\max_{x\in\mathcal{X}_h} \sigma_{l_h}(x)\big) \le \sqrt{\eta} \sum_{t=0}^{l_h-1} \sigma_t(x_{t+1}). \tag{79}$$

Finally, from Lemma 9, we have $\sum_{t=0}^{l_h-1} \sigma_t(x_{t+1}) \le \sqrt{(2\lambda+1)\gamma_{l_h} l_h}$. By combining this with Equation (79) and rearranging, we obtain the final result. □

Finally, we provide a result bounding the size of the set $\mathcal{S}_h$ in Algorithm 1.

Lemma 14. *For any epoch h and the corresponding set $\mathcal{S}_h$, we have*

$$|\mathcal{S}_h| \le \tfrac{2}{\ln\eta}\gamma_T. \tag{80}$$

Proof. By the algorithm design, the set $\mathcal{S}_h$ grows by at most one element after the condition in Line 6 is satisfied, i.e., when

$$\det(I_t + \lambda^{-1} K_t) > \eta \det(I_{t'} + \lambda^{-1} K_{t'}), \tag{81}$$

where t is the current iteration, and t' is iteration prior to t for which Line 6 held (or $t' = 0$). As before, let $\mathcal{T}_h = \{t \in [l_h] : \det(I_t + \lambda^{-1} K_t) > \eta \det(I_{t'} + \lambda^{-1} K_{t'})\}$ be the rounds in which this holds, ordered with respect to time. Thus, for consecutive t_i and t_{i-1} belonging to $\mathcal{T}_h$, we have

$$\det(I_{t_i} + \lambda^{-1} K_{t_i}) > \eta \det(I_{t_{i-1}} + \lambda^{-1} K_{t_{i-1}}). \tag{82}$$

By applying the previous relation recursively, it follows that

$$\begin{aligned} \det(I_{t_i} + \lambda^{-1} K_{t_i}) &> \eta \det(I_{t_{i-1}} + \lambda^{-1} K_{t_{i-1}}) \\ &> \eta^2 \det(I_{t_{i-2}} + \lambda^{-1} K_{t_{i-2}}) \\ &> \dots \\ &> \eta^{i+1} \det(1 + \lambda^{-1}) = \eta^{i+1}(1 + \lambda^{-1}). \end{aligned} \tag{83}$$

Using the definition of γ_{l_h} given in (8), and noting that the size of the set $\mathcal{T}_h$ is at least $|\mathcal{S}_h| - 1$, we obtain

$$\gamma_{l_h} \geq \tfrac{1}{2} \ln \det(I_{l_h} + \lambda^{-1} K_{l_h}) \geq \tfrac{1}{2} \ln(\eta^{|\mathcal{S}_h|}(1 + \lambda^{-1})) \geq \tfrac{1}{2} \ln(\eta^{|\mathcal{S}_h|}). \tag{84}$$

By rearranging, we obtain

$$|\mathcal{S}_h| \leq \tfrac{2}{\ln \eta} \gamma_{l_h}. \tag{85}$$

The result then follows since $\gamma_T \geq \gamma_{l_h}$ for every h.

□

D Regret Analysis

In this appendix, we prove our main result, Theorem 3. We first upper bound the regret of any point sampled in a given epoch.

Lemma 15. *With probability at least $1 - \delta$, we have for every epoch h and $x \in \mathcal{X}_h$ that*

$$\max_{x \in \mathcal{X}_h} f(x) - f(x) \leq 4\Big(\beta_{h-1} + \tfrac{C\sqrt{u_{h-1}}}{l_{h-1}\psi\lambda}\Big)\sqrt{\tfrac{\eta(2\lambda+1)\gamma_{l_{h-1}}}{l_{h-1}}}. \tag{86}$$

Proof. Recall that u_h denotes the epoch length, and let $x_h^* \in \arg\max_{x \in \mathcal{X}_h} f(x)$. By using the validity of the confidence bounds from the end of the previous epoch $h - 1$ (see Lemma 2), we have for all $x \in \mathcal{X}_h$ that

$$\begin{aligned} f(x_h^*) - f(x) \leq \widetilde{\mu}^{(h-1)}(x_h^*) + \big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x^*) \\ - \widetilde{\mu}^{(h-1)}(x) + \big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x), \end{aligned} \tag{87}$$

where in Lemma 2 we substitute $h - 1$ and set $u_{\min} = l_{h-1}\psi$ (since each action selected in epoch $h - 1$ in Algorithm 1 is played at least $\lceil l_{h-1}\psi \rceil$ times), to upper and lower bound $\max_{x \in \mathcal{X}_h} f(x)$ and $f(x)$, respectively.

Next, for any $x \in \mathcal{X}_h$, it holds that

$$\begin{aligned} &\widetilde{\mu}^{(h-1)}(x) + \big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x) \\ &\quad\geq \max_{x \in \mathcal{X}_{h-1}} \Big(\widetilde{\mu}^{(h-1)}(x) - \big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x)\Big) \qquad (88) \\ &\quad\geq \widetilde{\mu}^{(h-1)}(x_h^*) - \big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x_h^*), \qquad (89) \end{aligned}$$

where Eq. (88) follows from the elimination condition (see Line 15 in Algorithm 1), and Eq. (89) holds since $x_h^* \in \mathcal{X}_h \subseteq \mathcal{X}_{h-1}$.

Combining Eq. (89) with Eq. (87), we obtain

$$f(x_h^*) - f(x) \leq 2\big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x_h^*) + 2\big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\sigma^{(h-1)}(x). \tag{90}$$

$$\leq 4\big(\beta_{h-1} + \tfrac{C}{l_{h-1}\psi\lambda}\sqrt{u_{h-1}}\big)\max_{x\in\mathcal{X}_{h-1}}\sigma^{(h-1)}(x). \tag{91}$$

The desired result then follows by upper bounding $\max_{x\in\mathcal{X}_{h-1}}\sigma^{(h-1)}(x)$ according to Lemma 13. □

We are ready to prove our main theorem, which is restated as follows.

Theorem 3 (Main result)**.** *Under the preceding setup and Assumption 1, for any corruption budget $C \geq 0$, Algorithm 1 with a constant switching parameter $\eta > 1$ and truncation parameter $\psi = \frac{\ln\eta}{2\gamma_T}$ satisfies the following with probability at least $1-\delta$:*

$$R_T = O^*\big(\beta_{\bar{H}}\sqrt{T\gamma_T} + C\gamma_T^{3/2}\big). \tag{12}$$

Proof. Throughout the proof, we condition on the confidence bounds from Lemma 2 holding true. We use $u_h(x)$ to denote the number of times action x is played in epoch h, and bound the cumulative regret of Algorithm 1 as follows:

$$R_T = \sum_{h=0}^{H-1}\sum_{x\in\mathcal{S}_h}\big(f(x^*) - f(x)\big)u_h(x) \tag{92}$$

$$\leq u_0 B + \sum_{h=1}^{H-1}\sum_{x\in\mathcal{S}_h}\big(f(x^*) - f(x)\big)u_h(x) \tag{93}$$

$$\leq u_0 B + \sum_{h=1}^{H-1}\sum_{x\in\mathcal{S}_h} u_h(x)\cdot 4\Big(\beta_{h-1} + \tfrac{C\sqrt{u_{h-1}}}{l_{h-1}\psi\lambda}\Big)\sqrt{\tfrac{\eta(2\lambda+1)\gamma_{l_{h-1}}}{l_{h-1}}}. \tag{94}$$

Here, Eq. (92) follows since only points from $\mathcal{S}_h$ are queried by the algorithm (and each point $x\in\mathcal{S}_h$ is queried $u_h(x)$ times), Eq. (93) follows since the bound on the RKHS norm implies the same bound on the maximal function value when the kernel $k(\cdot,\cdot)$ is normalized (namely, $k(x,x)\leq 1$ for every x):

$$|f(x)| = |\langle f, k(x,\cdot)\rangle_k| \leq \|f\|_k\|k(x,\cdot)\|_k = \|f\|_k\langle k(x,\cdot), k(x,\cdot)\rangle_k^{1/2} \leq B\cdot k(x,x)^{1/2} \leq B, \tag{95}$$

and Eq. (94) follows from Lemma 15 and by noting that $f(x^*) = \max_{x\in\mathcal{X}_h} f(x)$ for every $h = 0,1,\ldots,H-1$ (i.e., since the confidence bounds of Lemma 2 are valid, the global maximizer never gets eliminated). Next, from Eq. (94), by noting that $\sum_{x\in\mathcal{S}_h} u_h(x) = u_h$, we have:

$$R_T \leq u_0 B + \sum_{h=1}^{H-1} 4u_h\Big(\beta_{h-1} + \tfrac{C\sqrt{u_{h-1}}}{l_{h-1}\psi\lambda}\Big)\sqrt{\tfrac{\eta(2\lambda+1)\gamma_{l_{h-1}}}{l_{h-1}}} \tag{96}$$

$$\leq u_0 B + \sum_{h=1}^{H-1} 4l_h(2+\psi|\mathcal{S}_h|)\Big(\beta_{h-1} + \tfrac{C\sqrt{l_{h-1}(2+\psi|\mathcal{S}_{h-1}|)}}{l_{h-1}\psi\lambda}\Big)\sqrt{\tfrac{\eta(2\lambda+1)\gamma_{l_{h-1}}}{l_{h-1}}} \tag{97}$$

$$\leq u_0 B + \sum_{h=1}^{H-1} 4l_h(2+\psi|\mathcal{S}_h|)\Big(\beta_{\bar{H}} + \tfrac{C\sqrt{l_{h-1}(2+\psi|\mathcal{S}_{h-1}|)}}{l_{h-1}\psi\lambda}\Big)\sqrt{\tfrac{\eta(2\lambda+1)\gamma_T}{l_{h-1}}} \tag{98}$$

$$= u_0 B + \sum_{h=1}^{H-1} 8(2+\psi|\mathcal{S}_h|)\Big(\beta_{\bar{H}}\sqrt{\eta(2\lambda+1)l_{h-1}\gamma_T} + \tfrac{C\sqrt{(2+\psi|\mathcal{S}_{h-1}|)\eta(2\lambda+1)\gamma_T}}{\psi\lambda}\Big) \tag{99}$$

$$\leq u_0 B + \sum_{h=1}^{H-1} 8(2+\psi|\mathcal{S}_h|)\Big(\beta_{\bar{H}}\sqrt{\eta(2\lambda+1)T\gamma_T} + \tfrac{C\sqrt{(2+\psi|\mathcal{S}_{h-1}|)\eta(2\lambda+1)\gamma_T}}{\psi\lambda}\Big) \tag{100}$$

$$\leq u_0 B + 8\bar{H}(2+\tfrac{2\psi}{\ln\eta}\gamma_T)\Big(\beta_{\bar{H}}\sqrt{\eta(2\lambda+1)T\gamma_T} + \tfrac{C\sqrt{(2+\frac{2\psi}{\ln\eta}\gamma_T)\eta(2\lambda+1)\gamma_T}}{\psi\lambda}\Big), \tag{101}$$

where Eq. (97) follows from the bound on u_h in Lemma 11, Eq. (98) from the monotonicity of β_h in $h \in \{1, \dots, \bar{H}\}$ and $\gamma_t \in \{1, \dots, T\}$ in t (see Lemma 10 for the statement that $h \leq \bar{H}$), Eq. (99) by rearranging and using $l_h = 2l_{h-1}$, Eq. (100) by upper bounding l_{h-1} by T, and Eq. (101) from the bound on $|\mathcal{S}_h|$ in Lemma 14.

By setting, $\psi = \frac{\ln \eta}{2\gamma_T}$ as in the theorem statement, it follows that

$$R_T \leq u_0 B + 24\bar{H}\Big(\beta_{\bar{H}}\sqrt{\eta(2\lambda+1)T\gamma_T} + C\sqrt{\tfrac{12\eta(2\lambda+1)\gamma_T^3}{\lambda^2(\ln\eta)^2}}\Big). \tag{102}$$

Treating $\lambda > 0$ as a constant, it suffices to set the switching parameter η to some constant value (above one), so we choose $\eta = e$ (Euler's number). Then, we note that $u_0 = O(1)$ by design in the algorithm (recall that $l_0 = 2$, and note that $\psi \leq 1$ except possibly when T is small), and we write our regret bound as

$$R_T \leq O\big(\bar{H}(\beta_{\bar{H}}\sqrt{T\gamma_T} + C\gamma_T^{3/2})\big). \tag{103}$$

By using the notation $O^*(\cdot)$ to hide the multiplicative $\bar{H} = \log_2 T$ factor, the final result then follows:

$$R_T \leq O^*\big(\beta_{\bar{H}}\sqrt{T\gamma_T} + C\gamma_T^{3/2}\big). \tag{104}$$

□

E Alternative Approach: Reduction to Linear Bandits

In this section, we introduce an alternative method for corrupted kernelized bandit optimization, and discuss its limitations. We reduce the kernelized bandit problem of dimension d to a linear bandit problem of dimension D[3] using techniques from [44], and then solve the corrupted linear bandit problem using a modified version of the Robust Phased Elimination algorithm [9].

We consider a finite set of D actions $\mathcal{X}_D = \{s_1, \dots, s_D\} \subseteq \mathcal{X}$,and denote by $V(\mathcal{X}_D)$ the vector subspace of $\mathcal{H}_k$ spanned by $\{k(\cdot, s_i) : s_i \in \mathcal{X}_D\}$. Following [44], we consider using the orthogonal projection $\Pi_D(f)$ of f onto $V(\mathcal{X}_D)$ as an approximation of f, where $\Pi_D(f)$ is also the unique interpolant of f on $\mathcal{X}_D$ in $V(\mathcal{X}_D)$, i.e., $\Pi_D(f)(s_i) = f(s_i)$ for $i = 1, \dots, D$. To design this set $\mathcal{X}_D$, we use Algorithm 2 (taken from [44]), which takes the kernel k, domain $\mathcal{X}$, and an admissible error e as input, and outputs $\mathcal{X}_D$ along with the Newton basis of $V(\mathcal{X}_D)$. Recalling that $\|f\|_k \leq B$, we run Algorithm 2 with admissible error $e = \Delta/B$ for some constant $\Delta > 0$. We will discuss the choice of Δ later.

Algorithm 2 Newton Basis Construction [44]

Input: Kernel k, domain $\mathcal{X}$, admissible error e
Output: $\mathcal{X}_D = \{s_1, \dots, s_D\} \subseteq \mathcal{X}$, Newton basis $N_1, \dots, N_D$ of $V(\mathcal{X}_D)$

$s_1 \leftarrow \arg\max_{x \in \mathcal{X}} k(x, x)$
$N_1(x) \leftarrow k(x, s_1)/\sqrt{k(s_1, s_1)}$
for $D \leftarrow 1, 2, \dots$ **do**
  Define $P_D^2(x) = k(x, x) - \sum_{i=1}^{D} N_i^2(x)$
  if $\max_{x \in \mathcal{X}} P_D^2(x) < e^2$ **then**
    return $\{s_1, \dots, s_D\}$ and $\{N_1, \dots, N_D\}$
  end if
  $s_{D+1} \leftarrow \arg\max_{x \in \mathcal{X}} P_D^2(x)$
  $u(x) \leftarrow k(x, s_{D+1}) - \sum_{i=1}^{D} N_i(s_{D+1}) N_i(x)$
  $N_{D+1}(x) \leftarrow u(x)/\sqrt{P_D^2(s_{D+1})}$
end for

By rearranging the equations in Theorem 6 of [44], we have that the number of points returned by the algorithm is $D = O\big((\log \frac{1}{\Delta})^d\big)$ for kernels with infinite smoothness (in particular, the SE kernel), and $D = O(\Delta^{-d/\nu})$ for kernels with finite smoothness ν (in particular, the Matérn-ν kernel).

[3] The notation D for the continuous domain $[0, 1]^d$ will not be used in this appendix, so it it safe to use D for this dimension quantity.

Since the Newton basis $\{N_1, \ldots, N_D\}$ returned is the Gram-Schmidt orthonormalization of the basis $\{k(\cdot, s_i) : s_i \in \mathcal{X}_D\}$, we have for any $f \in \mathcal{H}_k$ and $x \in \mathcal{X}$ that

$$\left| f(x) - \sum_{i=1}^{D} \langle f, N_i \rangle N_i(x) \right| \leq \|f\|_k \cdot e \leq \Delta \tag{105}$$

under the choice $e = \Delta/B$. Hence, for any fixed black-box $f \in \mathcal{H}_k$ with $\|f\|_k \leq B$, there exists a $\theta \in \mathbb{R}^D$ with $\|\theta\|_2 \leq B$ such that for any $x \in \mathcal{X}$,

$$|f(x) - \langle \theta, \widetilde{x} \rangle| \leq \Delta, \tag{106}$$

where for any given point x, we define $\widetilde{x} = [N_1(x), \ldots, N_D(x)]^T$. Now, we can reduce the corrupted kernelized bandit problem to a variant of the corrupted linear bandit problem [9] on the transformed domain $\widetilde{\mathcal{X}} = \{\widetilde{x} : x \in \mathcal{X}\}$ of dimension D, where $|\widetilde{y}_t - \langle \theta, \widetilde{x}_t \rangle - c_t - \epsilon_t| \leq \Delta$ for $t = 1, \ldots, T$.

E.1 A Variant of Robust Phased Elimination

We apply Algorithm 3, a variant of the Robust Phased Elimination algorithm for stochastic linear bandits [9], on the space $\widetilde{\mathcal{X}}$ of dimension D, where the only difference from the original algorithm is the confidence bound in the elimination rule.

Algorithm 3 Robust Phased Elimination

Input: Actions $\widetilde{\mathcal{X}} \subseteq \mathbb{R}^D$, kernel k, admissible error e, confidence $\delta \in (0, 1)$, truncation parameter $\alpha \in (0, 1)$, time horizon T

1: $h \leftarrow 0, m_0 \leftarrow 4D(\log\log D + 18), \mathcal{A}_0 \leftarrow \widetilde{\mathcal{X}}$.

2: Compute design $\zeta_h : \mathcal{A}_h \to [0, 1]$ such that

$$\max_{\widetilde{x} \in \mathcal{A}_h} ||\widetilde{x}||^2_{\Gamma(\zeta_h)^{-1}} \leq 2D, \text{ and } |\text{supp}(\zeta_h)| \leq m_0, \tag{107}$$

where $\Gamma(\zeta_h) = \sum_{\widetilde{x} \in \mathcal{A}_h} \zeta_h(\widetilde{x}) \widetilde{x}\widetilde{x}^T$ (e.g., using Frank-Wolfe [30])

3: $u_h(\widetilde{x}) \leftarrow 0$ if $\zeta(\widetilde{x}) = 0$, and $u_h(\widetilde{x}) \leftarrow \lceil m_h \max\{\zeta_h(\widetilde{x}), \alpha\} \rceil$ otherwise.

4: Take each action x such that $\widetilde{x} \in \mathcal{A}_h$ exactly $u_h(\widetilde{x})$ times, and get rewards $\{\widetilde{y}_t\}_{t=1}^{u_h}$, where $u_h = \sum_{\widetilde{x} \in \mathcal{A}_h} u_h(\widetilde{x})$.

5: Estimate the parameter vector $\widetilde{\theta}_h$:

$$\widetilde{\theta}_h = \Gamma_h^{-1} \sum_{t=1}^{u_h} \widetilde{x}_t u_h(\widetilde{x}_t)^{-1} \sum_{s \in \mathcal{T}(\widetilde{x}_t)} \widetilde{y}_s, \tag{108}$$

where $\Gamma_h^{-1} = \sum_{\widetilde{x} \in \mathcal{A}_h} u_h(\widetilde{x}) \widetilde{x}\widetilde{x}^T$ and $\mathcal{T}(\widetilde{x}) = \{s \in \{1, \ldots, u_h\} : \widetilde{x}_s = \widetilde{x}\}$.

6: Update the active set of actions:

$$\mathcal{A}_{h+1} \leftarrow \Bigg\{ \widetilde{x} \in \mathcal{A}_h : \max_{\widetilde{x}' \in \mathcal{A}_h} \langle \widetilde{\theta}_h, \widetilde{x}' - \widetilde{x} \rangle \leq 4\Delta\sqrt{D(1 + \alpha m_0)} + 4\sqrt{\frac{D}{m_h} \log\frac{1}{\delta}} + \frac{4C}{\alpha m_h}\sqrt{D(1 + \alpha m_0)} \Bigg\}. \tag{109}$$

7: $m_{h+1} \leftarrow 2m_h, h \leftarrow h + 1$ and return to step 3 (terminating after T actions are played).

The analysis of Algorithm 3 is very similar to that of [9], so we heavily rely on their auxiliary results and only focus on explaining the differences here. With $\widetilde{\theta}_h$ denoting the estimate of θ based on the corrupted observations $\{\widetilde{y}_t\}_{t=1}^{u_h}$ in the algorithm, and $\widehat{\theta}_h$ denoting the estimate of θ based on $\{\langle \theta, \widetilde{x}_t \rangle + c_t + \epsilon_t\}_{t=1}^{u_h}$ (i.e., the corrupted observations if the linear model were exact) in the original algorithm, we have for all $h \geq 0$ and $\widetilde{x} \in \mathcal{A}_h$ that

$$|\langle \widetilde{x}, \widetilde{\theta}_h - \widehat{\theta}_h \rangle| \leq \left| \widetilde{x}^T \Gamma_h^{-1} \sum_{t=1}^{u_h} \widetilde{x}_t \Delta \right| \leq \Delta \sum_{t=1}^{u_h} \left| \langle \widetilde{x}, \Gamma_h^{-1} \widetilde{x}_t \rangle \right| \overset{(a)}{\leq} \Delta\sqrt{u_h} ||\widetilde{x}||_{\Gamma_h^{-1}} \overset{(b)}{\leq} 2\Delta\sqrt{D(1 + \alpha m_0)}, \tag{110}$$

where (a) uses the definition of $\|\cdot\|_{\Gamma_h^{-1}}$ and the fact that the ℓ_1-norm is upper bounded by the ℓ_2-norm times the square root of the vector length, and (b) uses Lemmas 2 and 3 of [9]. Hence, in a fixed epoch h, we have for all $\widetilde{x} \in \mathcal{A}_h$ that

$$|\langle \widetilde{x}, \widetilde{\theta}_h - \theta \rangle| \leq |\langle \widetilde{x}, \widetilde{\theta}_h - \widehat{\theta}_h \rangle| + |\langle \widetilde{x}, \widehat{\theta}_h - \theta \rangle| \tag{111}$$

$$\leq 2\Delta\sqrt{D(1+\alpha m_0)} + 2\sqrt{\frac{D}{m_h}\log\frac{1}{\delta}} + \frac{2C}{\alpha m_h}\sqrt{D(1+\alpha m_0)}, \tag{112}$$

where the first term uses (110), and the remaining terms are obtained with probability at least $1 - 2|\mathcal{X}|\delta$ by Lemma 4 in [9].

Defining $\bar{x} = \arg\max_{\bar{x}\in\widetilde{\mathcal{X}}} \langle \theta, \bar{x} \rangle$, by a similar analysis to Section A.2 in [9], we can show that the elimination rule in (109) retains $\bar{x}$ in a given epoch with probability at least $1 - 2|\mathcal{X}|\delta$. Recalling that $x^* = \arg\max_{x\in\mathcal{X}} f(x)$, we have

$$f(x^*) = \langle \theta, \widetilde{x}^* \rangle + f(x^*) - \langle \theta, \widetilde{x}^* \rangle \leq \langle \theta, \widetilde{x}^* \rangle + \Delta \leq \langle \theta, \bar{x} \rangle + \Delta. \tag{113}$$

Hence, the cumulative regret can be upper bounded as follows

$$R_T = \sum_{t=1}^{T} f(x^*) - f(x_t) \leq \sum_{t=1}^{T} (\langle \theta, \bar{x} \rangle + \Delta) - (\langle \theta, \widetilde{x}_t \rangle - \Delta) = \sum_{t=1}^{T} \langle \theta, \bar{x} - \widetilde{x}_t \rangle + 2\Delta T. \tag{114}$$

Again following the analysis of Section A.2 in [9], using (112) and (109), we can then show that the cumulative regret is

$$R_T = O^*\left(\Delta T\sqrt{D} + \sqrt{DT\log\frac{|\mathcal{X}|}{\delta}} + CD^{3/2}\right) \tag{115}$$

with probability at least $1 - \delta$.

E.2 The choice of Δ

The only remaining step now is to find a proper choice of Δ, which is what dictates the choice of D (along with the kernel). The choice of Δ can be optimized with respect to the kernel parameters, and the optimal scaling is achieved by equating the first terms in (115) with one of the other two terms (whichever is larger). We first consider the choice $\Delta = \frac{1}{\sqrt{T}}$, which equates the first two terms (up to the $\log\frac{|\mathcal{X}|}{\delta}$ factor).

With $\Delta = \frac{1}{\sqrt{T}}$, it is known from [44, Corollary 7] that Algorithm 2 results in $D = O\big((\log T)^d\big)$ for the SE kernel and $D = O(T^{\frac{d}{2\nu}})$ for the Matérn kernel. Hence, the cumulative regret of our method is upper bounded as follows:

- For the SE kernel,

$$R_T = O^*\Big(\sqrt{T(\log T)^d \log\frac{|\mathcal{X}|}{\delta}} + C(\log T)^{\frac{3d}{2}}\Big). \tag{116}$$

- For the Matérn kernel,

$$R_T = O^*\Big(\sqrt{T^{\frac{d+2\nu}{2\nu}}\log\frac{|\mathcal{X}|}{\delta}} + CT^{\frac{3d}{4\nu}}\Big). \tag{117}$$

For the Matérn kernel, we can sometimes do better by equating the first and third terms in (115), whereas for the SE kernel this is never the case. The exact optimal choice depends on how C scales with respect to T, but to avoid unwieldy expressions, we focus here on the direct T dependence in (115) so treat C as a constant. Equating the first and third terms, and ignoring the $\log T$ term, we find that we should set $\Delta = 1/T^{\frac{\nu}{d+\nu}}$, which yields $D = O(T^{\frac{d}{d+\nu}})$ [44], and gives

$$R_T = O^*\Big(\sqrt{T^{\frac{2d+\nu}{d+\nu}}\log\frac{|\mathcal{X}|}{\delta}} + CT^{\frac{3d}{2(d+\nu)}}\Big). \tag{118}$$

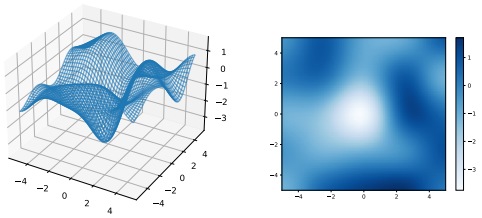

Figure 4: Illustration of 2D synthetic function.

We compare (117) and (118) for various (ν, d) pairs below.

For the SE kernel, the bound (116) turns out to be strong, matching our main result (Section 3), though we believe that our algorithm's feature of directly using the GP model (i.e., avoiding linear approximations) is still desirable.

For the Matérn kernel, however, the resulting bound is not as strong; in particular, the non-corrupted terms in both (117) and (118) are larger than the corresponding term $\sqrt{T\gamma_T} = O^*(T^{\frac{\nu+d}{2\nu+d}})$ in our main result.[4] The same goes for the corrupted terms, with the root cause for both terms being that either choice of D above is strictly higher than γ_T. For the corrupted term, this is further highlighted by comparing the regimes in which the bound remains sublinear:

- The term $T^{\frac{3d}{4\nu}}$ in (117) is sublinear when $\nu > \frac{3}{4}d$;
- The term $T^{\frac{3d}{2(d+\nu)}}$ in (118) is sublinear when $\nu > \frac{d}{2}$;
- The analogous term $\gamma_T^{3/2} = O^*(T^{\frac{3d}{4\nu+2d}})$ in the main body is sublinear under the milder condition $\nu > d/4$.

Note that in general, we have for constant C that (117) is a better bound than (118) when $\nu > d$, (118) is better than (117) when $\nu \in \left(\frac{d}{2}, d\right)$, and both fail to be sublinear when $\nu \leq \frac{d}{2}$.

We note that a slight caveat to the preceding findings is that it is unclear whether the choice of D arising from [44] is the best possible, but we are not aware of any similar results that are better for our purposes.

F Supplementary Experimental Results

Recall that the synthetic function f_1 is shown in Figure 4. This section contains the experimental results on f_1 with $C = 100$ (Figure 5), and on Robot3D with $C = 50$ (Figure 6). The overall findings are generally similar to those in the main text, and are not repeated here.

[4] For (117), this is seen by writing $T^{\frac{d+2\nu}{2\nu}} = T^{1+\frac{d}{2\nu}}$ and noting that $\frac{d}{2\nu}$ exceeds $\gamma_T = O^*(T^{\frac{d}{2\nu+d}})$. For (118), it is seen by writing $\sqrt{T^{\frac{2d+\nu}{d+\nu}}} = T^{\frac{2d+\nu}{2d+2\nu}}$, and noting that subtracting d from both the numerator and denominator makes the fraction smaller.

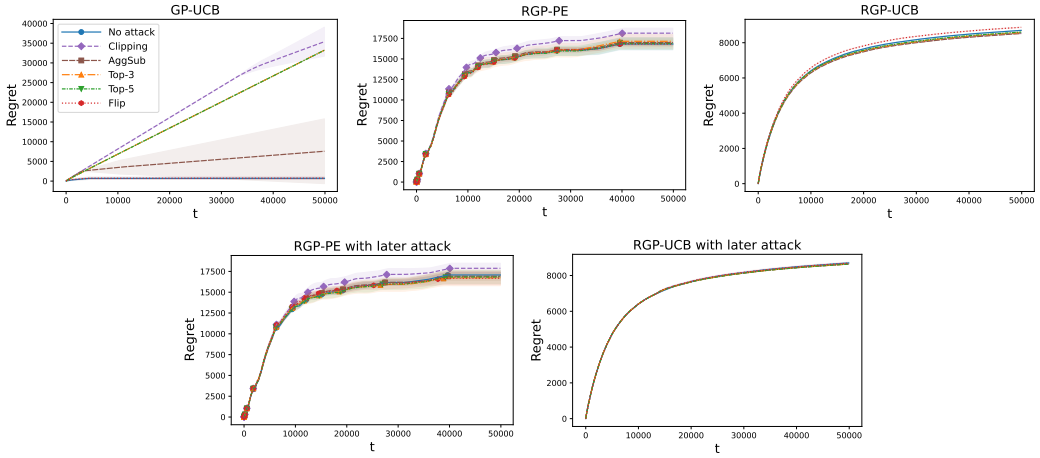

Figure 5: Performance on f_1 with $C = 100$. Note that for GP-UCB, the curves for Top-3 and Top-5 are indistinguishable, so only the latter is clearly visible. Similar trends are observed to the case $C = 50$ in Figure 2.

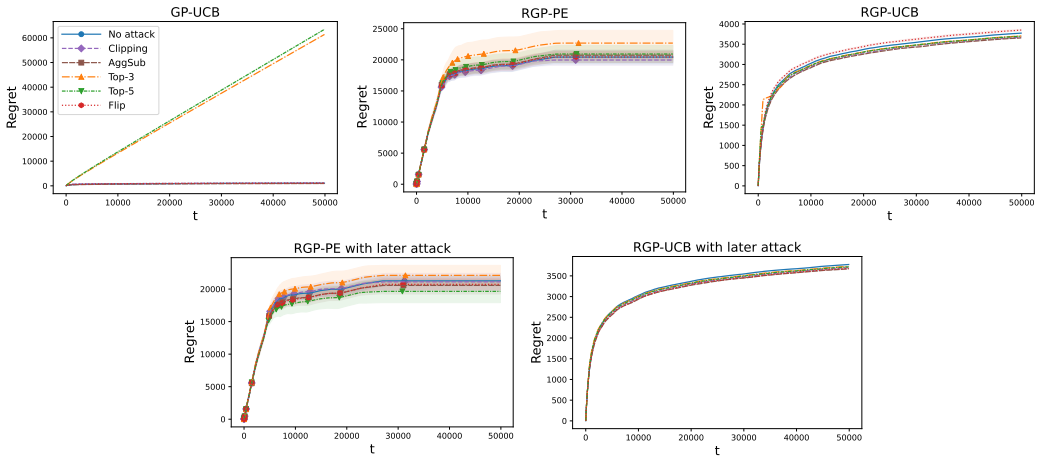

Figure 6: Performance on Robot3D with $C = 50$. Similar trends are observed to the case $C = 100$ in Figure 3.