# Supplementary Material

## Tight Regret Bounds for Bayesian Optimization in One Dimension (Jonathan Scarlett, ICML 2018)

## A. Doubling Trick for an Unknown Time Horizon

Suppose that we have an algorithm that depends on the time horizon $T'$ and achieves $\mathbb{E}[R_{T'}] \le C\sqrt{T' \log T'}$ for some $C > 0$. We show that we can also achieve $\mathbb{E}[R_T] = O\big(\sqrt{T \log T}\big)$ when $T$ is unknown.

To see this, fix an arbitrary integer $T_0 \in \big[1, \frac{T}{2}\big]$, and repeatedly run the algorithm with fixed time horizons $T_0$, $2T_0$, $4T_0$, etc., until $T$ points have been sampled. The number of stages is no more than $\ell_{\max} = \lceil \log_2 \frac{T}{T_0} \rceil$. Moreover, we have

$$\mathbb{E}[R_T] \le \sum_{\ell=1}^{\ell_{\max}} C\sqrt{2^{\ell-1} T_0 \log T} = C\sqrt{T_0 \log T} \sum_{\ell=0}^{\lceil \log_2 \frac{T}{T_0} \rceil - 1} \sqrt{2^\ell} \le C\sqrt{\log T} \cdot 4\sqrt{T} \tag{37}$$

where the first inequality uses $\log(2^{\ell-1}T_0) \le \log T$, and the last inequality uses $\sum_{\ell=0}^N 2^{\ell/2} \le 4 \cdot 2^{N/2}$. This establishes the desired claim.

## B. Proof of Theorem 1 (Upper Bound)

We continue from the auxiliary results given in Section 3, proceeding in several steps. Algorithm 2 gives a full description of the algorithm; the reader is encouraged to refer to this throughout the proof, rather than trying to understand all the steps therein immediately. Note that the constants $c_0$, $c_1$, $c_2$, and $\rho_0$ used in the algorithm come from Assumptions 2 and 3.

**Reduction to a finite domain.** Our algorithm only samples $f$ within a finite set $\mathcal{L} \subseteq D$ of pre-defined points. We choose these points to be regularly spaced, and close enough to ensure that the highest function value is within $\frac{1}{T}$ of the maximum $f(x^*)$. Under condition (8) in Assumption 2 (which implies that $f$ is $c_1$-Lipschitz continuous), it suffices to choose

$$\mathcal{L} = \bigg(\frac{1}{c_1 \cdot T}\mathbb{Z} \cap [0,1]\bigg) \cup \{1\}, \tag{38}$$

where $\mathbb{Z}$ denotes the integers. Here we add $x = 1$ to $\mathcal{L}$ because it will be notationally convenient to ensure that the endpoints $\{0, 1\}$ are both included in the set. Note that $\mathcal{L}$ satisfies $|\mathcal{L}| \le c_1 T + 1$, which we crudely upper bound by $|\mathcal{L}| \le 2c_1 T$.

Since $\max_{x\in\mathcal{L}} f(x) \ge \max_{x \in D} f(x) - \frac{1}{T}$, the cumulative regret $R_T^{(\mathcal{L})}$ with respect to the best point in $\mathcal{L}$ is such that

$$R_T \le R_T^{(\mathcal{L})} + 1. \tag{39}$$

Hence, it suffices to bound $R_T^{(\mathcal{L})}$ instead of $R_T$. For convenience, we henceforth let $x^*_{\mathcal{L}}$ denote an arbitrary input that achieves $\max_{x\in\mathcal{L}} f(x)$, and we define the *instant regret* as

$$r(x) = f(x^*) - f(x), \quad r_t = r(x_t) = f(x^*) - f(x_t), \quad r_t^{(\mathcal{L})} = f(x^*_{\mathcal{L}}) - f(x_t). \tag{40}$$

**Conditioning on high-probability events.** By assumption, the events in Assumptions 2 and 3 simultaneously hold with probability at least $1 - \delta_1 - \delta_2$. Moreover, by setting $\delta = \frac{1}{T}$ in Lemma 1 and letting $\mathcal{L}$ be as in (38) with $|\mathcal{L}| \le 2c_1 T$, we deduce that (14) holds with probability at least $1 - \frac{1}{T}$ when

$$\beta_T = 2\log\big(2c_1 T^3\big). \tag{41}$$

Denoting the intersection of all events in Assumptions 2 and 3 by $\mathcal{A}$, and the event in Lemma 1 by $\mathcal{B}$, we can write the average regret given $\mathcal{A}$ as follows:

$$\mathbb{E}[R_T|\mathcal{A}] = \mathbb{E}[R_T|\mathcal{A},\mathcal{B}] \cdot \mathbb{P}[\mathcal{B}|\mathcal{A}] + \mathbb{E}[R_T|\mathcal{A},\mathcal{B}^c] \cdot \mathbb{P}[\mathcal{B}^c|\mathcal{A}] \tag{42}$$

$$\le \mathbb{E}[R_T|\mathcal{A},\mathcal{B}] + \mathbb{E}[R_T|\mathcal{A},\mathcal{B}^c]\frac{1}{T(1-\delta_1-\delta_2)} \tag{43}$$

$$\le \mathbb{E}[R_T|\mathcal{A},\mathcal{B}] + \frac{2c_0}{1-\delta_1-\delta_2}, \tag{44}$$

**Algorithm 2** Full description of our algorithm, based on reducing uncertainty in epochs via repeated sampling.

**Require:** Domain $D$, GP prior $(\mu_0, k_0)$, time horizon $T$, constants $c_0, c_1, c_2, \rho_0$.

Set discrete sub-domain $\mathcal{L} = \big(\frac{1}{c_1 \cdot T}\mathbb{Z} \cap [0,1]\big) \cup \{1\}$, confidence parameter $\beta_T = 2\log(2c_1T^3)$, initial target confidence $\eta_{(0)} = c_0$, and initial potential maximizers $M_{(0)} = \mathcal{L}$.

Initialize time index $t = 1$ and epoch number $i = 1$.

**while** less than $T$ samples have been taken **do**

Set $\eta_{(i)} = \frac{1}{2}\eta_{(i-1)}$.

Define the interval

$$\mathcal{I}_{(i)} = \big[\min\{x \in M_{(i-1)}\}, \max\{x \in M_{(i-1)}\}\big] \cap \mathcal{L},$$

and its width

$$w_{(i)} = \max\{x \in M_{(i-1)}\} - \min\{x \in M_{(i-1)}\}.$$

Set the Lipschitz constant

$$L_{(i)} = \begin{cases} c_1 & w_{(i)} > \rho_0 \\ c_1 & w_{(i)} \le \rho_0 \text{ and either } 0 \in \mathcal{I}_{(i)} \text{ or } 1 \in \mathcal{I}_{(i)} \\ c_2 w_{(i)} & w_{(i)} \le \rho_0 \text{ and } \mathcal{I}_{(i)} \subseteq (0,1). \end{cases}$$

Construct a subset $\mathcal{L}_{(i)} \subseteq \mathcal{I}_{(i)}$ as follows:

- Initialize $\mathcal{L}_{(i)} \leftarrow \emptyset$.
- Construct $\widetilde{\mathcal{L}}_{(i)}$ (not necessarily a subset of $\mathcal{I}_{(i)}$ or $\mathcal{L}$) containing regularly-spaced points within the interval $\big[\min\{x \in \mathcal{I}_{(i)}\}, \max\{x \in \mathcal{I}_{(i)}\}\big]$, with spacing $\frac{\eta_{(i)}}{2L_{(i)}}$.
- For each $x \in \widetilde{\mathcal{L}}_{(i)}$, add its two nearest points in $\mathcal{I}_{(i)}$ to $\mathcal{L}_{(i)}$.

Sample each point in $\mathcal{L}_{(i)}$ repeatedly $K_{(i)}$ times, where

$$K_{(i)} = \left\lceil \frac{4\sigma^2\beta_T}{\eta_{(i)}^2} \right\rceil.$$

For each sample taken, increment $t \leftarrow t+1$, and terminate if $t > T$.

Update the posterior distribution $(\mu_{t-1}, \sigma_{t-1})$ according to (5)–(6), with $\mathbf{x}_{t-1} = [x_1, \dots, x_{t-1}]^T$ and $\mathbf{y}_{t-1} = [y_1, \dots, y_{t-1}]^T$ respectively containing all the selected points and noisy samples so far.

For each $x \in \mathcal{I}_{(i)}$, set

$$\mathrm{UCB}_t(x) = \mu_{t-1}(x') + \eta_{(i)}, \quad \mathrm{LCB}_t(x) = \mu_{t-1}(x') - \eta_{(i)},$$

where $x' = \arg\min_{x' \in \mathcal{L}_{(i)}} |x - x'|$.

Update the set of potential maximizers:

$$M_{(i)} = \Big\{ x \in M_{(i-1)} \,:\, \mathrm{UCB}_t(x) \ge \max_{x' \in M_{(i-1)}} \mathrm{LCB}_t(x') \Big\}.$$

Increment $i$.

**end while**

where (43) follows since $\mathbb{P}[\mathcal{B}|\mathcal{A}] \le 1$ and $\mathbb{P}[\mathcal{B}^c|\mathcal{A}] \le \frac{\mathbb{P}[\mathcal{B}^c]}{\mathbb{P}[\mathcal{A}]} \le \frac{1}{T(1-\delta_1-\delta_2)}$, and (44) follows since condition (8) in Assumption 2 ensures that $R_T \le T \cdot 2c_0$. By (44), in order to prove Theorem 1, it suffices to show that $R_T = O(\sqrt{T \log T})$ whenever the conditions of Assumptions 2–3 and Lemma 1 hold true. We henceforth condition on this being the case.

**Sampling mechanism.** Recall that $\eta_{(i)}$ represents the target confidence to attain by the end of the $i$-th epoch, and each such value is half of the previous value. For this interpretation to be valid, $\eta_{(0)}$ should be sufficient large so that the entire function is *a priori* known up to confidence $\eta_{(0)}$; by (8) in Assumption 2, the choice $\eta_{(0)} = c_0$ certainly suffices for this purpose.

In the $i$-th epoch, we repeatedly sample a *sufficiently fine* subset of $\mathcal{L}$ *sufficiently many times* to attain an overall confidence of $\eta_{(i)}$ within $M_{(i-1)}$ (with $M_{(0)} = \mathcal{L}$). Specifically:

- We sample each point $K_{(i)}$ times and average the resulting observations, yielding an *effective noise variance* of $\frac{\sigma^2}{K_{(i)}}$, and we choose $K_{(i)}$ large enough so that $\frac{\sigma^2}{K_{(i)}} \le \frac{\eta_{(i)}^2}{4\beta_T}$. Hence, $K_{(i)} = \lceil \frac{4\sigma^2\beta_T}{\eta_{(i)}^2} \rceil$ is sufficient.
- To design $\mathcal{L}_{(i)} \subseteq \mathcal{L}$, we consider the interval

$$\mathcal{I}_{(i)} = \big[ \min\{x \in M_{(i-1)}\}, \max\{x \in M_{(i-1)}\} \big] \cap \mathcal{L}, \tag{45}$$

which is the smallest interval (intersected with $\mathcal{L}$) containing $M_{(i-1)}$. We select a Lipschitz constant $L_{(i)}$ (to be specified later) such that $f$ is $L_{(i)}$-Lipschitz within $\mathcal{I}_{(i)}$, and then we choose $\mathcal{L}_{(i)} \subseteq \mathcal{I}_{(i)}$ to ensure the following:

$$\text{Each } x \in \mathcal{I}_{(i)} \text{ is within a distance } \frac{\eta_{(i)}}{2L_{(i)}} \text{ of the nearest } x' \in \mathcal{L}_{(i)}. \tag{46}$$

If we were sampling at arbitrary locations, it would suffice to choose $\lceil \frac{2w_{(i)}L_{(i)}}{\eta_{(i)}} \rceil$ equally-spaced points, where

$$w_{(i)} = \max\{x \in M_{(i-1)}\} - \min\{x \in M_{(i-1)}\} \tag{47}$$

is the width of the interval. With the restriction of sampling within the fine discretization $\mathcal{L}$, we can simply "round" to the two nearest points,[1] yielding a suitable set $\mathcal{L}_{(i)} \subseteq \mathcal{I}_{(i)}$ of cardinality at most $2\lceil \frac{2w_{(i)}L_{(i)}}{\eta_{(i)}} \rceil$

Combining these, the total number of samples $T_{(i)}$ is given by

$$\begin{aligned} T_{(i)} &= K_{(i)} \cdot |\mathcal{L}_{(i)}| && (48)\\ &\le 2 \cdot \Big\lceil \frac{4\sigma^2\beta_T}{\eta_{(i)}^2} \Big\rceil \cdot \Big\lceil \frac{2w_{(i)}L_{(i)}}{\eta_{(i)}} \Big\rceil. && (49) \end{aligned}$$

At the points that were sampled, we performed enough repetitions to attain a variance of at most $\frac{\eta_{(i)}^2}{4\beta_T}$ based on those samples alone. The information from any earlier samples only reduces the variance further, so the overall posterior variance[2] $\sigma_{t-1}^2(x)$ also yields $\beta_T^{1/2}\sigma_{t-1}(x) \le \frac{\eta_{(i)}}{2}$. Hence, Lemma 1 ensures that at these sampled points, we can set

$$\widetilde{\mathrm{UCB}}_t(x) = \mu_{t-1}(x) + \frac{\eta_{(i)}}{2}, \quad \widetilde{\mathrm{LCB}}_t(x) = \mu_{t-1}(x) - \frac{\eta_{(i)}}{2}. \tag{50}$$

For the points in $M_{(i-1)}$ that we didn't sample, we note that the following confidence bounds are valid as long as $f$ is $L_{(i)}$-Lipschitz continuous within $\mathcal{I}_{(i)}$:

$$\widetilde{\mathrm{UCB}}_t(x) = \mu_{t-1}(x') + \frac{\eta_{(i)}}{2} + L_{(i)}|x - x'|, \tag{51}$$

$$\widetilde{\mathrm{LCB}}_t(x) = \mu_{t-1}(x') - \frac{\eta_{(i)}}{2} - L_{(i)}|x - x'|, \tag{52}$$

[1] To give a concrete example, suppose that $\mathcal{L} = \{0, 0.01, \dots, 0.99, 1\}$, and that we seek a set of points such that each $x \in \mathcal{L}$ is within a distance $\frac{1}{3}$ of the nearest one. Without constraints, the points $\{\frac{1}{3}, \frac{2}{3}\}$ would suffice, but after rounding these to $\{0.33, 0.66\}$, the point $x = 1$ is at a distance $0.34 > \frac{1}{3}$. However, doubling up and constructing the set $\{0.33, 0.34, 0.66, 0.67\}$ clearly suffices.

[2] We consider $(\mu_{t-1}, \sigma_{t-1})$ instead of $(\mu_t, \sigma_t)$ because when the time index is $t$, we have only selected $t-1$ points.

where $x' = \arg\min_{x' \in \mathcal{L}_{(i)}} |x - x'|$ is the closest sampled point to $x$. If $x$ is itself in $\mathcal{L}_{(i)}$, these expressions reduce to (50).

Now, since we have ensured the condition (46), we find that we can weaken (51)–(52) to

$$\mathrm{UCB}_t(x) = \mu_{t-1}(x') + \eta_{(i)}, \quad \mathrm{LCB}_t(x) = \mu_{t-1}(x') - \eta_{(i)}. \tag{53}$$

That is, as long as the Lipschitz constant $L_{(i)}$ is valid, we have $\eta_{(i)}$-confidence at the end of the $i$-th epoch. As a result, by Lemma 2, the updated set of potential maximizers

$$M_{(i)} = \Big\{ x \in M_{(i-1)} \,:\, \mathrm{UCB}_t(x) \ge \max_{x' \in \mathcal{L}} \mathrm{LCB}_t(x) \Big\}, \tag{54}$$

with $t$ being the ending time of the epoch, must only contain points within $\mathcal{L}$ whose function value is within $4\eta_{(i)}$ of $f(x^*_{\mathcal{L}})$. Below, we will choose $L_{(i)}$ differently in different epochs, while still ensuring the required Lipschitz condition is valid.

**Analysis of early epochs.** Recall the following:

- By Assumption 1, the constant $\epsilon$ lower bounds the separation between $f(x^*)$ and the function value at the second highest local maximum (if any).
- By Assumption 3, we either have $x^*$ at an endpoint and the locally linear behavior (9), or we have $x^* \in (\rho_0, 1 - \rho_0)$ and the locally quadratic behavior (10).

In the epochs for which $w^{(i)} > \rho_0$, we choose $L_{(i)} = c_1$ (*cf.*, (8)), which is clearly a valid Lipschitz constant. We claim that *after a finite number of epochs*, all points $x \in M_{(i)}$ satisfy $f(x) > f(x^*) - \epsilon$ and $|x - x^*| \le \frac{\rho_0}{2}$, and therefore, $w_{(i)}$ ceases to be greater than $\rho_0$. We henceforth distinguish between the two cases using the terminology *early epochs* and *late epochs*.

To see that the preceding claim is true, we consider the two cases of Assumption 3:

- In the first case, all points satisfying $|x - x^*| > \rho_0$ are at least $\min\{\underline{c}_1 \rho_0, \epsilon\}$-suboptimal by the locally linear behavior (9) and the $\epsilon$ gap (7);
- In the second case, all points satisfying $|x - x^*| > \rho_0$ are at least $\min\{\underline{c}_2 \rho_0^2, \epsilon\}$-suboptimal by the locally quadratic behavior (9) and the $\epsilon$ gap (7).

Hence, in either case, all points satisfying $|x - x^*| > \rho_0$ are at least $\epsilon'$-suboptimal, where $\epsilon' = \min\{\underline{c}_1 \rho_0, \underline{c}_2 \rho_0^2, \epsilon\}$. As a result, to establish the desired claim, we only need to show that $M_{(i)}$ contains no points with instant regret $r(x) \ge \epsilon'$.

Since $f(x^*_{\mathcal{L}}) \ge f(x^*) - \frac{1}{T}$ (as stated following (38)), we find that as long as $T > \frac{2}{\epsilon'}$,[3] it suffices that $M_{(i)}$ only contains points such that $r_t^{(\mathcal{L})}(x) \le \frac{\epsilon'}{2}$. By Lemma 2, this happens as soon as $\eta_{(i)} < \frac{\epsilon'}{8}$. Since $\epsilon'$ is constant and we halve $\eta_{(i)}$ at the end of each epoch, it must be that only a finite number of epochs $i_{\max,1}$ pass before this occurs, with $i_{\max,1}$ depending only on $\eta_{(0)}$ and $\epsilon'$.

For these early epochs, we simply upper bound $w_{(i)}$ in (49) by one, meaning their overall cumulative time $T_{\mathrm{early}}$ satisfies

$$T_{\mathrm{early}} \le \sum_{i=1}^{i_{\max,1}} T_{(i)} \le 2 i_{\max,1} \bigg\lceil \frac{256 \sigma^2 \beta_T}{(\epsilon')^2} \bigg\rceil \cdot \bigg\lceil \frac{16 c_1}{\epsilon'} \bigg\rceil, \tag{55}$$

where we have used the fact that $\eta_{(i)} \ge \frac{\epsilon'}{8}$ and $L_{(i)} = c_1$ in these epochs.

**Analysis of late epochs.** Recall that we consider ourselves in a late epoch as soon as $w_{(i)} \le \rho_0$. This condition implies that all points in $M_{(i-1)}$ are within a distance $\rho_0$ of $x^*$,[4] yielding the locally linear behavior (9) if $x^*$ is an endpoint, and the locally quadratic behavior (10) otherwise. Moreover, Assumption 3 assumes $x^* \in (\rho_0, 1 - \rho_0)$ in the latter case, and as a result, the algorithm can identify which case has occurred: If $\mathcal{I}_{(i)}$ contains an endpoint, then we are in the first case, whereas if $\mathcal{I}_{(i)} \subseteq (0, 1)$, then we are in the second case.

[3] It is safe to assume that $T$ is sufficiently large, since the smaller values of $T$ can be handled by increasing $C$ in the theorem statement.

[4] Since we condition on the confidence bounds in Lemma 1 being valid, only points that are truly suboptimal are ever ruled out.

Accordingly, the algorithm can choose the Lipschitz constant $L_{(i)}$ differently in the two cases. In the first case, we simply continue to use the global choice $L_{(i)} = c_1$ from (8). In the second case, we observe that $f'(x^*) = 0$, and recall from (8) that $f'$ is $c_2$-Lipschitz continuous. Since the width of the interval of interest $\mathcal{I}_{(i)}$ is $w_{(i)}$, we conclude that $|f'(x)| \le c_2 w_{(i)}$ within $\mathcal{I}_{(i)}$, and accordingly, we can set

$$L_{(i)} = c_2 w_{(i)}. \tag{56}$$

We initially focus on this second case (which is the more interesting of the two), and later return to the first case.

Recall that within the $i$-th epoch, all points with $f(x) < f(x^*_{\mathcal{L}}) - 4\eta_{(i-1)}$ have already been removed from the potential maximizers (*cf.*, Lemma 2). This implies that the points sampled incur instant regret at most

$$r_t^{(\mathcal{L})} \le 4\eta_{(i-1)}, \tag{57}$$

and hence, since we have established that $f(x^*_{\mathcal{L}}) \ge f(x^*) - \frac{1}{T}$,

$$r_t \le 4\eta_{(i-1)} + \frac{1}{T}. \tag{58}$$

From this fact and the locally quadratic behavior (10), we deduce that the width $w_{(i)}$ defined in (47) satisfies $w_{(i)} \le \sqrt{\frac{4\eta_{(i-1)}+\frac{1}{T}}{\underline{c}_2}} = \sqrt{\frac{8\eta_{(i)}+\frac{1}{T}}{\underline{c}_2}}$ (since $\eta_{(i-1)} = 2\eta_{(i)}$), from which (49) and (56) yield

$$T_{(i)} \le 2\Big\lceil \frac{4\sigma^2\beta_T}{\eta_{(i)}^2} \Big\rceil \cdot \Big\lceil \frac{2c_2}{\underline{c}_2} \cdot \Big(8 + \frac{1}{T\eta_{(i)}}\Big) \Big\rceil. \tag{59}$$

Grouping all the constants together and writing $\lceil z \rceil \le 1 + z$, we can simplify this to

$$T_{(i)} \le c'\bigg(1 + \frac{1}{T\eta_{(i)}} + \frac{\sigma^2\beta_T}{\eta_{(i)}^2} + \frac{\sigma^2\beta_T}{T\eta_{(i)}^3}\bigg) \tag{60}$$

for suitably-chosen $c' > 0$.

**Bounding the cumulative regret.** In the early epochs, we crudely upper bound the regret at each time instant by $2c_0$ (*cf.*, (8)). Hence, since the total cumulative time of these epochs satisfies (55) for bounded $i_{\max,1}$, and $\beta_T = O(\log T)$ as per (41), the corresponding total cumulative regret $R^{(\mathcal{L})}_{\text{early}}$ is upper bounded by

$$R^{(\mathcal{L})}_{\text{early}} \le c''(1 + \sigma^2 \log T) \tag{61}$$

for some $c'' > 0$.

For the late epochs, we make use of the instant regret bound in (57), depending on the epoch index $i$. Since this upper bound is decreasing in $i$, and the epoch lengths satisfy (60), we can upper bound $R_T^{(\mathcal{L})}$ by considering the hypothetical case that the epoch lengths are *exactly* the right-hand side of (60), and the instant regret incurred at time $t$ is *exactly* $r_t^{(\mathcal{L})} = 4\eta_{(i-1)}$.

In this situation, we can easily upper bound the total number of epochs: The last epoch must certainly be no larger than $i_{\max,2}$, defined to be the smallest $i$ such that the term $c'\frac{\sigma^2\beta_T}{\eta_{(i)}^2}$ on the right-hand side of (60) is $T$ or higher. Substituting $\eta_{(i)} = \frac{\eta_{(0)}}{2^i}$ and re-arranging, we conclude that

$$i_{\max,2} \le \log_4 \frac{T\eta_{(0)}^2}{c'\sigma^2\beta_T} = \log_2 \sqrt{\frac{T\eta_{(0)}^2}{c'\sigma^2\beta_T}}. \tag{62}$$

For technical reasons, here and subsequently we can assume without loss of generality that $\sigma \le \kappa\sqrt{\frac{T}{\log T}}$ for arbitrarily small $\kappa > 0$ and sufficiently large $T$; otherwise, Theorem 1 states the trivial bound $\mathbb{E}[R_T] \le CT$. Since $\beta_T = \Theta(\log T)$, this technical condition means the right-hand side of (62) exceeds one.

Continuing, the total cumulative regret $R_{\rm late}^{(\mathcal{L})}$ from the late epochs is upper bounded as follows:

$$R_{\rm late}^{(\mathcal{L})} \le \sum_{i=1}^{i_{\max,2}} 4\eta_{(i-1)} T^{(i)} \tag{63}$$

$$\le 4c' \sum_{i=1}^{i_{\max,2}} \eta_{(i-1)} + 8c' \bigg( \sum_{i=1}^{i_{\max,2}} 1 \bigg) + 8c'\sigma^2\beta_T \sum_{i=1}^{i_{\max,2}} \frac{1}{\eta_{(i)}} + \frac{8c'\sigma^2\beta_T}{T} \sum_{i=1}^{i_{\max,2}} \frac{1}{\eta_{(i)}^2} \tag{64}$$

$$\le 4c' i_{\max,2}(\eta_{(0)} + 2) + 8c'\sigma^2\beta_T \sum_{i=1}^{i_{\max,2}} \frac{1}{\eta_{(i)}} + \frac{8c'\sigma^2\beta_T}{T} \sum_{i=1}^{i_{\max,2}} \frac{1}{\eta_{(i)}^2} \tag{65}$$

$$\le 4c' i_{\max,2}(\eta_{(0)} + 2) + \frac{8c'\sigma^2\beta_T}{\eta_{(0)}} \sum_{i=1}^{i_{\max,2}} 2^i + \frac{8c'\sigma^2\beta_T}{T\eta_{(0)}^2} \sum_{i=1}^{i_{\max,2}} 4^i \tag{66}$$

$$\le 4c' i_{\max,2}(\eta_{(0)} + 2) + \frac{16c'\sigma^2\beta_T}{\eta_{(0)}} 2^{i_{\max,2}} + \frac{16c'\sigma^2\beta_T}{T\eta_{(0)}^2} 4^{i_{\max,2}} \tag{67}$$

$$\le 4c'(\eta_{(0)} + 2) \log_4 \frac{T\eta_{(0)}^2}{c'\sigma^2\beta_T} + 16\sqrt{c'\sigma^2\beta_T T} + 16, \tag{68}$$

where (64) follows from (60) and the fact that $\eta_{(i-1)} = 2\eta_{(i)}$, (65) follows since $\eta_{(i-1)} \le \eta_{(0)}$, (66) follows since $\eta_{(i)} = \frac{\eta_{(0)}}{2^i}$, (67) follows since $\sum_{i=1}^N 2^i \le 2 \cdot 2^N$ and $\sum_{i=1}^N 4^i \le 2 \cdot 4^N$, and (68) follows by substituting the upper bound on $i_{\max,2}$ from (62). Using the fact that $\beta_T = O(\log T)$, and recalling that $\eta_{(0)} = c_0$ is constant, we simplify (68) to

$$R_{\rm late}^{(\mathcal{L})} \le c^\dagger \Big( 1 + \sigma\sqrt{T \log T} \Big) \tag{69}$$

for some $c^\dagger > 0$. Note that we can safely drop the $O\big(\log \frac{T}{\sigma^2\beta_T}\big) = O\big(\log \frac{T}{\sigma^2 \log T}\big)$ term in (68) due to the assumption $\sigma^2 \ge \frac{c_\sigma}{T^{1-\zeta}}$ in Theorem 1.

**Handling the first case in Assumption 3.** From (56) onwards, we focused only on the second case of Assumption 3. In the first case, we have a worse Lipschitz constant $L_{(i)} = c_1$, but the width also shrinks faster: By the locally linear behavior (9), achieving $\eta_{(i)}$-confidence not only brings the interval width $w_{(i)}$ down to at most $O(\sqrt{\eta_{(i)}})$, but also further down to $O(\eta_{(i)})$. Hence, we lose a factor of $\sqrt{\eta_{(i)}}$ in the Lipschitz constant, but we gain a factor of $\sqrt{\eta_{(i)}}$ in the upper bound on $w_{(i)}$. Since the number of points sampled in (49) contains the product of the two, the final result remains unchanged, i.e., we still have (69), possibly with a different constant $c^\dagger$.

**Completion of the proof.** Combining (39), (44), (61) and (69), we obtain

$$\mathbb{E}[R_T] \le C^\dagger \big( 1 + \sigma^2 \log T + \sigma\sqrt{T \log T} \big) \tag{70}$$

for some constant $C^\dagger$. As stated following (62), we can assume without loss of generality that $\sigma \le O\Big(\sqrt{\frac{T}{\log T}}\Big)$, which means that the third term of (70) dominates the second, and the proof is compete.

## C. Proof of Lemma 3

For the first part, we consider $\Delta$ sufficiently small so that $\underline{c}_2\Delta^2 < \epsilon$, for $\epsilon$ given in Assumption 2 and $\underline{c}_2$ in Assumption 4. Since all local maxima are at least $\epsilon$-suboptimal, achieving $r_+(x) < \underline{c}_2\Delta^2$ requires that $x$ lies within a small interval around $x_+^*$. Moreover, the locally quadratic behavior (12) in Assumption 4 yields $r_+(x) \ge \underline{c}_2(x - x_+^*)^2$ within this interval when $\Delta$ is sufficiently small. Combining this with $r_+(x) < \underline{c}_2\Delta^2$ gives $|x - x_+^*| < \Delta$, and since $|x_+^* - x_-^*| = 2\Delta$, the triangle inequality yields $|x - x_-^*| > \Delta$. Again using (12), we conclude that $r_-(x) > \underline{c}_2\Delta^2$, as required.

For the second part, we recall from (24)–(25) that $r_+(x) = r_0(x + \Delta)$ and $r_-(x) = r_0(x - \Delta)$, where $r_0(x) = f_0(x_0^*) - f_0(x)$. Again assuming $\Delta$ is sufficiently small (i.e., less than $\rho_0$), we can apply the general Taylor expansion according to (11) to obtain

$$|r_+(x) - r_0(x)| \le \Delta|r_0'(x)| + c_{2,\max}\Delta^2, \tag{71}$$

$$|r_-(x) - r_0(x)| \le \Delta|r_0'(x)| + c_{2,\max}\Delta^2, \tag{72}$$

where $c_{2,\max} = \max\{|\underline{c}_2'|, |\overline{c}_2'|\}$. Since $r_0'(x)$ is $c_2$-Lipschitz continuous (see (8)) and equals zero at $x_0^*$, we must have $|r_0'(x)| \le c_2|x - x_0^*|$. Hence, and using the triangle inequality along with (71)–(72), we have

$$|r_+(x) - r_-(x)| \le 2\Delta c_2|x - x_0^*| + 2c_{2,\max}\Delta^2, \tag{73}$$

which proves (29).

For the third part, we note that since $x_+^* = x_0^* - \Delta$ and $x_-^* = x_0^* + \Delta$, the conditions in (30) can be written as

$$r_+(x) \ge c''(x - x_+^*)^2, \qquad r_-(x) \ge c''(x - x_-^*)^2. \tag{74}$$

Using the locally quadratic behavior in (12), we deduce that (30) holds for all $x$ within distance $\rho_0$ of the respective function optimizer. On the other hand, if the distance from the optimizer is more than $\rho_0$, then a combination of (7) and (12) reveals that $r(x)$ is bounded away from zero. Since the quadratic terms in (30) are also bounded from above due to the fact that $x \in [0, 1]$, we conclude that (30) holds for sufficiently small $c''$.

## D. Proof of Theorem 2 (Lower Bound)

We continue from the reduction to binary hypothesis testing and auxiliary results given in Section 4. These results hold for an arbitrary given (deterministic) BO algorithm, which in general is simply a sequence of mappings that return the next point $x_t$ based on the previous samples $y_1, \dots, y_{t-1}$. Recall also that we implicitly condition on an arbitrary realization of $f_0$ satisfying the events in Assumptions 2 and 4, meaning that all expectations and probabilities are only with respect to the random index $V \in \{+, -\}$ and/or the noise. We proceed in two main steps.

**Bounding the mutual information.** To bound the mutual information term $I(V; \mathbf{x}, \mathbf{y})$ appearing in (31), we first apply the following tensorization bound for adaptive sampling, which is based on the chain rule for mutual information (e.g., see (Raginsky & Rakhlin, 2011)):[5]

$$I(V; \mathbf{x}, \mathbf{y}) \le \sum_{t=1}^T I(V; y_t | x_t). \tag{75}$$

It is well known that the conditional mutual information $I(V; y_t | x_t = x)$ is upper bounded by the maximum KL divergence $\max_{v,v'} D(P_{Y|V,X}(\cdot \,|\, v, x) \| P_{Y|V,X}(\cdot \,|\, v', x))$ between the resulting conditional output distributions $P_{Y|V,X}$ (e.g., see Eq. (31) of (Raginsky & Rakhlin, 2011)). In our setting, there are only two values of $v$, and since we are considering Gaussian noise, their conditional output distributions are $N(r_+(x), \sigma^2)$ and $N(r_-(x), \sigma^2)$. Using the standard property that the KL divergence between the $N(\mu_1, \sigma^2)$ and $N(\mu_2, \sigma^2)$ density functions is $\frac{(\mu_2 - \mu_1)^2}{2\sigma^2}$, we deduce that

$$I(V; y_t | x_t = x) \le \frac{(r_+(x) - r_-(x))^2}{2\sigma^2}. \tag{76}$$

Substituting property (29) in Lemma 3 gives

$$\begin{aligned} I(V; y_t | x_t = x) &\le \frac{(c')^2}{2\sigma^2}\big(\Delta|x - x_0^*| + \Delta^2\big)^2 && (77)\\ &\le \frac{3(c')^2}{2\sigma^2}\big(\Delta^2|x - x_0^*|^2 + \Delta^4\big), && (78) \end{aligned}$$

where (78) follows since $(a+b)^2 \le 3(a^2+b^2)$. Averaging over $x_t$, we obtain $I(X; y_t|x_t) \le \frac{3(c')^2}{2\sigma^2}\big(\Delta^2\mathbb{E}\big[|x_t - x_0^*|^2\big] + \Delta^4\big)$, and substitution into (75) gives

$$I(V; \mathbf{x}, \mathbf{y}) \le \frac{3(c')^2}{2\sigma^2}\bigg(\Delta^2\mathbb{E}\bigg[\sum_{t=1}^T |x_t - x_0^*|^2\bigg] + T\Delta^4\bigg). \tag{79}$$

[5]This form of the bound is not stated explicitly in (Raginsky & Rakhlin, 2011). However, Eq. (27) of (Raginsky & Rakhlin, 2011) states that $I(V; \mathbf{x}, \mathbf{y}) \le \sum_{t=1}^T I(V; y_t | x_1^t, y_1^{t-1})$, where $x_1^t = (x_1, \dots, x_t)$ and similarly for $y_1^{t-1}$. Letting $H(\cdot)$ denote the (differential) entropy function (Cover & Thomas, 2001), we obtain (75) by writing $I(V; y_t|x_1^t, y_1^{t-1}) = H(y_t|x_1^t, y_1^{t-1}) - H(y_t|x_1^t, y_1^{t-1}, V)$, applying $H(y_t|x_1^t, y_1^{t-1}) \le H(y_t|x_t)$ since conditioning reduces entropy, and applying $H(y_t|x_1^t, y_1^{t-1}, V) = H(y_t|x_t, V)$ since in our setting $y_t$ depends on $(x_1^t, y_1^{t-1}, V)$ only through $(x_t, V)$.

**Bounding the regret.** We consider the cases $\mathbb{E}[R_T] \geq c''T\Delta^2$ and $\mathbb{E}[R_T] < c''T\Delta^2$ separately, where $c''$ is defined in Lemma 3. In the former case, we immediately have a lower bound on the average cumulative regret, whereas in the latter case, the following lemma is useful.

**Lemma 5.** *If $\mathbb{E}[R_T] < c''T\Delta^2$ with $c''$ defined in Lemma 3, then $\mathbb{E}\big[\sum_{t=1}^T |x_t - x_0^*|^2\big] < 4T\Delta^2$.*

*Proof.* Since $V$ is equiprobable on $\{+,-\}$, we have

$$\mathbb{E}[R_T] = \sum_{v\in\{+,-\}} \mathbb{E}\bigg[\sum_{t=1}^T r_v(x_t)\,\bigg|\,V=v\bigg] \tag{80}$$

$$\geq c'' \sum_{v\in\{+,-\}} \mathbb{E}\bigg[\sum_{t=1}^T ((x_t - x_0^*) + v\Delta)^2\,\bigg|\,V=v\bigg] \tag{81}$$

$$\geq c''\mathbb{E}\bigg[\sum_{t=1}^T (x_t - x_0^*)^2 - 2\sum_{t=1}^T |x_t - x_0^*|\Delta + T\Delta^2\bigg], \tag{82}$$

where (81) follows from (30) in Lemma 3, and (82) follows by expanding the square and lower bounding the cross-term by its negative absolute value.

Substituting the assumption $\mathbb{E}[R_T] < c''T\Delta^2$ into (82), and canceling the term $c''T\Delta^2$ appearing on both sides, we obtain

$$\mathbb{E}\bigg[\sum_{t=1}^T (x_t - x_0^*)^2\bigg] < 2\Delta\mathbb{E}\bigg[\sum_{t=1}^T |x_t - x_0^*|\bigg] \tag{83}$$

$$\leq 2\Delta\sqrt{T}\mathbb{E}\Bigg[\sqrt{\sum_{t=1}^T (x_t - x_0^*)^2}\Bigg] \tag{84}$$

$$\leq 2\Delta\sqrt{T}\sqrt{\mathbb{E}\bigg[\sum_{t=1}^T (x_t - x_0^*)^2\bigg]}, \tag{85}$$

where (84) follows from the Cauchy-Schwartz inequality, and (85) follows from Jensen's inequality. Solving for $\mathbb{E}\big[\sum_{t=1}^T (x_t - x_0^*)^2\big]$ yields the desired claim. $\square$

In the case $\mathbb{E}[R_T] < c''T\Delta^2$, we claim that under the choice $\Delta = \big(\frac{\sigma^2}{\widetilde{C}T}\big)^{1/4}$ with a sufficiently large constant $\widetilde{C}$, it holds that $\mathbb{E}[R_T] \geq \widetilde{c}\sigma\sqrt{T}$ for some constant $\widetilde{c}$. Once this is established, combining the two cases with the choice of $\Delta$ gives

$$\mathbb{E}[R_T] \geq \min\bigg\{\frac{c''}{\sqrt{\widetilde{C}}}, \widetilde{c}\bigg\}\sigma\sqrt{T}, \tag{86}$$

which yields Theorem 2. We also note that by the assumption $\sigma^2 \leq c_\sigma T^{1-\zeta}$ in Theorem 2, we have for sufficiently large $T$ that $\Delta$ is indeed arbitrarily small under the above choice, as was assumed throughout the proof.[6]

It only remains to establish the claim stated above (86) when $\mathbb{E}[R_T] < c''\sigma\sqrt{T}$. By Lemma 5, we have $\mathbb{E}\big[\sum_{t=1}^T |x_t - x_0^*|^2\big] < 4T\Delta^2$, and substitution into (79) gives

$$I(V;\mathbf{x},\mathbf{y}) \leq \frac{15(c')^2}{2\sigma^2}T\Delta^4. \tag{87}$$

Since $\Delta^4 = \frac{\sigma^2}{\widetilde{C}T}$, we deduce that $I(V;\mathbf{x},\mathbf{y}) \leq \frac{\log 2}{4}$ (say) when $\widetilde{C}$ is sufficiently large. As a result, (31) gives $\mathbb{E}[R_T] \geq \underline{c}_2 T\Delta^2 H_2^{-1}\big(\frac{3\log 2}{4}\big)$ (note that $H_2^{-1}$ is an increasing function). Since $\Delta^2 = \sigma\sqrt{\frac{1}{\widetilde{C}T}}$, we deduce that $\mathbb{E}[R_T] \geq \widetilde{c}\cdot\sigma\sqrt{T}$, where $\widetilde{c} = \underline{c}_2\sqrt{\frac{1}{\widetilde{C}}}H_2^{-1}\big(\frac{3\log 2}{4}\big)$. This establishes the desired result.

[6] It is safe to assume that $T$ is sufficiently large, since the smaller values of $T$ can be handled by decreasing $C'$ in the theorem statement.