
Tight Regret Bounds for Bayesian Optimization in One Dimension

Jonathan Scarlett¹

Abstract

We consider the problem of Bayesian optimization (BO) in one dimension, under a Gaussian process prior and Gaussian sampling noise. We provide a theoretical analysis showing that, under fairly mild technical assumptions on the kernel, the best possible cumulative regret up to time T behaves as $\Omega(\sqrt{T})$ and $O(\sqrt{T \log T})$. This gives a tight characterization up to a $\sqrt{\log T}$ factor, and includes the first non-trivial lower bound for noisy BO. Our assumptions are satisfied, for example, by the squared exponential and Matérn- ν kernels, with the latter requiring $\nu > 2$. Our results certify the near-optimality of existing bounds (Srinivas *et al.*, 2009) for the SE kernel, while proving them to be strictly suboptimal for the Matérn kernel with $\nu > 2$.

1. Introduction

Bayesian optimization (BO) (Shahriari *et al.*, 2016) is a powerful and versatile tool for black-box function optimization, with applications including parameter tuning, robotics, molecular design, sensor networks, and more. The idea is to model the unknown function as a Gaussian process with a given *kernel function* dictating the smoothness properties. This model is updated using (typically noisy) samples, which are selected to steer towards the function maximum.

One of the most attractive properties of BO is its efficiency in terms of the number of function samples used. Consequently, algorithms with *rigorous guarantees* on the trade-off between samples and optimization performance are particularly valuable. Perhaps the most prominent work in the literature giving such guarantees is that of (Srinivas *et al.*,

2010), who consider the *cumulative regret*:

$$R_T = \sum_{t=1}^T \left(\max_x f(x) - f(x_t) \right), \quad (1)$$

where f is the function being optimized, and x_t is the point chosen at time t . Under a Gaussian process (GP) prior and Gaussian noise, it is shown in (Srinivas *et al.*, 2010) that an algorithm called Gaussian Process Upper Confidence Bound (GP-UCB) achieves a cumulative regret of the form

$$R_T = O^*(\sqrt{T \gamma_T}), \quad (2)$$

where $\gamma_T = \max_{x_1, \dots, x_T} I(\mathbf{f}; \mathbf{y})$ (with function values $\mathbf{f} = (f(x_1), \dots, f(x_T))$ and noisy samples $\mathbf{y} = (y_1, \dots, y_T)$) is known as the *maximum information gain*. Here $I(\mathbf{f}; \mathbf{y})$ denotes the mutual information (Cover & Thomas, 2001) between the function values and noisy samples, and $O^*(\cdot)$ denotes asymptotic notation up to logarithmic factors.

The guarantee (2) ensures sub-linear cumulative regret for many kernels of interest. However, the literature is severely lacking in *algorithm-independent lower bounds*, and without these, it is impossible to know to what extent the upper bounds, including (2), can be improved. In this work, we address this gap in detail in the special case of a one-dimensional function. We show that the best possible cumulative regret behaves as $\Theta^*(\sqrt{T})$ under mild assumptions on the kernel, thus identifying both cases where (2) is near-optimal, and cases where it is strictly suboptimal.

1.1. Related Work

An extensive range of BO algorithms have been proposed in the literature, typically involving the maximization of an acquisition function (Hennig & Schuler, 2012; Hernández-Lobato *et al.*, 2014; Russo & Van Roy, 2014; Wang *et al.*, 2016); see (Shahriari *et al.*, 2016) for a recent overview. As mentioned above, the most relevant algorithm to this work for the noisy setting is GP-UCB (Srinivas *et al.*, 2010), which constructs *confidence bounds* in which the function lies with high probability, and samples the point with the highest upper confidence bound. Several extensions to GP-UCB have also been proposed, including contextual (Krause & Ong, 2011; Bogunovic *et al.*, 2016a), batch (Contal *et al.*, 2013; Desautels *et al.*, 2014), and high-dimensional (Kandasamy *et al.*, 2015; Rolland *et al.*, 2018) variants.

¹Department of Computer Science & Department of Mathematics, National University of Singapore. Correspondence to: Jonathan Scarlett <scarlett@comp.nus.edu.sg>.

In the noiseless setting, it has been shown that it is possible to achieve *bounded* cumulative regret (de Freitas et al., 2012; Kawaguchi et al., 2015) under some technical assumptions. In (de Freitas et al., 2012), this is done by keeping track of a set of potential maximizers, and sampling increasingly finely in order to shrink that set and “zoom in” towards the optimal point. Similar ideas have also been used in the noisy setting for studying batch variants of GP-UCB (Contal et al., 2013), simultaneous online optimization (SOO) methods (Wang et al., 2014), and lookahead algorithms that use confidence bounds (Bogunovic et al., 2016b). Returning to the noiseless setting, upper and lower bounds were given in (Grünwälder et al., 2010) for kernels satisfying certain smoothness assumptions, with the lower bounds showing that bounded cumulative regret is not always to be expected.

Alongside the Bayesian view of the Gaussian process model, several works have also considered a *non-Bayesian* counterpart assuming that the function has a bounded norm in the associated reproducing kernel Hilbert space (RKHS). Interestingly, GP-UCB still provides similar guarantees to (2) in this setting (Srinivas et al., 2010). Moreover, lower bounds have been proved; see (Bull, 2011) for the noiseless setting, and (Scarlett et al., 2017) for the noisy setting. In the latter, the lower bounds nearly match the GP-UCB upper bound for the squared exponential (SE) kernel, but gaps remain for the Matérn kernel. For reference, we note that these kernels are defined as follows:

$$k_{\text{SE}}(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2l^2}\right) \quad (3)$$

$$k_{\text{Matérn}}(x, x') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}\|x - x'\|}{l}\right)^\nu \times B_\nu\left(\frac{\sqrt{2\nu}\|x - x'\|}{l}\right), \quad (4)$$

where $l > 0$ is a lengthscale parameter, $\nu > 0$ is a smoothness parameter, B_ν is the modified Bessel function, and Γ is the gamma function.

The multi-armed bandit (MAB) (Bubeck & Cesa-Bianchi, 2012) literature has developed alongside the BO literature, with the two often bearing similar concepts. The MAB literature is far too extensive to cover here, but it is worth mentioning that sharp lower bounds are known in numerous settings (Bubeck & Cesa-Bianchi, 2012), and the above-mentioned concept of “zooming in” to the optimal point has also been explored (Kleinberg et al., 2008). To our knowledge, however, none of the existing MAB results are closely related to our own.

1.2. Our Results and Their Implications

The main results of this paper are informally summarized as follows.

Main Results (Informal). *Under mild technical assumptions on the kernel, satisfied (for example) by the SE kernel and Matérn- ν kernel with $\nu > 2$, the best possible cumulative regret of noisy BO in one dimension behaves as $\Omega(\sqrt{T})$ and $O(\sqrt{T \log T})$.*

Our results have several important implications:

- To our knowledge, our lower bound is the first of any kind in the noisy Bayesian setting, and is tight up to a $\sqrt{\log T}$ factor under our technical assumptions.
- Our lower bound also establishes the order-optimality of the $O^*(\sqrt{T})$ upper bound of (Srinivas et al., 2010) applied to the SE kernel, up to logarithmic factors.
- On the other hand, our upper bound establishes that the upper bound of (Srinivas et al., 2010) for the Matérn- ν kernel, namely $O^*(T^{\frac{\nu+2}{2\nu+2}})$, is strictly suboptimal for $\nu > 2$. For example, if $\nu = 3$, then this is $O^*(T^{0.625})$, as opposed to our upper bound of $O^*(T^{0.5})$. (See also (Shekhar & Javidi, 2017) for recent improvements over (Srinivas et al., 2010) under the Matérn kernel in higher dimensions and/or with smaller ν).
- Another important implication for the Matérn kernel with $\nu > 2$ is that the Bayesian setting is provably harder than the non-Bayesian RKHS counterpart; the latter has cumulative regret $\Omega(T^{\frac{\nu+1}{2\nu+1}})$ (Scarlett et al., 2017), which is strictly worse than $O(\sqrt{T \log T})$.

Our upper bound is stated formally in Section 3, and its technical assumptions are given in Section 2.1. We build on the ideas of (de Freitas et al., 2012) for the noiseless setting, while addressing highly non-trivial challenges arising in the presence of noise.

Our lower bound is stated formally in Section 4, and its technical assumptions are given in Section 2.1. The analysis is based on a reduction to binary hypothesis testing and an application of Fano’s inequality (Cover & Thomas, 2001). This approach is inspired by previous work on lower bounds for stochastic convex optimization (Raginsky & Rakhlin, 2011), but the details are very different.

2. Problem Setup

2.1. Bayesian Optimization

We seek to sequentially optimize an unknown reward function $f(x)$ over the one-dimensional domain $D = [0, 1]$; note that any interval can be transformed to this choice via re-scaling. At time t , we query a single point $x_t \in D$ and observe a noisy sample $y_t = f(x_t) + z_t$, where $z_t \sim N(0, \sigma^2)$ for some noise variance $\sigma^2 > 0$, with independence across different times. We measure the performance using the cumulative regret R_T , defined in (1).

We henceforth assume f to be distributed according to Gaussian process (GP) (Rasmussen, 2006) having mean zero and kernel function $k(x, x')$. The posterior distribution of f given the points $\mathbf{x}_t = [x_1, \dots, x_t]^T$ and observations $\mathbf{y}_t = [y_1, \dots, y_t]^T$ up to time t is again a GP, with the posterior mean and variance given by (Rasmussen, 2006)

$$\mu_t(x) = \mathbf{k}_t(x)^T (\mathbf{K}_t + \sigma^2 \mathbf{I}_t)^{-1} \mathbf{y}_t \quad (5)$$

$$\sigma_t(x)^2 = k(x, x) - \mathbf{k}_t(x)^T (\mathbf{K}_t + \sigma^2 \mathbf{I}_t)^{-1} \mathbf{k}_t(x), \quad (6)$$

where $\mathbf{k}_t(x) = [k(x_i, x)]_{i=1}^t$, $\mathbf{K}_t = [k(x_i, x_j)]_{i,j}$, and \mathbf{I}_t is the $t \times t$ identity matrix.

2.2. Technical Assumptions

Here we introduce several assumptions that will be adopted in our main results, some of which were also used in the noiseless setting (de Freitas et al., 2012).

Assumption 1. We have the following:

1. The kernel k is stationary, depending on its inputs (x, x') only through $\tau = x - x'$;
2. The kernel k satisfies $k(x, x') \leq 1$ for all (x, x') , and $k(x, x) = 1$ for all $x \in D$;

Given the stationarity assumption, the assumptions $k(x, x') \leq 1$ and $k(x, x) = 1$ are without loss of generality, as one can always re-scale the function and adjust the noise variance σ^2 accordingly.

Next, we give some high-probability assumptions on the random function f itself.

Assumption 2. There exists a constant $\delta_1 \in (0, 1)$ such that, with probability at least $1 - \delta_1$, we have the following:

1. The function f has a unique maximizer x^* such that

$$f(x^*) \geq f(x') + \epsilon \quad (7)$$

for any *local maximum* x' that differs from x^* , for some constant $\epsilon > 0$.

2. The function f is twice differentiable;
3. The function f and its first two derivatives are bounded:

$$|f(x)| \leq c_0, \quad |f'(x)| \leq c_1, \quad |f''(x)| \leq c_2 \quad (8)$$

for all $x \in D$ and some constants (c_0, c_1, c_2) . This implies that f is c_1 -Lipschitz continuous, and f' is c_2 -Lipschitz continuous.

The assumption of a unique maximizer holds with probability one in most non-trivial cases (de Freitas et al., 2012), and

(7) simply formally defines the gap to the second-highest peak. Moreover, given twice differentiability, the remaining conditions in (8) are very mild, only requiring that the function value and its derivatives are bounded, and formally defining the corresponding constants.

Next, we provide assumptions regarding the derivatives of f and the resulting Taylor expansions (typically around the optimizer x^*). We adopt slightly different assumptions for the upper and lower bounds, starting with the former.

Assumption 3. There exist constants $\delta_2 \in (0, 1)$ and $\rho_0 \in (0, \frac{1}{2})$ such that conditioned on the events in Assumption 2, we have with probability at least $1 - \delta_2$ that *one of the following* is true:

1. The maximizer is at an endpoint (i.e., $x^* = 0$ or $x^* = 1$), and satisfies the following *locally linear* behavior: For all $\xi \in [0, \rho_0]$ (if $x^* = 0$) or $\xi \in [-\rho_0, 0]$ (if $x^* = 1$), it holds that

$$f(x^*) - \underline{c}_1 |\xi| \geq f(x^* + \xi) \geq f(x^*) - \bar{c}_1 |\xi| \quad (9)$$

for some constants $\bar{c}_1 \geq \underline{c}_1 > 0$.

2. The maximizer satisfies $x^* \in (\rho_0, 1 - \rho_0)$, and f satisfies the following *locally quadratic* behavior: For all $\xi \in [-\rho_0, \rho_0]$, we have

$$f(x^*) - \underline{c}_2 \xi^2 \geq f(x^* + \xi) \geq f(x^*) - \bar{c}_2 \xi^2 \quad (10)$$

for some constants $\bar{c}_2 \geq \underline{c}_2 > 0$.

This assumption is near-identical to the main assumption adopted in the noiseless setting (de Freitas et al., 2012), and is also mild given the assumption of twice differentiability. Indeed, (9) and (10) amount to standard Taylor expansions, with the assumptions $\underline{c}_1 > 0$ and $\underline{c}_2 > 0$ only requiring non-vanishing gradient at the endpoint (first case) or non-vanishing second derivative at the function maximizer (second case). These conditions typically hold with probability one (de Freitas et al., 2012).

The following assumption will be used for the lower bound.

Assumption 4. There exists constants $\delta'_2 \in (0, 1)$ and $\rho_0 \in (0, \frac{1}{2})$ such that conditioned on the events in Assumption 2, *both of the following* hold with probability at least $1 - \delta'_2$:

1. For any $x \in D$ and $\xi \in [-\rho_0, \rho_0]$ for which $x + \xi \in D$, we have

$$\xi \cdot f'(x) + \underline{c}'_2 \xi^2 \leq f(x + \xi) - f(x) \leq \xi \cdot f'(x) + \bar{c}'_2 \xi^2 \quad (11)$$

for some (possibly negative) constants $\bar{c}'_2, \underline{c}'_2$.

2. The maximizer satisfies $x^* \in (\rho_0, 1 - \rho_0)$, and f satisfies the following for all $\xi \in [-\rho_0, \rho_0]$:

$$f(x^*) - \underline{c}_2 \xi^2 \geq f(x^* + \xi) \geq f(x^*) - \bar{c}_2 \xi^2 \quad (12)$$

for some constants $\bar{c}_2 \geq \underline{c}_2 > 0$.




Figure 1. Illustration of some of the main assumptions: The function is bounded within $[-c_0, c_0]$ and its derivative within $[-c_1, c_1]$, the gap to the second highest peak is at least ϵ , and the function is locally quadratic for points within a distance ρ_0 of the maximizer.

The first part is similar to (10), but performs a Taylor expansion around an arbitrary point rather than the specific point x^* , and the second part is precisely (10). Note, however, that here we are assuming *both* of two conditions to hold, rather than one of two. Hence, we are implicitly assuming that the first item of Assumption 3 does *not* have a significant probability of occurring. For stationary kernels, the only situations where an endpoint has a high probability of being optimal are those where f varies very slowly (e.g., the SE kernel with a larger lengthscale than the domain width). Such functions are of limited practical interest.

Similarly to the noiseless setting (de Freitas et al., 2012), all of the above assumptions hold for the SE kernel, as well as the Matérn- ν kernel with $\nu > 2$, with the added caveat that δ'_2 in Assumption 4 is a function of the lengthscale and cannot be chosen arbitrarily. Specifically, a smaller lengthscale implies a smaller value of δ'_2 . In contrast, δ_1 and δ_2 in Assumptions 2 and 3 can be made arbitrary small by suitably changing the constants ϵ , c_0 , c_1 , c_2 , ρ_0 , and so on.

An illustration of some of the main assumptions and their associated constants is given in Figure 1.

3. Upper Bound

Our upper bound is formally stated as follows.

Theorem 1. (Upper Bound) *Consider the problem of BO in one dimension described in Section 2.1, with time horizon T and noise variance σ^2 satisfying $\sigma^2 \geq \frac{c_\sigma}{T^{1-\zeta}}$ for some $c_\sigma > 0$ and $\zeta > 0$. Under Assumptions 1, 2, and 3, there exists an algorithm satisfying the following: With probability at least $1 - \delta_1 - \delta_2$ (with respect to the Gaussian process f), the average cumulative regret (averaged over the noisy*

samples) satisfies

$$\mathbb{E}[R_T] \leq C \left(1 + \sigma \sqrt{T \log T}\right). \quad (13)$$

Here δ_1 and δ_2 are defined in Assumptions 2 and 3, and C depends only on the constants therein and (c_σ, ζ) .

The assumption that $\sigma^2 \geq \frac{c_\sigma}{T^{1-\zeta}}$ for some (c_σ, ζ) is very mild, since typically σ^2 is constant with respect to T . The proof of Theorem 1 extends immediately to a high probability guarantee with respect to both f and the noisy samples (i.e., holding with probability $1 - \delta_1 - \delta_2 - \delta$ for δ in Lemma 1 below). We have stated the above form for consistency with the lower bound, which will be given in Section 4.

3.1. High-Level Description of the Algorithm

The algorithm considered in the proof of Theorem 1 is described informally in Algorithm 1; the details will be established throughout the proof of Theorem 1, and a complete description is given in Appendix B.

Algorithm 1 Informal description of our algorithm, based on reducing uncertainty in epochs via repeated sampling.

Require: Domain D , GP prior (μ_0, k_0) , discrete subdomain $\mathcal{L} \subseteq D$, time horizon T .

- 1: Initialize $t = 1$, epoch number $i = 1$, potential maximizers $M_{(0)} = \mathcal{L}$, and target confidence $\eta_{(0)}$.
- 2: **while** less than T samples have been taken **do**
- 3: Set $\eta_{(i)} = \frac{1}{2}\eta_{(i-1)}$.
- 4: Sample each point within a subset $\mathcal{L}_{(i)} \subseteq \mathcal{L}$ repeatedly $K_{(i)}$ times, where $\mathcal{L}_{(i)}$ and $K_{(i)}$ are chosen such that after this sampling, all points $x \in M_{(i-1)}$ satisfy upper and lower confidence bounds of the form

$$\text{LCB}_t(x) \leq f(x) \leq \text{UCB}_t(x),$$

with the gap between the two bounded by $|\text{UCB}_t(x) - \text{LCB}_t(x)| \leq 2\eta_{(i)}$.

- 5: Update the set of potential maximizers:

$$M_{(i)} = \left\{ x \in M_{(i-1)} : \text{UCB}_t(x) \geq \max_{x' \in M_{(i-1)}} \text{LCB}_t(x') \right\}.$$

- 6: Increment i .
- 7: **end while**

As in the noiseless setting (de Freitas et al., 2012), the idea is to operate in *epochs* and sample a set of increasingly closely-packed points $\mathcal{L}_{(i)}$ to reduce the posterior variance, but only within a set of *potential maximizers* that are updated according to the confidence bounds. As a simple means of bringing the effective noise level down, we perform *resampling*, i.e.,

sampling the same point $K_{(i)}$ times consecutively. In each epoch, we sample enough to be able to produce upper and lower confidence bounds $\text{UCB}_t(x)$ and $\text{LCB}_t(x)$ that differ by at most a target value $2\eta_{(i)}$ within $M_{(i-1)}$, and then the target is halved for the next epoch.

We do not expect our algorithm to perform well in practice by any means, but it still suffices for our purposes in establishing $O(\sqrt{T \log T})$ regret. Indeed, we have made no attempt to optimize the corresponding constant factors, and doing so would require more sophisticated techniques. Moreover, the quantities $\mathcal{L}_{(i)}$, $K_{(i)}$, UCB_t , and LCB_t in Algorithm 1 are chosen as functions of both the kernel and the constants appearing in our assumptions, which limits the algorithm’s practical utility even further. Note, however, that these constants are merely a function of the kernel, and that suitable bounds suffice in place of exact values (e.g., lower bound on ρ_0 , upper bound on c_0 , etc.).

While our algorithm assumes a known time horizon T (which is used when selecting $K_{(i)}$; see Appendix B), this assumption can easily be dropped via a standard doubling trick. The details are given in Appendix A.

3.2. Auxiliary Lemmas

Here we present two very standard auxiliary lemmas. We begin with a simpler version of the conditions of Srinivas *et al.* (Srinivas *et al.*, 2010) guaranteeing that the posterior mean and variance provide valid confidence bounds with high probability. The reason for being slightly simpler is that we are considering a fixed time horizon.

Lemma 1. Fix $\delta \in (0, 1)$. For any finite set of points $\mathcal{L} \subseteq D$ and time horizon T , under the choice $\beta_T = 2 \log \frac{|\mathcal{L}| \cdot T}{\delta}$, it holds that

$$|f(x) - \mu_t(x)| \leq \beta_T^{1/2} \sigma_t(x), \quad \forall x \in \mathcal{L}, t = 1, \dots, T, \quad (14)$$

with probability at least $1 - \delta$.

Proof. It was shown in (Srinivas *et al.*, 2010) that for fixed x and t , the event $|f(x) - \mu_t(x)| \leq \beta_T^{1/2} \sigma_t(x)$ holds with probability at least $1 - e^{-\beta_T/2}$. The lemma follows by substituting the choice of β_T and taking the union bound over the $|\mathcal{L}| \cdot T$ values of x and t . \square

The following lemma is also standard, and has been used (implicitly or explicitly) in the study of multiple algorithms that eliminate suboptimal points based on confidence bounds (de Freitas *et al.*, 2012; Contal *et al.*, 2013; Bogunovic *et al.*, 2016b). For completeness, we give a short proof.

Lemma 2. Suppose that at time t , for all x within a set of points $\tilde{\mathcal{L}} \subseteq D$, it holds that

$$\text{LCB}_t(x) \leq f(x) \leq \text{UCB}_t(x) \quad (15)$$

for some bounds UCB_t and LCB_t such that

$$\max_{x \in \tilde{\mathcal{L}}} |\text{UCB}_t(x) - \text{LCB}_t(x)| \leq 2\eta. \quad (16)$$

Then any point $x \in \tilde{\mathcal{L}}$ satisfying $f(x) < \max_{x' \in \tilde{\mathcal{L}}} f(x') - 4\eta$ must also satisfy

$$\text{UCB}_t(x) < \max_{x \in \tilde{\mathcal{L}}} \text{LCB}_t(x). \quad (17)$$

That is, any 4η -suboptimal point can be ruled out according to the confidence bounds (15).

Proof. We have

$$\text{UCB}_t(x) \leq \text{LCB}_t(x) + 2\eta \quad (18)$$

$$\leq f(x) + 2\eta \quad (19)$$

$$< \max_{x' \in \tilde{\mathcal{L}}} f(x') - 2\eta \quad (20)$$

$$\leq \max_{x' \in \tilde{\mathcal{L}}} \text{UCB}_t(x') - 2\eta \quad (21)$$

$$\leq \max_{x' \in \tilde{\mathcal{L}}} \text{LCB}_t(x'), \quad (22)$$

where (18) and (22) follow from (16), (19) and (21) follow from the confidence bounds in (15), and (20) follows from the assumption $f(x) < \max_{x' \in \tilde{\mathcal{L}}} f(x') - 4\eta$. \square

3.3. Outline of Proof of Theorem 1

Here we provide a high-level outline of the Proof of Theorem 1; the details are given in Appendix B.

Algorithm 1 only samples on a discrete sub-domain \mathcal{L} . This set is chosen to be a set of regularly-spaced points that are fine enough to ensure that the cumulative regret with respect to $\max_{x \in \mathcal{L}} f(x)$ is within a constant value of the cumulative regret with respect to $\max_{x \in D} f(x)$. Working with the finite set \mathcal{L} helps to simplify the subsequent analysis.

We split the epochs into two classes, which we call *early epochs* and *late epochs*. The late epochs are those in which we have shrunk the potential maximizers down enough to be entirely within the locally quadratic region, cf., Figure 1; here we only discuss the second case of Assumption 3, which is the more interesting of the two. Since the width of the locally quadratic region is constant, we can show that this occurs after a *finite number of epochs*, each lasting for at most $O(\log T)$ time. Hence, even if we naively upper bound the instant regret by $2c_0$ according to (8), the overall regret incurred within the early epochs is insignificant.

In the later epochs, we exploit the locally quadratic behavior to show that the set of potential maximizers shrinks rapidly, i.e., by a constant factor after each epoch. As a result, we can let the repeatedly-sampled set $\mathcal{L}_{(i)}$ in Algorithm 1 lie within a given interval that similarly shrinks, thereby controlling the number of samples we need to take in the epoch.

By Lemma 2, after we attain uniform $\eta_{(i)}$ -confidence, the instant regret incurred at each time thereafter is at most $4\eta_{(i)}$. Using the fact that $\eta_{(i)} = \eta_{(0)}2^{-i}$ and summing over the epochs, we find that the overall regret behaves as in (13).

A notable difficulty that we omitted above is how we attain the confidence bounds in order to update the potential maximizers $M_{(i)}$. While we directly apply Lemma 1 for the points that were repeatedly sampled, we found it difficult to do this for the non-sampled points. For those, we instead use Lipschitz properties of the function. In the early epochs, we use the global Lipschitz constant c_1 from Assumption 2, whereas in the later epochs, we find a considerably smaller Lipschitz constant due to the locally quadratic behavior.

4. Lower Bound

Our lower bound is formally stated as follows.

Theorem 2. (Lower Bound) *Consider the one-dimensional BO problem from Section 2.1, with time horizon T and noise variance σ^2 satisfying $\sigma^2 \leq c'_\sigma T^{1-\zeta'}$ for some $c'_\sigma > 0$ and $\zeta' > 0$. Under Assumptions 1, 2, and 4, any algorithm must yield the following: With probability at least $1 - \delta_1 - \delta'_2$ (with respect to the Gaussian process f), the average cumulative regret (averaged over the noisy samples) satisfies*

$$\mathbb{E}[R_T] \geq C'(1 + \sigma\sqrt{T}). \quad (23)$$

Here δ_1 and δ'_2 are defined in Assumptions 2 and 4, and C' depends only on the constants therein and (c'_σ, ζ') .

The assumption that $\sigma^2 \leq c'_\sigma T^{1-\zeta'}$ for some (c'_σ, ζ') is very mild, since typically σ^2 is constant with respect to T . The assumption is required to avoid (23) contradicting the trivial $O(T)$ upper bound. We also note that Theorem 2 immediately implies an $\Omega(1 + \sigma\sqrt{T})$ lower bound on the expected regret $\mathbb{E}[R_T]$ with respect to *both* f and the noisy samples, as long as $1 - \delta_1 - \delta'_2 > 0$. As discussed following Assumption 4, the latter condition is mild.

In the remainder of the section, we introduce some of the main tools and ideas, and then outline the proof. We note that $\mathbb{E}[R_T] = \Omega(1)$ is trivial, as the average regret of the first sample alone is lower bounded by a constant. As a result, we only need to show that $\mathbb{E}[R_T] = \Omega(\sigma\sqrt{T})$.

4.1. Reduction to Binary Hypothesis Testing

Recall that f is a one-dimensional GP on $[0, 1]$ with a stationary kernel $k(x, x')$. We fix $\Delta > 0$, and think of the GP as being generated by the following procedure:

1. Generate a GP f_0 with the same kernel on the larger domain $[-\Delta, 1 + \Delta]$;
2. Randomly shift f_0 along the x -axis by $+\Delta$ or $-\Delta$ with equal probability, to obtain \tilde{f} ;




Figure 2. Examples of functions f_+ and f_- considered in the lower bound. The two are identical up to a small horizontal shift.

3. Let $f(x) = \tilde{f}(x)$ for $x \in [0, 1]$.

Since the kernel is stationary, the shifting does not affect the distribution, so the induced distribution of f is indeed the desired GP on $[0, 1]$.

We consider a genie argument in which f_0 is revealed to the algorithm. Clearly this additional information can only help the algorithm, so any lower bound still remains valid for the original setting. Stated differently, the algorithm knows that f is either f_+ or f_- , where

$$f_+(x) = f_0(x + \Delta), \quad (24)$$

$$f_-(x) = f_0(x - \Delta). \quad (25)$$

See Figure 2 for an illustrative example.

This argument allows us to reduce the BO problem to a *binary hypothesis test with adaptive sampling*, as depicted in Figure 3. The hypothesis, indexed by $v \in \{-, +\}$, is that the underlying function is f_v . We show that under a suitable choice of Δ , achieving small cumulative regret means that we can construct a decision rule $\hat{V}(\mathbf{x})$ such that $\hat{V} = v$ with high probability, i.e., the hypothesis test is successful. The contrapositive statement is then that *if the hypothesis test cannot be successful, we cannot achieve small cumulative regret*, from which it only remains to prove the former. This idea was used previously for stochastic convex optimization in (Raginsky & Rakhlin, 2011).

In the remainder of the analysis, we implicitly condition on an arbitrary realization of f_0 , meaning that all expectations and probabilities are only with respect to the random index V and/or the noise. We also assume that f_0 satisfies the conditions in Assumptions 1, 2, and 4, which holds with probability at least $1 - \delta_1 - \delta'_2$. For sufficiently small Δ , the same assumptions are directly inherited by f_+ and f_- .

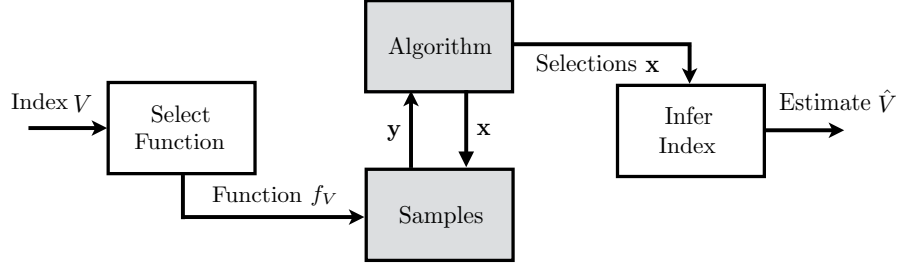


Figure 3. Illustration of reduction from optimization to binary hypothesis testing. The gray boxes are considered to be fixed, whereas the white boxes are introduced for the purpose of proving the lower bound.

We henceforth assume that Δ is indeed sufficiently small; we will verify that this is the case when we set its value.

We introduce some further notation. Letting x_+^* , x_-^* , and x_0^* denote the maximizers of f_+ , f_- and f_0 (which are unique by Assumption 2), we see that Assumption 4 ensures these are in the interior $(0, 1)$, and hence the optimal values coincide: $f_+(x_+^*) = f_-(x_-^*) = f_0(x_0^*) =: f^*$. To simplify some of the notation, instead of working with these functions directly, we consider the equivalent problem of *minimizing the corresponding regret functions*:

$$r_+(x) = f^* - f_+(x), \quad (26)$$

$$r_-(x) = f^* - f_-(x). \quad (27)$$

Indeed, since we assume the algorithm knows f_0 and hence also the optimal value f^* , it can always choose to transform the samples as $y \rightarrow f^* - y$. In this form, we have the convenient normalization $r_+(x_+^*) = r_-(x_-^*) = 0$.

4.2. Auxiliary Lemmas

We first state the following useful properties of r_+ and r_- .

Lemma 3. *The functions r_+ and r_- constructed above satisfy the following for sufficiently small Δ under the conditions in Assumptions 2 and 4:*

1. We have for all $x \in D$ that

$$r_+(x) < c_2 \Delta^2 \implies r_-(x) > c_2 \Delta^2, \quad (28)$$

where c_2 is defined in Assumption 4.

2. There exists a constant $c' > 0$ such that, for all $x \in D$,

$$|r_+(x) - r_-(x)| \leq c' (\Delta |x - x_0^*| + \Delta^2). \quad (29)$$

3. There exists a constant $c'' > 0$ such that, for all $x \in D$,

$$\begin{aligned} r_+(x) &\geq c'' ((x - x_0^*) + \Delta)^2, \\ r_-(x) &\geq c'' ((x - x_0^*) - \Delta)^2. \end{aligned} \quad (30)$$

Proof. See Appendix C. \square

The first part states that any point can be better than $c_2 \Delta^2$ -optimal for at most one of the two functions, the second part shows that the two functions are close for points near x_0^* , and the third part shows that the instant regret is lower bounded by a quadratic function.

The first part of the Lemma 3 allows us to bound the cumulative regret using Fano's inequality for binary hypothesis testing with adaptive sampling (Raginsky & Rakhlin, 2011). This inequality lower bounds the success probability of such a hypothesis test in terms of a mutual information quantity (Cover & Thomas, 2001). The resulting lower bound on regret is stated in the following; it is worth noting that the consideration of *cumulative* regret here provides a distinction from the analogous bound on the instant regret in (Raginsky & Rakhlin, 2011).

Lemma 4. *Under the preceding setup, we have*

$$\mathbb{E}[R_T] \geq c_2 T \Delta^2 \cdot H_2^{-1}(\log 2 - I(V; \mathbf{x}, \mathbf{y})), \quad (31)$$

where V is equiprobable on $\{+, -\}$, and (\mathbf{x}, \mathbf{y}) are the selected points and samples when the minimization algorithm is applied to r_V . Here $H_2^{-1} : [0, \log 2] \rightarrow [0, \frac{1}{2}]$ is the functional inverse of the binary entropy function $H_2(\alpha) = \alpha \log \frac{1}{\alpha} + (1 - \alpha) \log \frac{1}{1 - \alpha}$.

Since this result is particularly fundamental to our analysis, we provide a proof at the end of this section.

4.3. Outline of Proof of Theorem 2

Here we provide a high-level outline of the proof of Theorem 2; the details are given in Appendix D.

Once the lower bound in Lemma 4 is established, the main technical challenge is upper bounding the mutual information. A useful property called *tensorization* (e.g., see (Raginsky & Rakhlin, 2011)) allows us to simplify the mutual information with the vectors (\mathbf{x}, \mathbf{y}) to a sum of mutual informations containing only a single pair (x_t, y_t) : $I(V; \mathbf{x}, \mathbf{y}) \leq \sum_{t=1}^T I(V; y_t | x_t)$.

Each such mutual information term $I(V; y_t | x_t)$ can further be upper bounded by the KL divergence (Cover & Thomas,

2001) between the conditional output distributions corresponding to r_+ and r_- , which in turn equals $\frac{(r_+(x)-r_-(x))^2}{2\sigma^2}$ when $x_t = x$. By substituting the property (29) given in Lemma 3, we find that $I(V; \mathbf{x}, \mathbf{y})$ is upper bounded by a constant times $\frac{1}{\sigma^2}(\Delta^2 \mathbb{E}[\sum_{t=1}^T |x_t - x_0^*|^2] + T\Delta^4)$. If we can further upper bound $I(V; \mathbf{x}, \mathbf{y})$ by a constant in $(0, \log 2)$, then (31) establishes an $\Omega(T\Delta^2)$ lower bound.

We proceed by considering the cases $\mathbb{E}[R_T] \geq c''T\Delta^2$ and $\mathbb{E}[R_T] < c''T\Delta^2$ separately, with c'' given in (30). The former case will immediately give the lower bound in Theorem 2 when we set Δ , whereas in the latter case, we can use (30) to show that $\mathbb{E}[\sum_{t=1}^T |x_t - x_0^*|^2]$ is upper bounded by a constant times $T\Delta^2$, which means that the desired mutual information upper bound (see the previous paragraph) is attained under a choice of Δ scaling as $(\frac{\sigma^2}{T})^{1/4}$. Under this choice, the lower bound $\mathbb{E}[R_T] = \Omega(T\Delta^2)$ evaluates to $\Omega(\sigma\sqrt{T})$, as required.

4.4. Proof of Lemma 4

As mentioned above, the proof of Lemma 4 follows along the lines of (Raginsky & Rakhlin, 2011), which in turn builds on previous works using Fano's inequality to establish minimax lower bounds in statistical estimation problems; see for example (Yu, 1997).

In the following, we use $R_{T,+} = \sum_{t=1}^T r_+(x_t)$ and $R_{T,-} = \sum_{t=1}^T r_-(x_t)$ to denote the cumulative regret associated with r_+ and r_- , respectively, and we generically write $R_{T,v}$ to denote one of the two with $v \in \{+, -\}$.

We first use Markov's inequality to write

$$\mathbb{E}[R_T] \geq (1 - \alpha)c_2T\Delta^2 \cdot \mathbb{P}[R_T \geq (1 - \alpha)c_2T\Delta^2] \quad (32)$$

for any $\alpha \in (0, 1)$. We proceed by analyzing the probability on the right-hand side.

Recall that V is equiprobable on $\{+, -\}$, and (\mathbf{x}, \mathbf{y}) are generated by running the optimization algorithm on r_V . Given the sequence of inputs \mathbf{x} , let \hat{V} be the index $\hat{v} \in \{+, -\}$ with the lower cumulative regret $R_{T,\hat{v}} = \sum_{t=1}^T r_{\hat{v}}(x_t)$. By Lemma 3, R_T can be less than $c_2T\Delta^2$ for at most one of the two functions, and hence, if $R_{T,v} \leq (1 - \alpha)c_2T\Delta^2$ then we must have $\hat{V} = v$. Therefore,

$$\mathbb{P}_v[R_T \geq (1 - \alpha)c_2T\Delta^2] \geq \mathbb{P}_v[\hat{V} \neq v], \quad (33)$$

where, here and subsequently, \mathbb{P}_v and \mathbb{E}_v denote probabilities and expectations when the underlying instant regret function is r_v (i.e., the underlying function that the algorithm seeks to maximize is f_v).

Continuing, we can lower bound the probability appearing

in (32) as follows:

$$\begin{aligned} \mathbb{P}[R_T \geq (1 - \alpha)c_2T\Delta^2] \\ = \frac{1}{2} \sum_{v \in \{+, -\}} \mathbb{P}_v[R_T \geq (1 - \alpha)c_2T\Delta^2] \end{aligned} \quad (34)$$

$$\geq \frac{1}{2} \sum_{v \in \{+, -\}} \mathbb{P}_v[\hat{V} \neq v] \quad (35)$$

$$\geq H_2^{-1}(\log 2 - I(V; \mathbf{x}, \mathbf{y})), \quad (36)$$

where (35) follows from (33), and (36) follows from Fano's inequality for binary hypothesis testing with adaptive sampling (see Eq. (22) and (24) of (Raginsky & Rakhlin, 2011)). The proof is completed by combining (32) and (36), and recalling that α can be arbitrarily small.

5. Conclusion and Discussion

We have established tight scaling laws on the regret for Bayesian optimization in one dimension, showing that the optimal scaling is $\Omega(\sqrt{T})$ and $O(\sqrt{T \log T})$ under mild technical assumptions on the kernel. Our results highlight some limitations of the widespread upper bounds based on the information gain, as well as providing cases where the noisy Bayesian setting is provably more difficult than its non-Bayesian RKHS counterpart.

An immediate direction for further work is to sharpen the constant factors in the upper and lower bounds, and to establish whether the upper bound is attained by any algorithm that can also provide state-of-the-art performance in practice. We re-iterate that our algorithm is certainly not suitable for this purpose, as its cumulative regret contains large constant factors, and the algorithm makes use of a variety of specific constants present in the assumptions (though they are merely a function of the kernel).

We expect our techniques to extend to any constant dimension $d \geq 1$; the main ideas from the noiseless upper bound still apply (de Freitas et al., 2012), and in the lower bound we can choose an arbitrary single dimension and introduce a random shift in that direction as per Section 4.1. While these extensions may still yield \sqrt{T} poly(log T) regret, the dependence on d would be exponential or worse in the upper bound, but constant in the lower bound, with the latter dependence certainly being suboptimal. Multi-dimensional lower bounding techniques based on Fano's inequality (Raginsky & Rakhlin, 2011) may improve the latter to poly(d), but overall, attaining a sharp joint dependence on T and d appears to require different techniques.

Acknowledgments. I would like to thank Ilija Bogunovic for his helpful comments and suggestions. This work was supported by an NUS startup grant.

References

- Bogunovic, I., Scarlett, J., and Cevher, V. Time-varying Gaussian process bandit optimization. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2016a.
- Bogunovic, I., Scarlett, J., Krause, A., and Cevher, V. Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation. In *Conf. Neur. Inf. Proc. Sys. (NIPS)*, 2016b.
- Bubeck, S. and Cesa-Bianchi, N. *Regret Analysis of Stochastic and Nonstochastic Multi-Armed Bandit Problems*. Found. Trend. Mach. Learn. Now Publishers, 2012.
- Bull, A. D. Convergence rates of efficient global optimization algorithms. *J. Mach. Learn. Res.*, 12(Oct.):2879–2904, 2011.
- Contal, E., Buffoni, D., Robicquet, A., and Vayatis, N. *Machine Learning and Knowledge Discovery in Databases*, chapter Parallel Gaussian Process Optimization with Upper Confidence Bound and Pure Exploration, pp. 225–240. Springer Berlin Heidelberg, 2013.
- Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. John Wiley & Sons, Inc., 2001.
- de Freitas, N., Zoghi, M., and Smola, A. J. Exponential regret bounds for Gaussian process bandits with deterministic observations. In *Int. Conf. Mach. Learn. (ICML)*, 2012.
- Desautels, T., Krause, A., and Burdick, J. W. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *J. Mach. Learn. Res.*, 15(1):3873–3923, 2014.
- Grünewälder, S., Audibert, J.-Y., Opper, M., and Shawe-Taylor, J. Regret bounds for Gaussian process bandit problems. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, pp. 273–280, 2010.
- Hennig, P. and Schuler, C. J. Entropy search for information-efficient global optimization. *J. Mach. Learn. Research*, 13(1):1809–1837, 2012.
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. In *Adv. Neur. Inf. Proc. Sys. (NIPS)*, pp. 918–926, 2014.
- Kandasamy, K., Schneider, J., and Póczos, B. High dimensional Bayesian optimisation and bandits via additive models. In *Int. Conf. Mach. Learn.*, 2015.
- Kawaguchi, K., Kaelbling, L. P., and Lozano-Pérez, T. Bayesian optimization with exponential convergence. In *Conf. Neur. Inf. Proc. Sys. (NIPS)*, 2015.
- Kleinberg, R., Slivkins, A., and Upfal, E. Multi-armed bandits in metric spaces. In *Proc. ACM Symp. Theory Comp. (STOC)*, pp. 681–690, 2008.
- Krause, A. and Ong, C. S. Contextual Gaussian process bandit optimization. In *Conf. Neur. Inf. Proc. Sys. (NIPS)*, pp. 2447–2455. Curran Associates, Inc., 2011.
- Raginsky, M. and Rakhlin, A. Information-based complexity, feedback and dynamics in convex programming. *IEEE Trans. Inf. Theory*, 57(10):7036–7056, Oct. 2011.
- Rasmussen, C. E. *Gaussian processes for machine learning*. MIT Press, 2006.
- Rolland, P., Scarlett, J., Bogunovic, I., and Cevher, V. High-dimensional Bayesian optimization via additive models with overlapping groups. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2018.
- Russo, D. and Van Roy, B. Learning to optimize via information-directed sampling. In *Conf. Neur. Inf. Proc. Sys. (NIPS)*, 2014.
- Scarlett, J., Bogunovic, I., and Cevher, V. Lower bounds on regret for noisy Gaussian process bandit optimization. In *Conf. Learn. Theory (COLT)*. 2017.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE*, 104(1):148–175, 2016.
- Shekhar, S. and Javidi, T. Gaussian process bandits with adaptive discretization. <http://arxiv.org/abs/1712.01447>, 2017.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Int. Conf. Mach. Learn. (ICML)*, 2010.
- Wang, Z., Shakibi, B., Jin, L., and de Freitas, N. Bayesian multi-scale optimistic optimization. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, pp. 1005–1014, 2014.
- Wang, Z., Zhou, B., and Jegelka, S. Optimization as estimation with Gaussian processes in bandit settings. In *Int. Conf. Art. Intel. Stats. (AISTATS)*, 2016.
- Yu, B. Assouad, Fano, and le Cam. In *Festschrift for Lucien Le Cam*, pp. 423–435. Springer, 1997.